

University of Southampton Research Repository

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.

Thesis: Author (Year of Submission) "Full thesis title", University of Southampton, name of the University Faculty or School or Department, PhD Thesis, pagination.

Data: Author (Year) Title. URI [dataset]



University of
Southampton

University of Southampton

Faculty of Engineering and Physical Sciences

School of Electronics and Computer Science

Capacitor Design for High-Speed High-Precision Charge Redistribution SAR ADCs

DOI [10.5258/SOTON/PG/T202](https://doi.org/10.5258/SOTON/PG/T202)

by

William Christopher Webb

ORCID ID [0000-0003-3568-6279](https://orcid.org/0000-0003-3568-6279)

Thesis for the degree of Doctor of Philosophy

May 2026

University of Southampton

Abstract

Faculty of Engineering and Physical Sciences

School of Electronics and Computer Science

Doctor of Philosophy

Capacitor Design for High-Speed High-Precision Charge Redistribution SAR ADCs

by

William Christopher Webb

Analog-to-Digital Converters (ADC) are a key component in the majority of instrumentation and imaging applications, and are often the limiting factor in their system performance. In the high-precision (>16b), high-speed (>100MHz) space, successive approximation register (SAR) based ADCs are favoured. These ADCs usually contain a capacitive DAC (CDAC) whose output is matched to the input signal through charge redistribution. The CDAC is built in the back-end-of-line (BEOL) layers of a typical CMOS process with the Cu metal tracks and low-k dielectric insulation forming lateral coupling capacitors. The resultant precision and performance of the CDAC SAR ADC is frequently limited by the errors developed during fabrication of the nominally identical capacitors in the CDAC.

Improving or correcting for the matching of the capacitors is an expensive, difficult, and time-consuming procedure. Improvement in the variation of capacitors is hard to achieve as both systematic and stochastic errors are introduced within the CMOS fabrication process. These processes are optimised for the performance of the underlying transistors and cannot be adjusted to benefit a BEOL capacitor array. Hence, typically, each capacitor on each die is measured after fabrication, followed by additional expensive calibration, supporting circuitry, and the implementation of power and time-consuming linearisation techniques to compensate for such variations.

This thesis investigates in detail the influence of the geometric design and circuit layout of the capacitor array on its matching properties in a 40nm CMOS process, with the aim of achieving improved matching performance of the DAC to minimise the requirement for additional supporting circuitry or calibration routines. This is achieved through the use of dedicated test chips designed and optimised for the measurement of the capacitors in large arrays, as well as development of an ultra-precise methodology for doing so. This has resulted in a significantly enhanced understanding of both the systematic and correlated errors in the process as well as the random error due to statistical variation. From this work three main conclusions are drawn.

Firstly, it is shown that variation in capacitive value scales inversely proportional to the square root of the length of the capacitor fingers and number of metal layers as would have been expected from a typical statistical (Pelgrom) model. However, it is also shown that the variation due to scaling of finger width and spacing, is primarily determined by the critical dimension of the process, and that matching is significantly improved through relaxing these parameters slightly. Together, these observations allow an optimised design which balances capacitor matching against area use.

Secondly, the measurement and analysis of the capacitor matching in three different 40 nm CMOS designs reveals that despite all care and attention in the layout design, a large part of the variation in capacitor values is derived from the most subtle limitations of process technology leading to unpredictable but systematic errors across the die which are repeated on all other die. This means that by careful separation of the matching errors into systematic index-to-index variations and random die-to-die variations, it is possible to compensate for the systematic errors by only measuring a very small subset of all available dies. This approach will lead to a significant time and cost saving in the test and calibration phase of each die.

Lastly, the research reveals that, due to the presence of correlated error sources, considerable matching improvements may be attained through proper understanding of the relationships between the key capacitor geometry parameters, and the aggressive use of layout folding schemes. Additional benefit can also be achieved through use of the correct amount, and type, of dummy structures on arrays of all sizes.

The developed theory and findings in the earlier sections are subsequently applied to a simulated model of a single-stage 18-bit SAR ADC, demonstrating that the capacitor mismatch error is now surpassed by noise as the dominant source of error in the system, and that the effectiveness of other linearisation techniques is also further enhanced by the optimised capacitor design.

Table of Contents

Abstract	2
Table of Contents	4
Research Thesis: Declaration of Authorship.....	7
Acknowledgements.....	8
Chapter 1 Introduction	9
1.1 Motivation	9
1.2 Research Goals, Scope, and Contributions	9
Chapter 2 An Overview of Analog-to-Digital Converters	12
2.1 Introduction	12
2.2 Target Applications.....	12
2.3 High-Speed Analog-to-Digital Converters	13
2.4 ADC Architectures and Figures of Merit	13
2.5 Successive Approximation Register (SAR) ADC Architecture	17
2.6 Linearisation Techniques for SAR ADCs	20
2.7 Conclusion.....	26
Chapter 3 Design and Fabrication of a Test Chip for Capacitor Measurement in a 40nm Process	27
3.1 Introduction	27
3.2 Process Choice	27
3.3 Introduction to the 40nm CMOS Process	29
3.4 Overview of Key 40nm Back-End-of-Line Processes	31
3.5 Lateral Fringing Metal-Oxide-Metal (MOM) Capacitor Layout	34
3.6 Design of an Ultra-Precise Capacitor Measurement Architecture for Large Arrays	44
3.7 Setup and Evaluation of the Capacitor Measurement Die	55
3.8 Conclusion.....	61

Chapter 4 Measurement and Analysis of Single Finger Lateral Coupling	
Capacitors	62
4.1 Introduction	62
4.2 Capacitor Weighting Analysis.....	62
4.3 Conclusion.....	71
Chapter 5 Detailed Analysis of Capacitor Errors and Trends	72
5.1 Introduction	72
5.2 Derivation of Matching Statistics from Array Measurements.....	72
5.3 Verification of the Gaussian Statistical Model for Matching Errors	79
5.4 Analysis of Observed Systematic Errors	81
5.5 Trends in Matching Error with Capacitor Geometry Variations	92
5.6 Conclusion.....	108
Chapter 6 Capacitor Errors and Trends in Large Arrays	109
6.1 Introduction	109
6.2 Matching Within Differential DAC Arrays	109
6.3 Impact of Dummy Structures on Capacitor Matching	110
6.4 Impact of Array Shielding on Capacitor Matching	114
6.5 Localised Errors in Large BEOL Capacitor Arrays	115
6.6 Sources of Systematic Errors	124
6.7 Conclusion.....	126
Chapter 7 Simulated Performance Impacts of Improved Matching in a	
Behavioural 18-bit SAR ADC	127
7.1 Introduction	127
7.2 Behavioural ADC Architecture.....	127
7.3 Simulation Methodology	130
7.4 Incorporation of Measured Arrays for Simulation	131
7.5 Impact of Capacitor Interdigitation and Folding in Layout	132
7.6 Simulated ADC Performance	136

Table of Contents

7.7 Conclusion.....	139
Chapter 8 Conclusion and Outlook	141
Appendix A Tables of Capacitor Variations	144
References	147

Research Thesis: Declaration of Authorship

Print Name: William Christopher Webb

Title of Thesis: Capacitor Design for High-Speed High-Precision Charge Redistribution SAR ADCs

I declare that this thesis and the work presented in it are my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:
William Webb, Rares Bodnar, Kees De Groot, "Matching Improvements in Capacitor Arrays for Precision Charge Redistribution SAR ADCs.", International Conference on Electronic Circuits and Systems (ICECS), Nancy, France, 2024, doi: 10.1109/ICECS61496.2024.10848714.
William Webb, Rares Bodnar, Kees De Groot, "Optimisation of BEOL Capacitor Matching in Charge Redistribution DACs for Precision and High-Speed ADCs", Transactions on Circuits and Systems I: Regular Papers (TCAS-I), Early Access 2026, doi:10.1109/TCAS.2026.3664777

Signature: Date: 2026-05-25.....

Acknowledgements

I would like to express my deepest gratitude to Dr. Rares Bodnar for his exceptional guidance and mentorship. His deep knowledge of the subject, patience, passion, and his continual support over the past four years have been instrumental to my growth and to overcoming interdisciplinary challenges at every stage of this project.

Thank you also to Prof. Kees De Groot for his supervision throughout this project. His comments have kept this work grounded and purposeful towards its result.

I have to thank my colleagues at Analog Devices all for their unique insights and expertise, without whom substantial roadblocks that developed through this investigation I believe would have taken even longer to overcome. Their support on the development of the test chips across the 3 tape-outs for this work has been invaluable. The collaboration on the work of this project with every one of them has been a pleasure: Henry Kennedy, Luke Smithers, Peter Hurrell, Mark Vickery, Peter Strong, Ahmet Suna, Junming Zeng, Eric Thompson, Archana Patil.

I would also like to acknowledge the comments and guidance received from Prof. Bill Redman-White, which helped to cement the usefulness and context for a lot of this works' discussion. His extensive academic and industrial experience has helped shape the presentation of my work and supporting papers significantly.

Lastly, but by no means least, I would like to thank my close family and friends, particularly my parents, and my partner Oana. They have always inspired and pushed me to pursue the things I love. Their continuous support, encouragement, and love are what have made this achievement possible.

Chapter 1 Introduction

“We demand rigidly defined areas of doubt and uncertainty” – Vroomfondel

1.1 Motivation

A significant limitation in the development of modern instrumentation equipment is the precision and linearity of their Analog-to-Digital Converters (ADCs). These instruments frequently demand high sample rates of their ADCs without compromising the signal-to-noise (SNR) ratio or linearity (INL).

At low to medium throughput, Successive-Approximation (SAR) ADCs are generally preferred due to their instantaneous sample accuracy, bandwidth, and overall power efficiency inherent to their architecture. Usually, the SAR ADC is found within more complex hybrid architectures combining pipelining, interleaving, or noise-shaping techniques to achieve the highest throughput converters at a given precision. This has led to SAR ADCs attaining up to 24-bit resolution; however, the matching of their charge redistribution DACs has not followed suit. This is a fundamental limitation of the SAR ADC architecture that has been commonly solved through development of complex calibration and Dynamic-Element-Matching (DEM) algorithms, which are costly, and do not address the core of the problem: the matching of the capacitors in the capacitive DAC (CDAC). Therefore, it is essential to address the uncharacterised matching properties and dependencies of the common capacitor structures for SAR ADCs, such that continued improvements in SNR and linearity may be achieved, without the increasing cost of calibration and other correction techniques.

1.2 Research Goals, Scope, and Contributions

This research project seeks answers to answer the questions: What are the dominant contributing factors to, and the nature of, the matching errors found in large arrays of lateral coupling capacitors; and are there techniques and methods possible in the capacitor design, within the boundaries of the process, that may allow achieving of higher matching capacitors. While this work is executed within the field of high-speed, high-precision SAR ADCs, as will be seen there are undoubtedly many areas that benefit from this investigation. Execution within this space is performed primarily due to the prominence of the issue within the SAR ADC space. The work of this thesis aims to:

- Understand the likely sources of matching errors in BEOL lateral coupling capacitors

Chapter 1

- Verify the presence of systematic errors across large capacitor arrays
- Identify methods or geometries in layout that can be used to reduce the formation of errors in large arrays of capacitors
- Where avenues for improvement are found, demonstrate their effectiveness with comparison to the foundry capacitor cells in the context of a SAR ADC CDAC.

The scope of this project's experimental investigation therefore covers the 40nm process due to availability of the PDK and its applicability to high-speed, high-precision SAR ADCs. This is however only one of many suitable processes for this investigation. The project considers primarily what can be done within the design stages of the capacitors or capacitor arrays which lead, through understanding, or use of another technique, to better matching capacitors, and thereby better linearity and performance of the attached circuit. The precise implementation of calibration algorithms, or other linearisation techniques are considered out of the scope of this project, as are the other supporting components of the circuits which may lead to other detriments to performance. For most commercial processes, the precise process information is protected, and even still it is extremely limited. As a result, this thesis focuses predominantly on the relationships and relative improvements possible, as opposed to the specific figures. Where process is discussed, it is done using only the publicly available information, which may limit discussion.

This thesis is divided into several chapters targeted to tackle the above question and goals. Chapter 2 provides an overview of the existing high-speed high-precision ADC space, justifying the desire for high linearity, and the current research aimed in achieving it. In Chapter 3, test structures for a series of test chips are designed, and an ultraprecise measurement strategy is developed to characterise the matching dependencies of the capacitors. Chapter 4 explores the relationship between the capacitor geometry parameters and the resulting capacitor value to ascertain the likely mismatch contributors. This is used to inform the work of Chapter 5, where the statistical analysis is applied to the capacitor weightings and the nature of the capacitor matching is analysed, and relationships extracted. Chapter 6 explores the larger scale matching effects, comparing the observations with traditional techniques, and discusses the likely causes of the novel relationships that are observed prior. The thesis then demonstrates the resulting matching performance improvement through simulation of a novel, fractional reference schemed, calibration-free, 16-bit linear SAR ADC, in Chapter 7. Significant improvements over the foundry capacitor structure and performance are demonstrated, and the dependencies of their performance are characterised. The main contributions from this project are summarised below.

Chapter 1

- Design and evaluation of bespoke capacitor arrays for characterising the matching dependencies of 40nm BEOL lateral MOM capacitors.
- Development of a high-precision on-chip capacitor measurement technique, and theory allowing separation of the systematic and randomly sourced errors.
- Characterisation of capacitor matching error dependencies on the main capacitor geometry parameters, and identification of significant correlated matching effects.
- Characterisation of the requirement for array-level dummy structures, and their impact on the capacitor matching errors.
- Characterisation of the possible improvements through aggressive use of layout folding techniques in large capacitor arrays.
- Demonstration of the possible performance improvement for a calibration-free 18-bit SAR ADC performing with 16-bit linearity using a behavioural 18-bit ADC and optimised capacitor structures.

Chapter 2 An Overview of Analog-to-Digital Converters

2.1 Introduction

The scope of this project covers predominantly the use of the large arrays of laterally coupling capacitors within charge-redistribution SAR ADCs. This chapter provides an overview of the typical uses, figures of merit, and configurations of ADC under which is work is considered, as well as common linearisation techniques, for comparison of what this work seeks to improve upon, or ultimately remove the need for entirely. The issues with the use of calibration are also discussed, and the impact of highly linear performance is noted.

2.2 Target Applications

The issues presented by badly matched capacitors are very widespread across a great range of applications. For example: in differential structures, in the same way as device mismatch leads to circuit offset and bad common mode rejection, mismatches in the capacitance of the nodes, can also lead to the same sensitivity. Most high-precision circuits are trimmed during production or during background calibration routines (as will be discussed), and the capacitor arrays that are used work best when they are monotonic. Where control loops are introduced, non-monotonicity leads to unstable behaviour and is considered unacceptable. Achieving of high dynamic range in any application will frequently require matching of high ratios of devices, be they transistors, resistors, or capacitors. Such circuits can be found in offset trimming within voltage-controlled oscillators (VCOs). When matching of passive devices is concerned, capacitors are generally chosen due to their ability to be used in switched capacitor circuits which are frequently lower power, as well as their standard integration into the process and lower area costs compared with resistors.

The most prominent application for large arrays of capacitors which must have high levels of matching due to the high dynamic range that is required, is in the application of charge redistribution analog-to-digital converters (ADCs). This is the application that is focused on within this work, as an exemplar of both the impact on performance that the capacitor matching can have, but also as a guide to the trade-offs that a designer might have to make and consider when consulting different options for improving the structure's performance. Across each of the applications, there are a broad range of considerations that become applicable. Some applications such as VCOs require good control of the absolute capacitor value, whereas this is not of great importance with analog-to-digital converters. The sizes of capacitors can vary

considerably too. Power trimming applications might use arrays in the μF range, whereas the chosen application will consider itself with arrays of picofarads of capacitance. Due to these wide ranges, it is important that this study constrains to a single area to focus the investigation. It will however be found in later sections, that much of the results of this study, as hoped, will be highly transferrable to other applications and processes by careful managing of the trade-offs.

2.3 High-Speed Analog-to-Digital Converters

Analog-to-Digital Converters (ADCs) provide the interface between advanced digital processing techniques, and the high-precision sensors in modern instrumentation. Fundamentally they take an input, usually a voltage, of a real analog signal from a sensor and convert it to a digital code for a processor or system-on-chip to interpret. The requirements and demands of these ADCs can vary dramatically between applications with high-speed communications systems requiring sample rates of more than 1 billion samples per second (1 GSPS) with resolutions as low as 5-bit, and bespoke instrumentation requiring sometimes less than 1kHz bandwidth at up to 30-bit resolution. In both applications, the ADC is frequently the limiting factor in the system performance. The focus of this study revolves around the typical performance expected of an ADC in an imaging type instrumentation scenario, where both high precision (>16 -bit) and high speed ($>100\text{MHz}$) are common specifications.

In imaging applications, an extreme amount of post processing is performed on the digitised signal in order to infer and reconstruct the source being imaged. This is often the nature of CT and MRI type scanner systems. As a result, the noise and other distortion elements in the output of the ADC can be exacerbated quite easily by the post processing algorithms. The complexity of this processing is often large enough in its own right that system integrators are not able to implement individual compensation for each ADC (of which there may easily be hundreds) in addition to image-resolving post-processing algorithms, and so it is up to the ADC designer to ensure that the off-the-shelf performance of the ADC is as high as possible for the customer. This means integration of such compensations to the ADC itself.

2.4 ADC Architectures and Figures of Merit

2.4.1 Figures of Merit

Figure 2.1 illustrates the extraction of several key figures of merit that may be used to quantify the dynamic performance of a given ADC. The plot shows the extraction of these metrics from a fast-Fourier transform (FFT) of a single fixed-frequency input at f_{in} . This is the preferred method of evaluation of AC characteristics.

The simplest metrics relating to ADC performance, are the signal-to-noise ratio (SNR), and the dynamic range (DR). The SNR is as its name suggests, the ratio of the applied input signal, to the noise floor of the ADC. Note that this excludes any harmonic spurs. The dynamic range is the ratio of the full-scale input level of the ADC to the noise floor. Full-scale refers to the maximum input range of the ADC.

Imperfections in the ADC operation, arising from various sources such as input signal dependency, mismatch, internal clocks and crosstalk, generally manifest themselves as spurs in the frequency domain. Only HD2 and HD3 are shown in Figure 2.1 but subsequent spurs may exist. These are the result of tones and distortion in the ADC response. In applications where the distortion of the ADC is of concern, the signal-to-noise and distortion ratio (SNDR) of a converter better captures the real-world performance than the SNR. Similarly, the spurious-free dynamic-range (SFDR), which describes the ratio between the ADC full-scale, and the largest spur present in the frequency response.

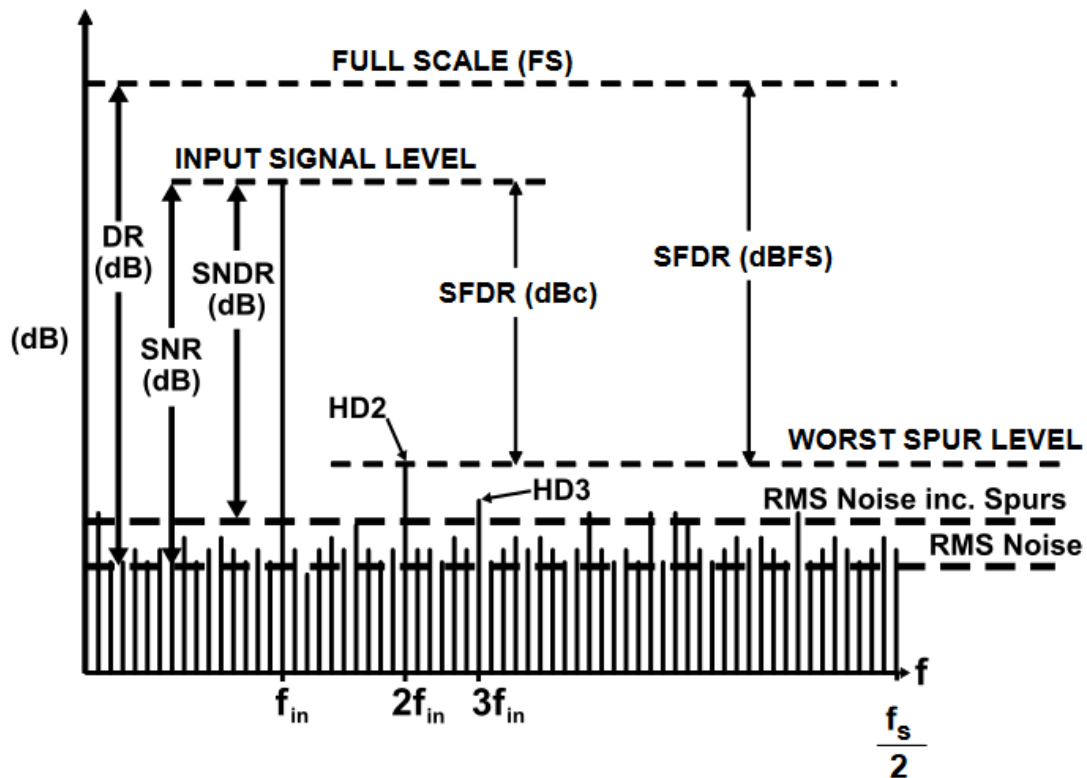


Figure 2.1. Illustration of the commonly used figures of merit for an ADC with reference to the ADC output FFT. Adapted from Ref [1].

Where DC characteristics are of concern, and frequently in applications reliant on low-distortion, the linearity of the converter is represented predominantly by the integrated-non-linearity (INL), and differential-non-linearity (DNL). These are illustrated in Figure 2.2. For an ADC to be monotonic (a requirement in the majority of applications, but particularly in control or regulation loops) means for the DNL to be less than one least significant bit (LSB). The linearity of an ADC directly impacts its precision. To consider an 18-bit ADC linear, it would be expected

that its INL would be less than 1LSB at 18-bit level. These are the most important metrics which will be used for further discussion of ADC performance in the subsequent sections, where the product of capacitor mismatch will most significantly impact the INL, and DNL of a converter.

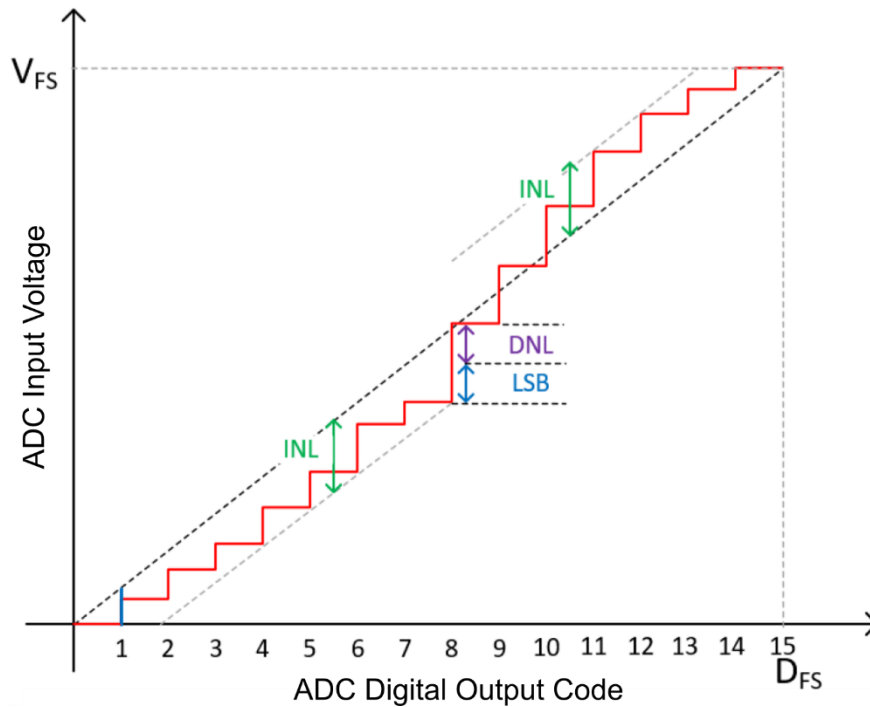


Figure 2.2. Visualisation of ADC INL and DNL non-linearities between the input voltage and the post-conversion output code of the ADC.

2.4.2 ADC Architectures

Figure 2.3 demonstrates the state-of-the-art performance of various converter architectures from the past decade. To understand the motivation behind the use of the SAR ADC as the basis for many hybrid converter architectures it is instructive to have a brief look at the other converter architectures.

Flash ADCs use a chain of resistors or capacitors to divide up the input voltage and a series of comparators tapped into said chain to provide a near-instantaneous conversion result in a single clock cycle. They excel in providing very high-speed conversion results with low-latency and are commonly utilised in wireless and wireline applications where sampling frequency and bandwidth requirements are the highest. Whilst their simplicity allows very high sampling speeds (>1 GSPS) the flash architecture does not scale linearly with increased resolution, requiring a near doubling of area and power, with each additional bit of resolution. This places a general limit of 8-bits resolution on flash ADCs, although integration of low-resolution flash converters within other ADC architectures is commonplace.

Sigma-Delta (SD) based converters harness error shaping techniques to extract very high precision from a lower precision quantifier. As a result, they are continuous in nature and hence

require aggressive oversampling to produce adequate conversion results. As a result, sigma-delta converters are the highest resolution ADCs on the market, being able to attain higher resolution performance through oversampling and noise shaping (NS) techniques. Their suitability for simultaneous sampling, or multiplexed applications is however compromised by the required large oversampling ratio (OSR), and their effective conversion latency is unsuitable for many instruments, particularly imaging, where tight sampling timing constraints exist. Sigma-delta ADCs are often operating with less than 5MHz bandwidth and dominate within the audio sector.

The substantially reduced bandwidth compared with other architectures and the high oversampling ratio makes sigma-delta ADCs unsuitable for multiplexed applications commonly seen in instrumentation, although sometimes this is still acceptable. A true Nyquist ADC by comparison has an oversampling ratio of 1 making other architectures more suitable, although oversampling techniques are frequently applied to most ADCs to improve SNR. Due to their aggressive use of oversampling, sigma-delta ADCs are able to achieve particularly high figure of merit (FoM) in the <1MHz region.

Pipeline ADCs achieve higher precision by breaking the conversion into multiple steps. They chain a series of quantisers together, feeding the quantisation error of the previous stage typically through an amplifier to the next stage for further quantisation. This allows for fast sampling speeds as the maximum sampling speed is determined by the throughput of the slowest quantiser in the pipeline, as all quantisers are usually able to operate independently. This architecture can increase the complexity of the ADC architecture and brings with it another set of challenges for designing low-noise, high-speed amplifiers between each stage. Any of the discussed architectures may be used for pipeline ADC quantisers, however pipelined SAR ADCs are by far the most common in the instrumentation space and generally have the highest Nyquist performance.

A comprehensive range of competitive ADC converters recently published in ISSCC and VLSI conferences are shown in Figure 2.3, highlighting the trade-off between precision and speed. In the range of 100 kHz to 1MHz the pipelined SAR ADC is generally the superior ADC. As a result, the pipelined, or 2-stage, SAR converter is the typical architecture of choice for instrumentation and imaging applications, where a high-precision (16-24 bit), high speed (10-500 MHz) ADC is required.

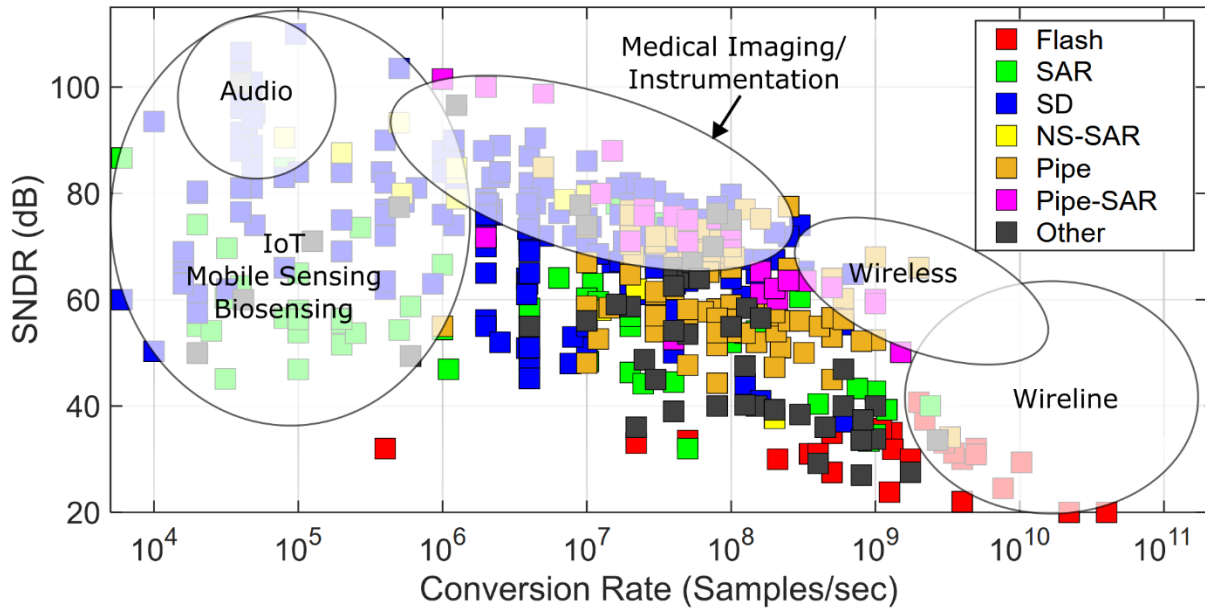


Figure 2.3. Example converter performance by architecture, overlaid with each application's SNDR demands. Assembled from [2]. Data collated in [3]. (Time-interleaved converters are omitted from this plot)

2.5 Successive Approximation Register (SAR) ADC Architecture

A diagram of a typical SAR ADC is shown in Figure 2.4, and illustrates the fundamental components on which its performance relies. In its simplest form, the ADC consists of 3 parts. A digital to analog converter (DAC), a comparator, and the successive approximation register (SAR) logic.

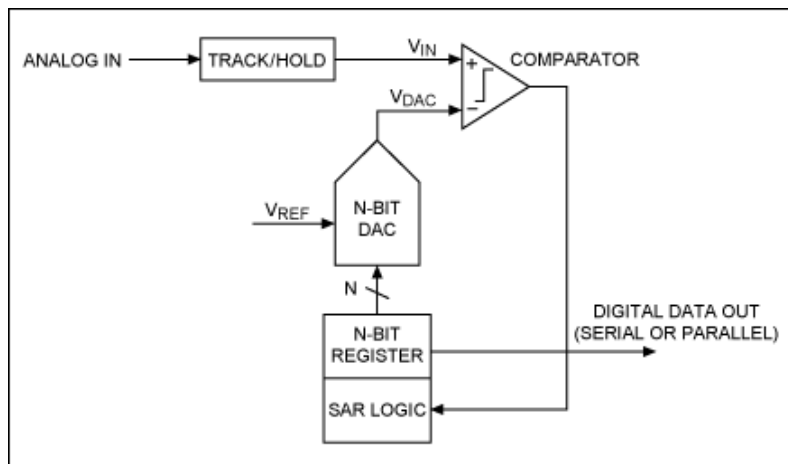


Figure 2.4. Simple architecture of an N-bit SAR ADC from [4], taking an analog voltage input and providing a digital representation of its output after a number of clock cycles.

A capacitive charge-redistribution DAC (CDAC) consists of an array of binary weighted capacitors with their 'top plates' all connected to the comparator, while their bottom plates are connected to switches to tie the node between two references and the ADC input voltage. This is shown in Figure 2.5. The comparator output feeds into the SAR logic which controls the

capacitor bottom plate switches based on whether the top plate voltage at the end of each comparison is above or below the reference. The SAR algorithm is an iterative process which functionally serves to equalise the top plate voltage to the comparator reference. This is achieved by trialling the binary weighted capacitors successively from largest to smallest as shown in Figure 2.6 to redistribute the charge across the CDAC and manipulate the top plate voltage. Given the final switch positions, and the capacitor weights, a digital value for the conversion result is constructed. Any difference between the designed value of the capacitor weighting, and the fabricated weighting, results in the possibility of incorrect decisions being made during the conversion, producing an incorrect final conversion results. This simple algorithm is highly efficient in its power usage and scales exceptionally well compared with other ADC topologies. Multiple of these SAR stages are often chained together to form a pipeline or 2-stage SAR ADC. Usually operating between 5 to 9-bit level per stage for multi-stage and pipelined SAR ADCs. The residue (final error on the top plate at the end of a conversion) would be amplified with a residue-amplifier and taken as the input to the proceeding stage. This is done generally required to enhance the dynamic range of the DAC beyond 16-bit, and to increase speed by operating each stage simultaneously.

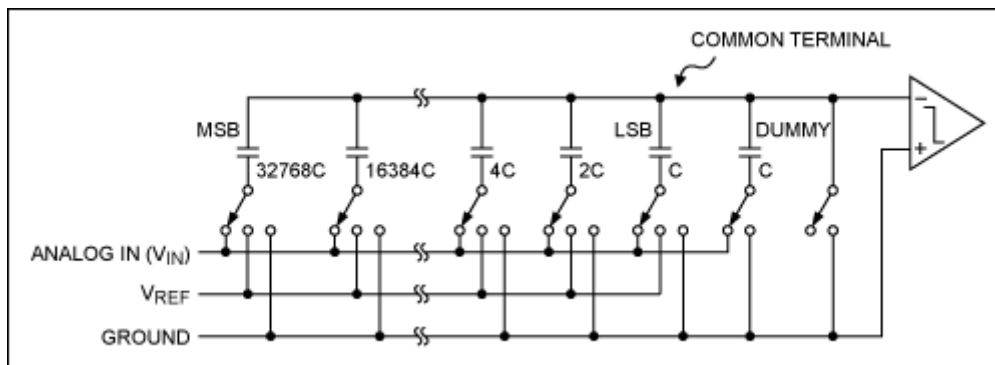


Figure 2.5. Architecture of a 16-bit charge-redistribution capacitive DAC [4]. The capacitor switches are manipulated successively to manipulate the top plate voltage to match the reference (ground here) via the redistribution of charge.

It should be noted that the DAC is not necessarily a charge redistribution capacitive DAC. Current or voltage scaling DACs are also in use. Current steering DAC alternatives are in use but are more commonly found in direct current ADCs or in low frequency discrete DACs where the switching nature of a CDAC renders it unsuitable for continuous time operation. The switching involved in a typical SAR ADC by comparison occurs quickly enough such that leakage does not degrade performance. An 'R-2R' style DAC could also be appropriate here, however the use of resistors as the critical elements introduces additional thermal noise to the system, and is unlikely to meet the noise requirements for a precision (>16-bit) SAR ADC DAC.




Figure 2.6. The capacitor array 'top plate' voltage during a simple 4-bit successive approximation conversion [4].

Whilst comparator offset is one error source contributing to the performance of the ADC, this is a singular quantity which is easily trimmed out, and is typically signal independent. If we consider now the contributions of the matching errors of the capacitors, with each capacitor having its own unique, ideally decorrelated, error that is only present when the capacitor is switched, then it is easy to understand the potentially signal dependent nature of the errors. Figure 2.7 illustrates this through simulation of an ideal 18-bit SAR ADC that was designed and is presented later in Chapter 6. The SNR degrades significantly when random errors are applied to its capacitors in a Monte-Carlo simulation [5]. For the capacitor values used in a precision application, typical matching errors are well in excess of 1000ppm, resulting in an immediate degradation of 20 dB SNR in this instance before any other error contributors are considered. This makes correction of the capacitor matching a significant topic of interest for those developing high-speed, precision SAR ADCs. It is the purpose of this thesis to design the capacitors in such a way to ensure that the capacitor error variation is limited ideally to 100ppm such that its effect on performance is small compared to the quantisation noise.

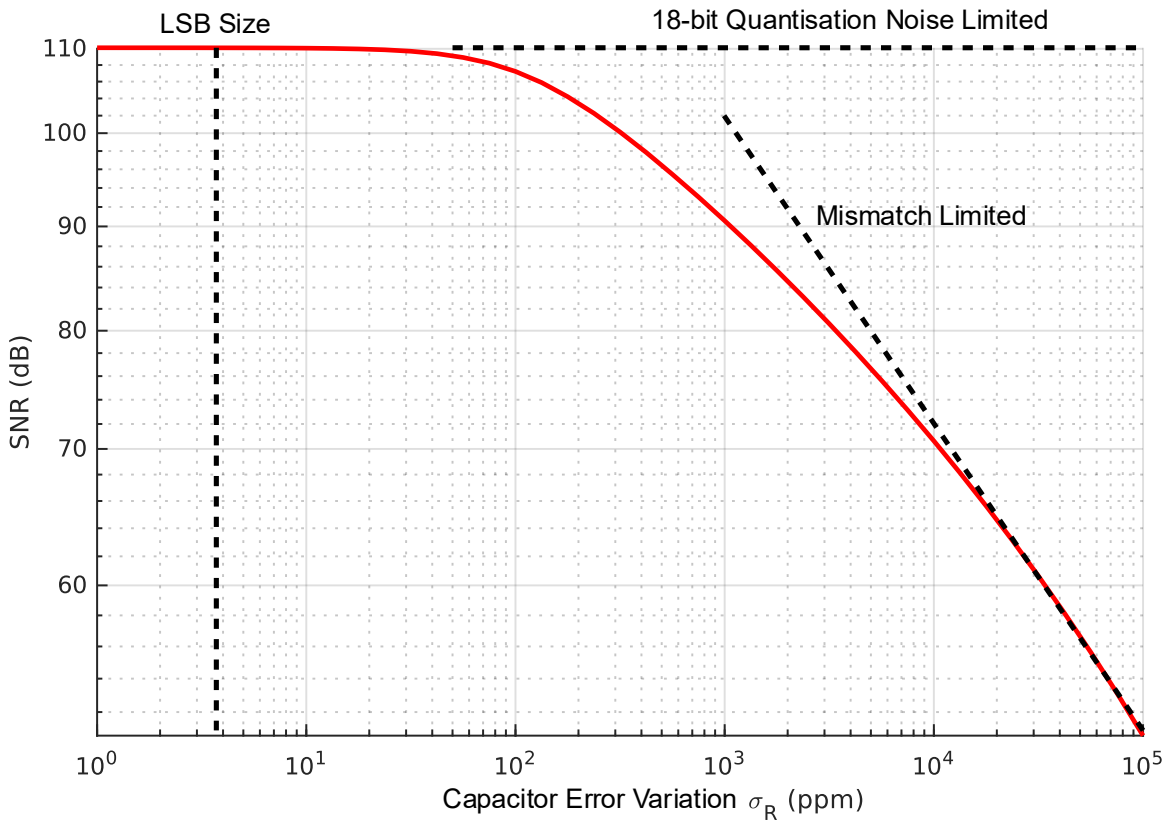


Figure 2.7. Simulated impact of capacitor matching error on the performance of an 18-bit SAR ADC which uses redundancy to mitigate the impact of capacitor mismatch. The design and optimisation of this ADC is the subject of Chapter 6 of this thesis.

2.6 Linearisation Techniques for SAR ADCs

2.6.1 DEM and Dither

Dynamic Element Matching (DEM), frequently referred to as shuffling, serves as one of the more fundamental techniques frequently used when tackling the presence of unwanted mismatch in capacitor DACs [6], [7], [8], [9], [10], [11]. It effectively provides the ability to trade-off the capacitor errors, and the subsequent spurs introduced in the frequency domain, for a wider band noise that recovers the converter's SFDR and SNDR, and enhances linearity. This is achieved by cycling of capacitor allocations in thermometrically weighted capacitor arrays each conversion as in Figure 2.8, such that the errors of each capacitor are averaged over multiple conversions. This assumes the errors are approximately gaussian and zero mean. The result is a DAC that is "*linear: on average*"[8], [11]. DEM finds application in a broad range of architectures, including SARs, and Sigma-Delta ADCs, as well as other applications, but the general rules and principles that ensure its effectiveness remain the same.

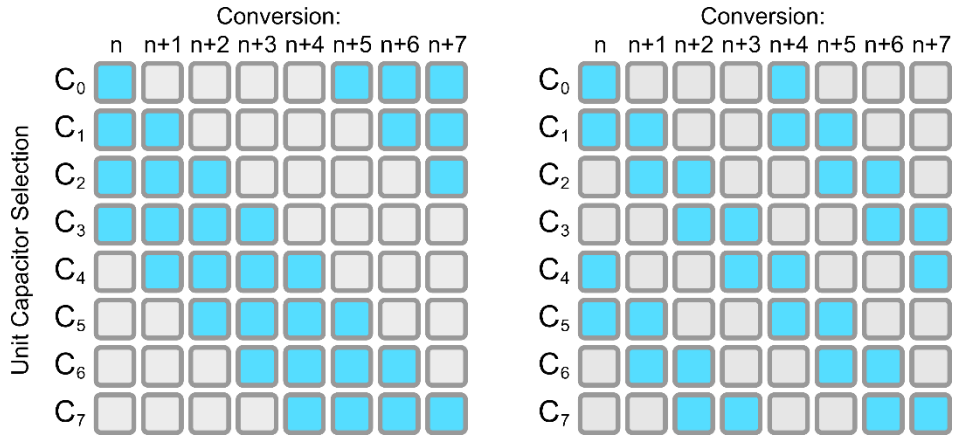


Figure 2.8. Capacitor allocation changes each conversion cycle through use of a barrel shuffler for DEM implementations. Shown here are two valid shuffling schemes for a constant DC input [8].

The real impact of a well-performing DEM scheme in an SD ADC is illustrated in Ref [9], where mismatch derived frequency spurs are notably suppressed, and the achievable SNR of numerous matching schemes is illustrated. In order to provide these SNDR improvements, oversampling is required. This is often a key factor against the use of DEM in particular instantaneous-sample type applications, as multiple samples must be taken and averaged to produce a DC, or at least Nyquist, linear result. As previously mentioned, whilst the SNDR of a converter is maintained through use of DEM, due to the spreading of the errors into the noise floor, a very slight deterioration in SNR is expected.

In the context of a SAR ADC, DEM may only be applied to unary sections of the DAC as it requires elements of equal weighting. Hence, it is usually seen limited to the 4 or 5 most significant bits (MSBs) of the ADC, while the remaining sections of the DAC are binary weighted. The main factor limiting usage to 4 or 5 bits in high-speed SAR ADCs (>100MHz) is the propagation delay through the shuffler which implements the DEM scheme, as the load and parasitics of the shuffler scale exponentially. This also removes the typically observed step in INL from the MSB transition of the DAC without compromising the dynamic range, simply by reduction of the MSB unit element size. Some DEM derived techniques are referred to as Data Weighted Averaging (DWA), but these techniques still rely on the same fundamental principles of DEM for their operation.

DEM is generally found combined with the other most impactful linearisation technique, Dither [10], [12], [13], [14], [15]. When combined, the ADC transfer function can be made impressively linear, at less than 0.1ppm INL error [10]. Dither further linearises the ADC transfer function through random modulation of the ADC input, and thereby the quantisation error (the residue). This is achieved through addition of a random, but known, signal to the ADC input, generally through added sampling capacitors, which is then converted in addition to the user's input to the ADC. This dither signal is then digitally removed post-conversion as in Figure 2.9. The result

is a quantisation error that is decorrelated from the capacitor errors of the DAC, effectively smoothing out the MSB transitions in the INL and DNL as in Figure 2.10, and resulting in a lower overall DNL due to the more gradual transition.

The use of dither also provides the opportunity for mismatch calibration through dither injection [14], [16], [17], [18]. This typically requires additional capacitors and usually introduces more capacitors which have their own errors that must be calibrated out, albeit easily achieved through similar means. The main downfall of dither is that it consumes the input range of the ADC, which must then be combatted with the use of further redundancy as the dither signal is summed with the input signal in operation, thereby requiring additional range for the bit trials to ensure a residue $<|1/2LSB|$. Whilst this in itself may not constitute a detriment to performance, it poses challenging design decisions as the DAC must be extended to account for the further redundant capacitors, increasing the occupied area, which may be significant if near-MSB sized redundant capacitors are required.

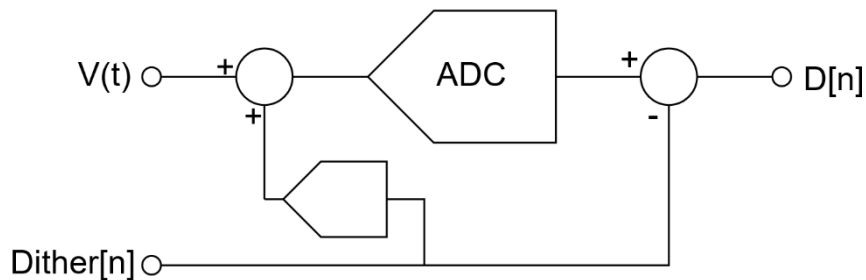


Figure 2.9. This diagram shows how dither is introduced on top of an existing ADC architecture by introduction of a random dither signal 'Dither[n]' and subsequent removal of the signal post-conversion.




Figure 2.10. Illustration of the operation of dither. The input voltage at which code transitions occur is artificially varied, resulting in linearisation of the average INL profile, and the removal of large steps in bit transitions.

Unlike calibration which directly compensates for the capacitor errors, DEM and dither can only distribute the frequency at which these errors are spread into the noise. In both cases the objective is to decorrelate the mismatch induced noise from the input signal. Signal dependent errors are not absorbed into the noise floor and create tones and harmonics in the ADC conversion result, whereas the random nature of both dither and DEM decorrelate these errors and subsequently increase the converter's SFDR, and linearity.

2.6.2 Redundancy

Redundancy introduces additional trials into the SAR algorithm that are repetitions of trials of the same weighting [19], [20], [21]. This is operated exactly the same as a normal bit-trial would be, resulting in a second voltage change on the top plate of the same magnitude when switched. If the capacitors are of ideal value, then the redundant capacitor bit trial result would be the opposite of the previous bit trial. However, if there are errors accumulated in the previous capacitor values resulting in an incorrect bit-trial result, then this allows some level of recovery by adding extra range to the DAC, and allowing the final residue to be returned to within ± 0.5 LSB. This additional input range also allows for absorption of other error sources, such as over-ranging of the preceding residue amplifier, or mismatch between the DAC and a supporting-flash ADC trip point.

Redundancy is also critical in the use of many techniques, particularly calibration routines, as it allows multiple paths for reaching a particular ADC output code, with different combinations of capacitors switched. This is the result of errors in the bit weightings cause a wrong bit-trial result. This is shown in Figure 2.11 where the presence of mismatch causes a false bit-trial result in the dashed-blue pattern.

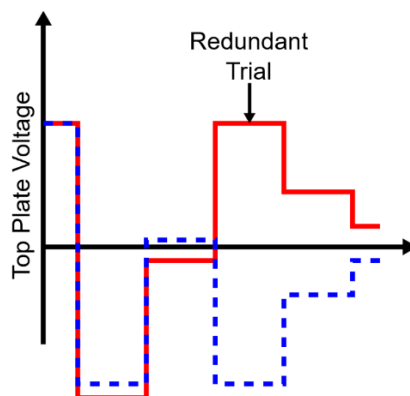


Figure 2.11. Demonstration of how redundancy can enable multiple paths to achieve the same ADC result in the presence of mismatch. The dashed blue line shows a conversion with ideal capacitors, and the solid red line shows a conversion with mismatch errors, resulting in a different decision whilst maintaining the residue within the prescribed range.

2.6.3 Calibration

Possibly the most prominent, and direct, method for addressing the inherent capacitor matching errors, is through digital calibration of the capacitor weightings. Such calibration routines have become increasingly popular due to their low area and power cost in geometries below $0.18\mu\text{m}$, and sort themselves into two types: foreground calibration, and background calibration.

Foreground calibration algorithms, sometimes known as factory calibration, requires the stopping of the normal conversion cycle in order to detect, and correct calibration errors, through use of an internal or external reference or known stimulus. These calibration routines are seen in the higher precision SAR ADCs [5], [10], [13], [21], as more control is afforded over the ADC input, and switching sequence. This allows more precise determination of any correction coefficients. This is of great benefit in pipeline ADCs utilising aggressive dither and shuffling techniques, as these additional capacitors must also undergo calibration in order to be effective. Ye [12] details a technique for generation of the calibration matrix that exploits the redundant capacitor usage in order to build the matrix for estimation of the capacitor weightings. This method demonstrates 4.4dB SNDR, and 12.5dB SFDR improvement without relying on a calibrated external source. However, this still requires a 'dead' time in operation of the ADC. To clarify, these routines in real production test scenarios may consume a whole number of seconds to capture enough data in order to fully calibrate the ADC and then generate correction coefficients to fuse to the part. This is due in part to the large amount of additional shuffle, dither, and redundant capacitors that are used in high-precision ADCs, which all also require calibration in order to achieve the best nominal performance of the part.

By comparison, background calibration algorithms run in parallel with the standard conversion operation. This enables them to continually monitor and correct for any temperature dependent mismatch errors, comparator offset or gain drift that may occur, as well as aging effects which are an important consideration within commercial designs [16], [22], [23]. Most often these calibration routines involve injection of additional, deliberate, errors into the ADC conversion to infer the capacitor weightings and offsets through algorithms such as a simple least-mean-squared (LMS) algorithm. Sometimes it is appropriate, and effective, to use the least significant bit (LSB) or existing DAC elements to generate such a signal, which has been proven to be highly effective [14], [16], [17], [18].

In both types of calibration routines, the actual correction is generally implemented in the digital domain whereby the ADC result may undergo some transformations based on the specific bit pattern determined by the SAR algorithm to correct for the non-ideal weightings of the capacitors. The two distinguishing characteristics of a background calibration algorithm to a

foreground one, are the lack of dedicate calibration time as the process runs continually in the background during standard conversions. Both foreground and calibration algorithms may easily be extended to account for other errors within the DAC architecture. For example, residue-amplifier gain and offset errors.

Overall, calibration algorithms, whilst effective, require power consuming logic which area on the die, as well as adding latency to the conversion result. In imaging applications, conversion latency can be a key requirement. As sampling rates of ADCs increase, allocating time for the calibration of the result becomes difficult, hence why calibration routines are not able to be utilised in wireline and wireless applications. Calibration-free performance is generally limited to 14-bit performance [24], [25], [26].

2.6.4 Noise and Mismatch Error Shaping Techniques

Whilst DEM and dither move the linearity errors into the noise floor of the converter, noise shaping (NS) and mismatch error shaping (MES) techniques serve to not only randomise and decorrelate these errors from the input signal like DEM and dither does, they typically also move or shape the resulting noise outside the band of interest [24], [27], [28]. This is contrary to the wideband noise resulting from say dither which occupies all signal frequencies and hence cannot be easily removed. The techniques are usually found derived from those used in sigma-delta architectures which rely upon them more heavily. The most basic type of first order noise shaping takes the residue at the end of the SAR conversion cycle and adds it to the input in the following cycle. The implementation of this technique can significantly reduce the mismatch derived spurs in the conversion results FFTs as in Figure 2.12 which results in significantly increased SFDR and SNDR where the ADC was previously mismatch limited in these regards.




Figure 2.12. Significant improvement in SFDR may be attained through use of mismatch error shaping (MES) techniques to reduce or remove the mismatch-derived spurs shown in red. Reproduced from [28].

2.7 Conclusion

The above discussed linearisation and correction techniques are well established and frequently used in academia and industry to negate the impact of the fundamental capacitor mismatch errors and non-linearities present in all SAR ADCs. Combinations of the discussed techniques are used to achieve outstanding performance metrics in both linearity and SNR, with figures of merit being regularly pushed by increasingly complex architectures, and calibration algorithms. Nevertheless, all these techniques come at a cost and the capacitor mismatch error remains the major limitation in the performance of CDAC SAR ADCs without extensive implementation of the aforementioned techniques. This underpins the research question for this thesis: what are the dominant contributing factors to, and the nature of, the matching errors found in large arrays of lateral coupling capacitors; and are there techniques and methods possible in the capacitor design, within the boundaries of the process, that may allow achieving of higher matching capacitors.

The next chapters of this thesis will demonstrate design, fabrication, measurement, and analysis of the CDAC capacitors, seeking to analyse the aspects that may produce these errors in the first place, including the capacitor geometry and the initial sources of the errors. Any improvements made by reduction of these errors could potentially serve to reduce the need for other linearisation techniques. Furthermore, a better understanding of the nature of the errors which are corrected, can only serve to enhance and empower the existing techniques in the hands of an experienced designer. The analysis during earlier chapters then leads to the design and simulation of an optimised 18-bit SAR ADC as shown in Figure 2.13 which, as will be shown in the final chapters, achieves over 16-bit linearity without the use of any calibration techniques.

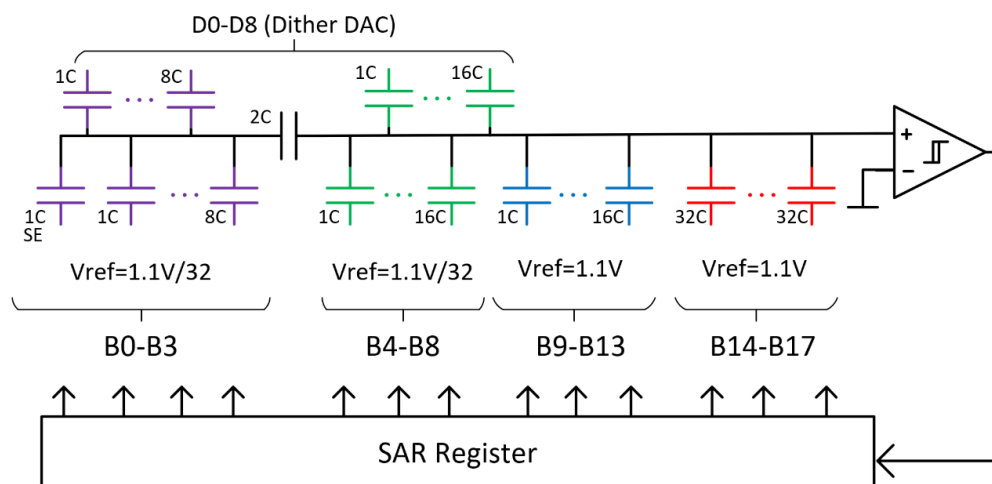


Figure 2.13. Overview of the chosen straight 18-bit SAR ADC DAC architecture. Using measurements from optimised capacitor structures this architecture achieves 18-bit linearity without the use of calibration techniques. This architecture is further explained in Chapter 7.

Chapter 3 Design and Fabrication of a Test Chip for Capacitor Measurement in a 40nm Process

3.1 Introduction

To enable analysis of the intrinsic matching properties of large capacitor arrays, it is necessary to do so with a more focused lens, and the development of exemplary test structures for that of DACs found within SAR ADCs. This chapter introduces the fundamental aspects of the chosen process for the test structures for capacitor measurement. This leads into the development of the test structures themselves and the variables for study, followed by development of the measurement strategy to enable ultra-precise measurements of the capacitor errors. The chapter concludes by optimisation of the measurement strategy, and then by its evaluation.




Figure 3.1. A survey of ADCs by their Schreier figures of merit published in ISSCC and VLSI publications in the last 5 years [3], organised by sampling frequency and grouped by the technology node used (expressed in nm).

3.2 Process Choice

Due to the investment in masks required for a modern tapeout, this project had to limit investigation to a single process node. The methodology and results later presented however,

are likely to be transferrable across most nodes between 16nm and ~90nm due to their common BEOL processes, and similar lithography techniques.

A selection of trade-offs that are commonly made when choosing a suitable process for a design are summarised in Table 3.1. As technology progresses, the demand for significant increases in speed of CMOS circuits increases. Whilst once transistor gate lengths may be measured in microns, processes for new state-of-the-art designs are frequently found at the 16nm node and below in pursuit of this. Pure analog circuits, however, tend to suffer when implemented in these kinds of processes, due to the massively increased complexity of the process, giving way to transistors that whilst exceptionally fast, are no longer able to turn off fully, have a more limited operating voltage range, and a host of non-linear behaviours that make the design of amplifiers and precision circuitry a significant challenge. As a result, it would be found that most precision reference and amplifier circuits sit in pure bipolar or bi-CMOS processes where possible, and fully CMOS amplifiers, switches and other analog functions may sit in at around the 0.18 μ m node.

As well as this, higher precision applications where speed and digital integration is not as much of a concern favour the 0.18 μ m processes, due to its low cost and ease of migration between fabs. It is considered a very mature process, where models are well established and off-the-shelf intellectual property is readily available. However, it becomes difficult to achieve higher speeds (>10MSPS) in this process due to the large device sizes and high layout parasitics. On the flipside, the 5 and 10nm processes are wholly unsuitable for precision converter design due to their very high mask cost (approximately 10x that of a 40nm mask set) and strict DRC rules. They do however allow unparalleled density and power efficiency for digital implementations and enable the use of an onboard microcontrollers for the use of advanced background calibration routines and other signal processing techniques. It is the need for digital density, as consumers expect more and more functionality to be integrated to a single device to ease integration, and reduce system cost, that is primarily pushing industry and therefore research to utilise these lower geometry nodes at the expense of analog circuit performance.

Copper metals are generally available in nodes below 90nm, providing designs with decreased interconnect resistance, and the possibility of higher density interconnects. In these processes, bespoke dielectrics for use in MIM capacitors are typically unavailable due to compatibility of the dielectrics with the copper processing. Where available (>65nm) the high-k dielectrics allow implementation of high-precision, low-drift capacitors between the metal layers. Larger geometry nodes were not chosen as MIM caps are not commonly found in modern high-speed ADCs due to the focus on the lower geometry nodes.

Were free choice of process have been allowed for this project, the 16nm node would likely be the preferred choice. This process has high prevalence amongst the top performing SAR ADCs published in the last 5 years in Figure 3.1, and provides designers with further reduced parasitics and higher intrinsic gain than the 40nm process due to the use of fin-FET transistors. The use of this process node would likely have had significant impact on the development time of the test-chips within this thesis, due to the lack of support and experience with this node available to the author. Instead, the 40nm process node was settled upon due to its ready-availability, as well as its common use in high-precision, high-speed designs within industry, and the support for this node available from industrial partners of this project. 55nm was also considered, however, this process was lacking in digital IP to ensure the test chips could be completed in a timely manner.

Table 3.1. Scorecard for common process technologies, illustrating the areas in which each process excels.

Category	0.18um	55nm	40nm	16nm	5nm
Digital Density	1	3	4	7	10
Operating Frequency	1	4	5	8	10
Power Efficiency	2	4	5	8	10
Dies Per Wafer (given fixed design)	1	3	4	7	10
Analog/Mixed-Signal Quality	10	8	7	5	3
Transistor Intrinsic Gain (gm/gds)	10	7	6	5	2
Supply Voltage Headroom	10	7	6	4	2
Leakage	10	7	6	4	2
Mask Cost	10	7	6	3	1
DRC Complexity	10	7	6	3	1
Ease of Design	9	7	6	4	1

3.3 Introduction to the 40nm CMOS Process

An overview of the process stack available in a typical 40nm process is illustrated in Figure 3.2. The chosen process is a deep n-well process, allowing for isolation of blocks from the substrate for both isolation from substrate noise, well-well coupling, as well as better control over device behaviour and performance using ‘hot wells’. The process provides several specialised types of devices for high and low threshold through additional doping options. Core devices in the 40nm process allow design within a 1.1V supply, with the additional possibility of using thick-oxide devices which allow up to 3.3V at the expense of increased device size, parasitic coupling, and speed. All of these factors are key in the high-speed and high-precision areas where noise is also a key trade-off.

The lowest metal (M1) in this process is for both the taps and polysilicon gate contacts. M1 allows smaller contact and metal-metal spacing than the rest of the metals, although is very

strict in enforcement of metal density rules. M1's use is usually limited to short range interconnects and local wiring between transistors. Metals M1 to M7 are all copper and fabricated through a variation of the copper damascene process discussed later in this section. The same low-k dielectric for reduced parasitics is likely found between all metals, contacts, and vias. There are no high-k MIM cap layers in the 40nm process, and the lateral spacing allowed is approximately 3-4x smaller than the vertical spacing between metal layers. Copper is also used for a higher metal (M7) intended for local power distribution, featuring a higher thickness, and larger vias. A thicker aluminium is used for the redistribution layer (M8/AP). This layer is also used for routing between pads for bonding, and is good for long range routing of low-impedance signals and power. A capping layer is typically found on the die with opening for the bond pads. This serves as both protection during handling and assembly, as well as a barrier layer for moisture, and, depending on the material used, for distribution of stresses (heat and packaging stress).

The 40nm process also brings with it a particular sensitivity to the density of metals, as well as transistors and doped regions, with respect to the matching and offset of circuit elements. Whilst the design rules from the foundry that are in place to address these effects are not as aggressive as that of the quite popular 28nm process node, they still pose significant challenges when layout of large analog circuit blocks is performed. This can have a significant impact on the choices of a designer during layout. Overall, the designer is very much at the mercy of the process, and conventional theories on matching and layout stress reduction must be questioned when using nodes at 40 or below due to the inherent differences in fabrication which are discussed in the following sections.

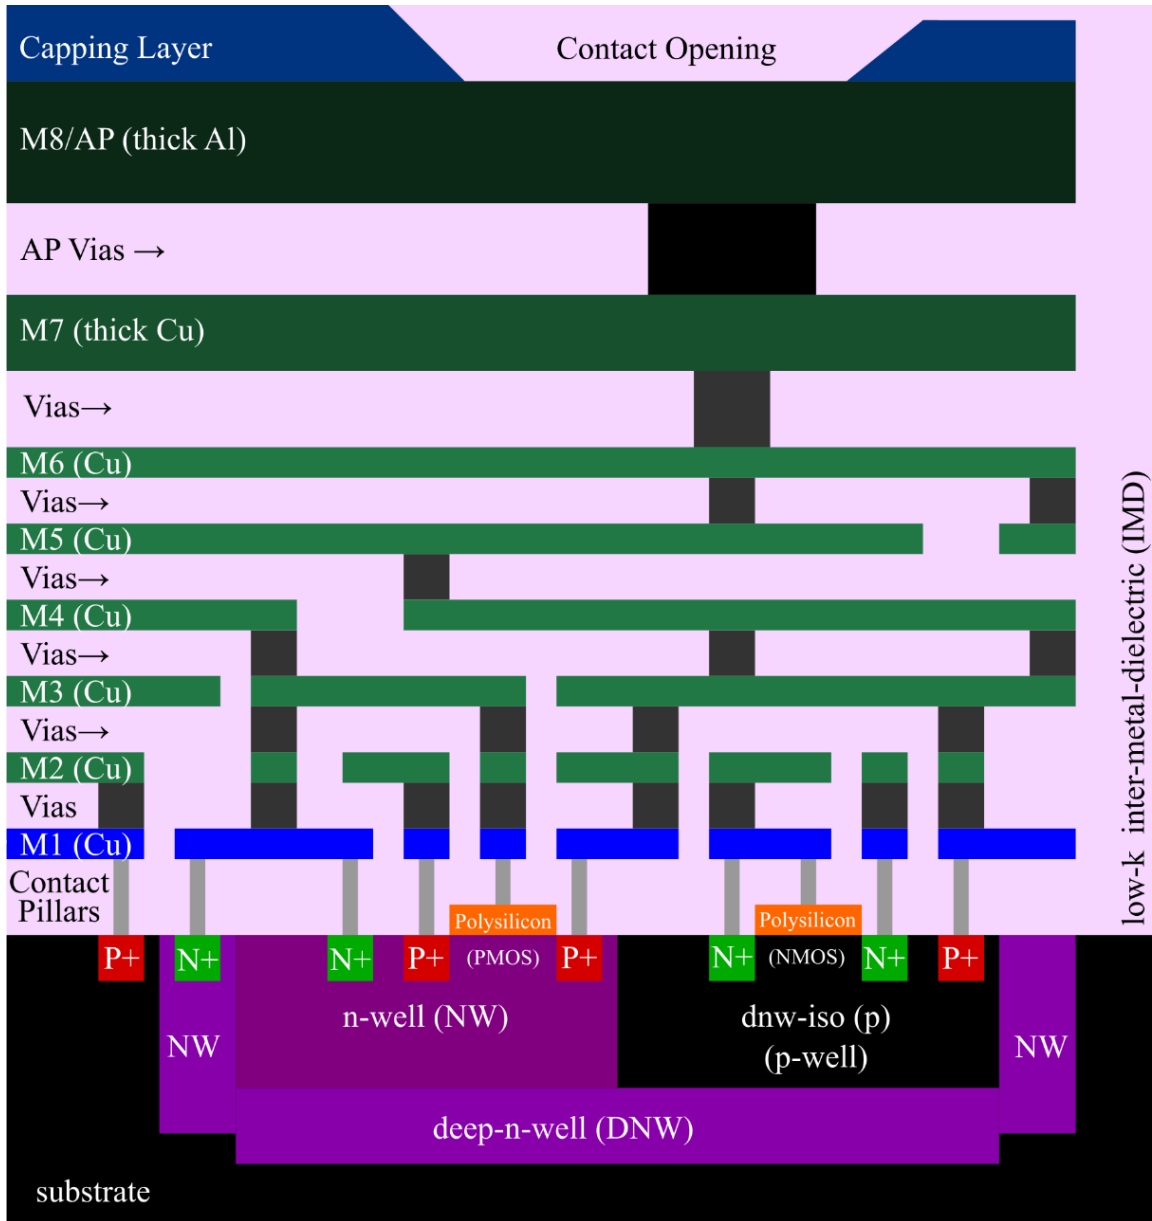


Figure 3.2. A cross-sectional overview of the 40nm process.

3.4 Overview of Key 40nm Back-End-of-Line Processes

While the transistors undergo significant scaling down to 40nm, so must the metals that connect them, otherwise the parasitics and subsequent interconnect delays will prevent any meaningful benefit from being drawn. The term back-end-of-line (BEOL) refers to the metal, dielectric, and supporting capping layers integrated above the transistors. A typical process at the 40nm node utilises between 6 and 10 copper metal layers for routing, with special intermetal dielectrics (IMD) used to reduce interconnect capacitance. Without a change in interconnect processing and design, they quickly become the limiting factor defining the speed of the circuits, due to the interconnect resistance, and parasitic capacitances, which serve to slow the signals. This has led to the shift from the use of aluminium to copper metal layers for at least the core metals (M1-M6) in processes below ~90nm due to its drastically lower sheet

resistance (~36% [29]). This is combined then with the use of low-k ($k < 3$) dielectrics for reduced parasitic capacitance, allowing significant increase in interconnect signal speed. In order to understand the complications that this has on the process, and more specifically the capacitors made from these BEOL metals, the process must first be understood.

3.4.1 Copper Dual-Damascene Process

The copper damascene process describes the means by which copper interconnects and vias are fabricated on-chip. The key steps in the copper-damascene process are illustrated in Figure 3.3.

Unlike the typical subtractive aluminium metal process where the aluminium sheet is first deposited, and the spaces between the metal tracks are etched away, the damascene process begins with deposition of the intermetal dielectric (Figure 3.3a). This is due to the difficulties in patterning the copper itself with traditional reactive ion etching (RIE) processes. The dielectric is then masked with a photo-resist, and patterned using UV lithography techniques to achieve fine pitch in patterning. The trenches, which will later contain the copper forming the interconnects, are then etched into the dielectric using RIE. This part of the process is generally repeated once for the high aspect-ratio via etch (Figure 3.3a), and again for the tracks (Figure 3.3b). It is then necessary to deposit a barrier, or liner layer such as silicon-fluoride or silicon-chloride that will separate the copper interconnect and prevent diffusion into the likely porous dielectric (Figure 3.3c). The liner and subsequent copper seed layers are formed using sputter deposition (Figure 3.3d). The copper seed enables growth of the bulk of the copper interconnect and via through electroplating (Figure 3.3e). The excess copper is then removed through chemical-mechanical polishing (CMP) and is responsible for the final surface finish and definition of the final copper wire (Figure 3.3f). Altogether, the final shape of the copper interconnect is defined by the quality of the lithography, etching (RIE), and polishing (CMP) processes. Interconnects usually end up trapezoidal due to under-etching at the bottom of the defined trenches. The precise recipe used to produce any part of the dies during fabrication is a closely guarded secret by the foundry. It therefore becomes difficult to predict the secondary aspects arising from fabrication, such as density and aspect ratio related dependencies to the final metal shape, or gauge the quality or reliability of any of the aforementioned process stages. Any of these stages may potentially impact the performance of the metal track within a sensitive circuit such as the dedicated BEOL capacitor discussed within this thesis.




Figure 3.3. A brief overview of the key steps in the copper-damascene process. Modified from [30].

3.4.2 Extreme Low-K Dielectrics

Alongside the use of copper metals comes a low- k dielectric. This can be made from a whole range of materials having a lower permittivity than that of the typical SiO_2 dielectric used in the larger geometry processes like $0.18\mu\text{m}$. Typical values for the relative constant k can be around 3 or even below, compared with that of SiO_2 of 3.9 [31]. To achieve this, the foundry combines the use of more complex materials, and introduces porosity to them in order to drastically reduce the effective dielectric constant [29], [32], [33]. The result of new material and porosity is a dielectric that offers significantly reduced parasitic coupling between adjacent metals due to the reduced k . This is massively desirable for increasing the speed of circuits in low geometry processes. The capacitance to the substrate is also reduced. Whilst all foundries keep the real figures for their dielectric a closely guarded secret, for the purposes of analysis 2.5 is taken as a reasonable figure for k [34].

The porosity of the dielectric can be particularly problematic with respect to metal structure performance and formation, as it creates a drastically different density than the metals. As with any mechanical system, this gives opportunity for concentration of transmission of packaging stresses when compared with a conventional, solid dielectric, which may be further exacerbated by temperature fluctuations. This is as the thermal expansion coefficients of the dielectric and the metals are likely to be further separated. The intermetal dielectric makes up a significant proportion of the BEOL layers of a fabricated die, and hence any stress will have notable bearing on the surrounding structures, be they the BEOL capacitors, or surrounding

active circuitry. There are also further issues between the interaction of the copper, barrier layers, and the dielectric. Clearly, adequate solutions have been successfully used by foundries for numerous years [32]. However, these issues are frequently observed in fabricated die within industry.

Overall, the combination of low-k dielectric and copper metals introduces a host of issues that would be preferably avoided. This is not always possible due to subsequent increase in achievable speed that these processes can allow in-circuit which would otherwise be unattainable. However, simply understanding the presentation of the issues through characterisation usually proves sufficient that circuits may be designed in such a way to mitigate the impact.

3.5 Lateral Fringing Metal-Oxide-Metal (MOM) Capacitor Layout

This section intends to provide an overview of the typical implementation and nomenclature associated with the capacitors, and their arrangement in large arrays as used within the smaller geometry processes below 65nm. Attention is also paid to the possible sources of error resulting from defects and variances in their fabrication. It is by no means a complete guide of every implementation and it must also be acknowledged that industry likely has its own variations on the design of capacitors that the author is not aware, nor is able to discuss.

3.5.1 Lateral and Vertical Coupling Capacitors

A capacitor in a CMOS process may be implemented in various ways. Vertical coupling capacitors are often implemented in the larger processes as metal-insulator-metal (MIM) capacitors, allowing for attaining particularly high density through use of bespoke high-k dielectrics. A typical, modern, low-geometry process does not allow for the use of such additional specialised layers due to the high costs of manufacture, and compatibility of the different processing techniques. Also, whilst they are generally very stable with respect to temperature fluctuations, the reliance on the vertical coupling can leave the capacitor susceptible to vertical packaging stresses when used in commercial products. MOS capacitors are another option, although they are generally discarded where signals are concerned as whilst they offer good density their voltage dependency introduces significant non-linearities if used in CDAC applications [35], [36]. Instead, the normal back-end-of-line (BEOL) metal routing layers are utilised to form the capacitor plates, with the standard low-k intermetal-dielectrics serving as the insulator. While these processes may not always use oxides for their dielectric, this capacitor is usually referred to as a metal-oxide-metal (MOM) BEOL capacitor.

The BEOL MOM capacitors can be implemented vertically or lateral, as shown in Figure 3.4. In larger geometry processes, similar capacitive density (C/m^2) and matching performance is attainable [37] between structures reliant on the vertical and horizontal coupling [37]. In smaller geometries optimisation is achieved by utilising the lateral coupling between drawn capacitor fingers in order make the most of the high-resolution lithography. Theoretically, the lateral approach yields decreased matching errors as the planar resolution from lithography is high, and well defined when compared with the vertical etching and deposition which defines the thickness of the capacitor plates and the dielectric between them. This also avoids the additional vertical stress errors that are expected with vertical coupling capacitors.

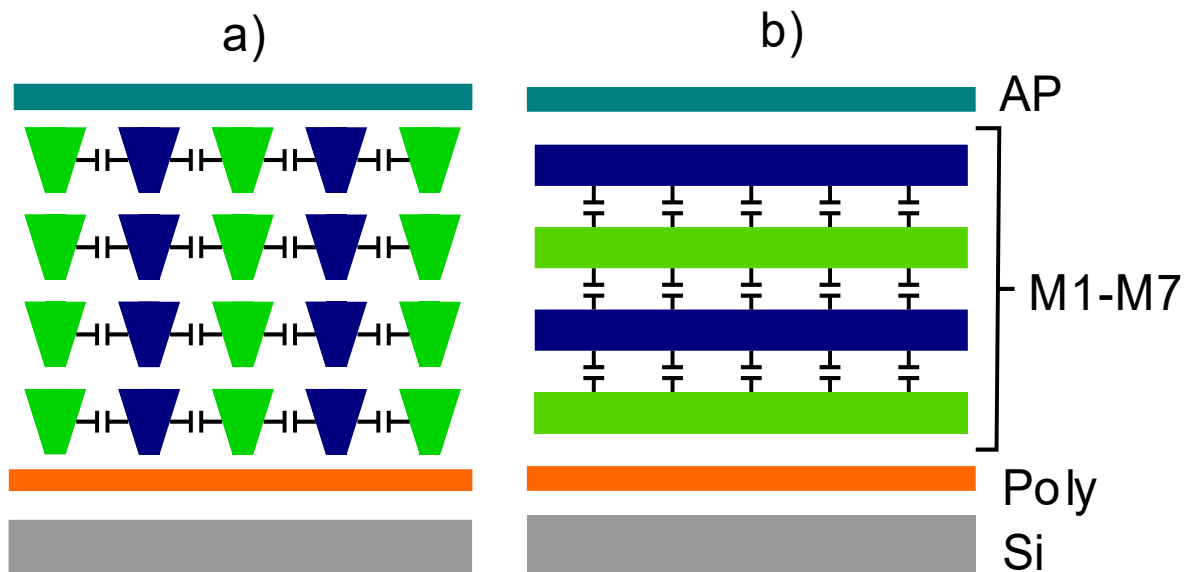


Figure 3.4. Cross-section of typical a) lateral coupling, and b) vertical coupling BEOL capacitor implementations. A combination of the two may also be used to increase density.

3.5.2 Foundry Cell Library Capacitor Structure

The specific top-view layout of the lateral coupling capacitor varies somewhat between various process design kits (PDK) and individual designer. As a reference, the standard layout chosen by the foundry for this style of capacitor is first described. The simplified geometry of the standard library cell capacitor is shown in Figure 3.5. This geometry is implemented identically in metals 2 to 5, with 8 bottom plate fingers each side of the top plate spine for each unit capacitor as shown. As a result, the unit capacitor sits on a pitch of $1.54\mu\text{m}$. Each capacitor is designed to target 32 fF each, yielding a σ_T of 1278 ppm according to the foundry-provided simulation model.




Figure 3.5. Illustration of the basic foundry capacitor geometry used for comparison in this work. This capacitor reflects the typical implementation with no established grounds for optimisation. The top plate is joined to adjacent cells in another metal not shown.

Each array of foundry capacitors uses 3 dummy capacitors on the top and bottom edge of the array, as well as a single unit capacitor either side to provide lateral dummyming. This amount of dummyming reflects the guideline approach to dummy structures for capacitor arrays within the 40nm process, with a limited view to compensate for edge effects from CMP deformation and uneven polishing. These are the predominant error sources. A single one of these foundry based 40fF capacitors measures $12.76 \times 1.47\mu\text{m}$ and occupies an area of $18.75\mu\text{m}^2$.

3.5.3 Typical Design and Layout of Single Finger Capacitors

A compilation of the favoured implementations of laterally coupling MOM capacitors is shown in Figure 3.6. These structures exclusively utilise the lateral coupling between metal fingers to form the capacitor and duplicate the patterning and connections across numerous layers to increase the overall capacitance as required. They all share similarities in their comb-like arrangement where multiple fingers are used, with broad differences in finger length, spacing, and width being found across processes, and between authors. There is insufficient data presented in these cases to draw any conclusion relating to favoured trends in geometry or sizing, as these choices are often left undiscussed. Despite this, no extensive report has yet been undertaken to justify the decisions in design of such capacitors. The layouts vary in both their specific sizes, as well as design decisions pertaining to external routing, via placement,

and use of dummy structures, preventing objective judgements being made in respect of each factor impact on the matching performance of each structure. In most cases the dummy structures used in these instances is not specified. A generalisation of the most common arrangement of laterally coupling capacitor is presented in and Figure 3.7, and is the general geometry that is taken forwards for investigation within this thesis.

Particular attention may be brought to the images in Figure 3.6a, c and e. These layouts utilise fractional length fingered capacitors in search of attaining higher dynamic ranges for their arrays without further increase in the number of unit capacitors. These novel arrangements in capacitors are also not further discussed in their impact on their matching, or other circuit performance in the context of their layout, although it is clear that this is one avenue for improving dynamic range with reasonable performance [38].

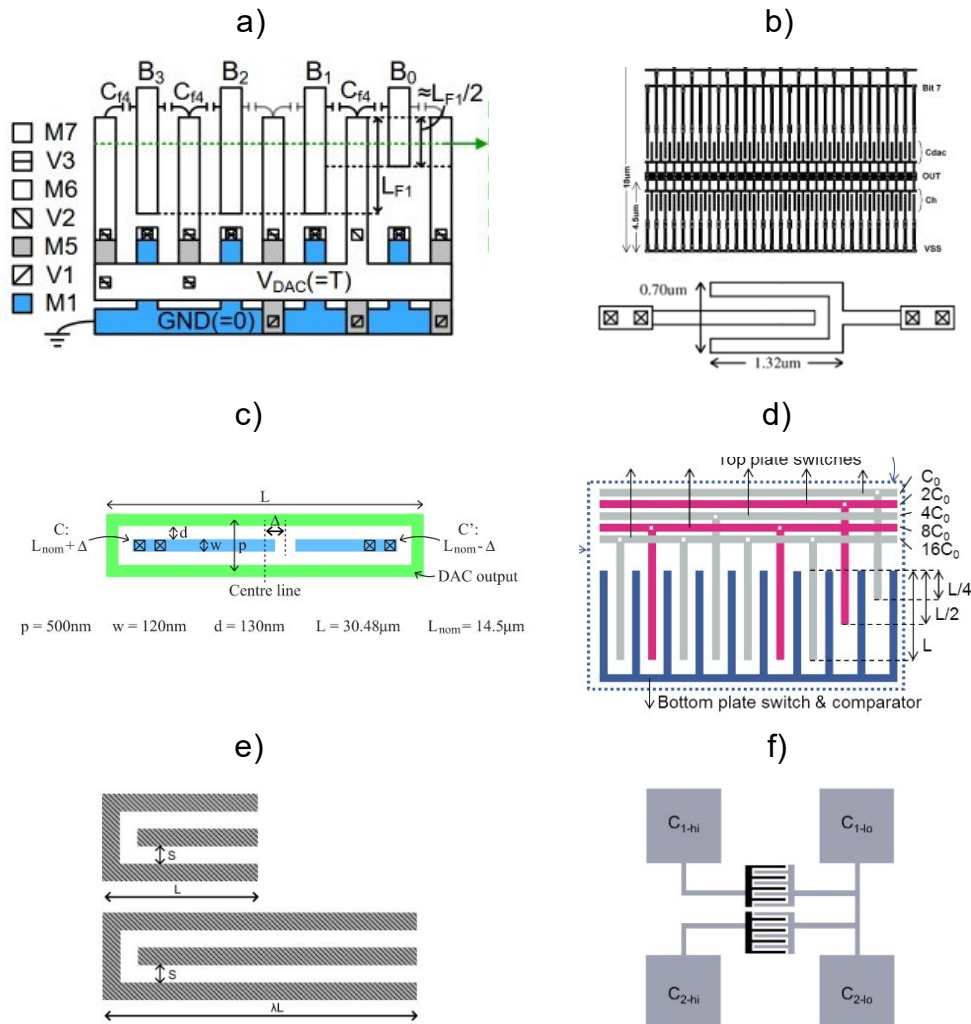


Figure 3.6. Assortment of lateral coupling capacitors from the literature for illustration of the consistencies in implementation. Reproduced from: a) S. Lee 2024 [39], b) P. Harpe 2011 [40], c) P. Harpe 2019 [38], d) N. Le Dortz 2014 [23], e) V. Tripathi, B. Murmann 2014 [41], f) H. Tuinhout 2009.

The basic capacitor structure studied in this thesis is shown in Figure 3.6 and Figure 3.8. The figure shows to approximate scale the capacitor geometry, and the spacing between the used 4

layers, and does not utilise any vertical coupling, aside from the expected diagonal coupling of Figure 3.6. As discussed, the dielectric between the copper capacitor fingers is the standard BEOL low-k dielectric. The bottom plates located directly above and below each other are connected with vias externally to the structure. For the metal layers directly above and below the capacitor structure, there is no metal density fill, nor shielding, except where this was the subject of investigation in later chapters. No poly shielding was used below the structures and thin stripes perpendicular to the capacitor fingers were used on larger designs in order to meet DRC poly and metal density requirements. Exact thickness for the vertical metal and dielectric stack is not available due to confidentiality agreements. From experience with similar processes, the metals are likely to be formed as slightly trapezoidal due to the high aspect ratio, however there is no data available in 40nm to confirm this.




Figure 3.7. A 3D representation of the BEOL capacitor structures that this thesis concerns. Four capacitors are shown. The top (common/shared) plate is marked in blue and the actuated bottom plates are shown in green.

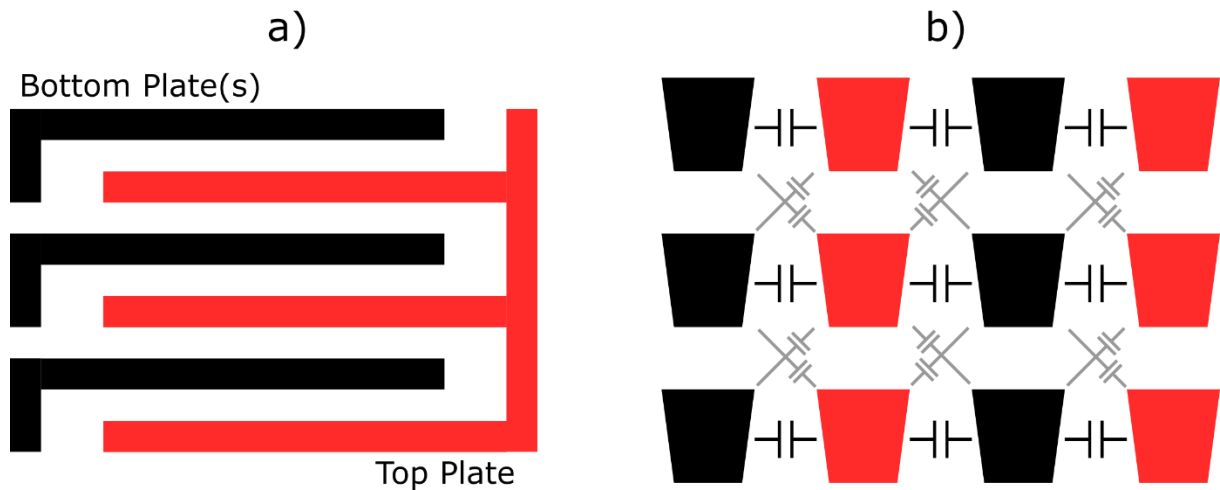


Figure 3.8. The typical lateral coupling BEOL MOM capacitor implementation. a) Top view. b) Cross-section.

This geometry forms the foundation for characterisation and optimisation within this thesis.

3.5.4 Design and Layout of Capacitor Arrays

An n -bit differential DAC requires minimally $(2^n - 1)$ unit capacitors, split between a P and N-DAC, which are typically mirror images of each other for layout symmetry. Precision SAR architectures utilise between 8 and 12-bits per stage, to achieve their speed and bandwidth requirements, as well as minimising the area of each capacitive DAC. Figure 3.9 shows the implementation considered for variation in this work. It is a differential 8-bit DAC of 512 total elements, split into two mirrored arrays of 256 single finger unit capacitors. Each unit capacitor has its own individual set of bottom plate switches and is operated independently. A single dimensioned array is utilised in this investigation. While this is typical of a high-speed SAR ADC which is the initial reason for this choice, it also allows better analysis of any systematic error across the die given the small capacitor pitch it produces. While it is possible to implement a two-dimensional array [31], where each capacitor cell is more rectangular and comprised of more fingers, in applications where both precision and speed are being targeted as measures of performance, the capacitors are often in the 5-30fF range. Implementation of capacitors of this size in a 2D grid would require multiple fingers per capacitor. This substantially increases the parasitics due to the necessity to remotely locate the drivers away from the capacitor, not to mention the increased complexity in routing of the capacitors in a 2d array. The gradient requirement measurement is quite often acknowledged by the literature and is usually used as an argument towards the use of 2D CDAC array where further dimensional folding and cross coupling can be used to average out the die gradient effects and maximise the proximity of instances of all capacitors. However, such gradients have been shown to be substantially more complex in studies that consider the larger scale ($>100\mu\text{m}$)[42], with no current data for the smaller scale which may affect such DACs that are considered in this work.

The capacitor cells are directly abutted and the top plates joined in the centre of the DAC, with their respective drivers positioned towards the outside of the structure as shown in Figure 3.10. This layout is optimal for connection to further stages, such as to a residue amplifier and comparator in a practical ADC design, due to providing the shortest possible routing between the top plate and the adjacent block. Above and below the operational capacitor and driver sections, vertical dummy structures are placed. These are replicas of the operational section of the DAC (capacitor and bottom plate driver) but are never operated, and are hardwired to a fixed state in the design. Similarly, a lateral dummy structure designed to imitate the patterning of the capacitor fingers is implemented either side of the array between the capacitors and their drivers. Where possible, the fingers of the lateral dummyming are interleaved with the capacitor bottom plate terminations to ensure maximum continuity in the metal patterning and density. A total of $30\mu\text{m}$ of vertical dummyming is implemented above and below each array, and a nominal $5\mu\text{m}$ of lateral dummy structure is included either side.

The routing for the capacitor drivers is done in a very similar fashion to that of a typical SAR ADC implementation, albeit with a significant number more individual wires being run due to the lack of binary weighting or grouping of unit capacitors. This also ensures consistent parasitics and delays in the lines feeding the drivers, as the toggling of the capacitors is operated externally to this block. The wires are run vertically either side of the array for the P and N arrays respectively. A total of 264 wires was necessary either side of the array and implemented in metals 2 and 4.

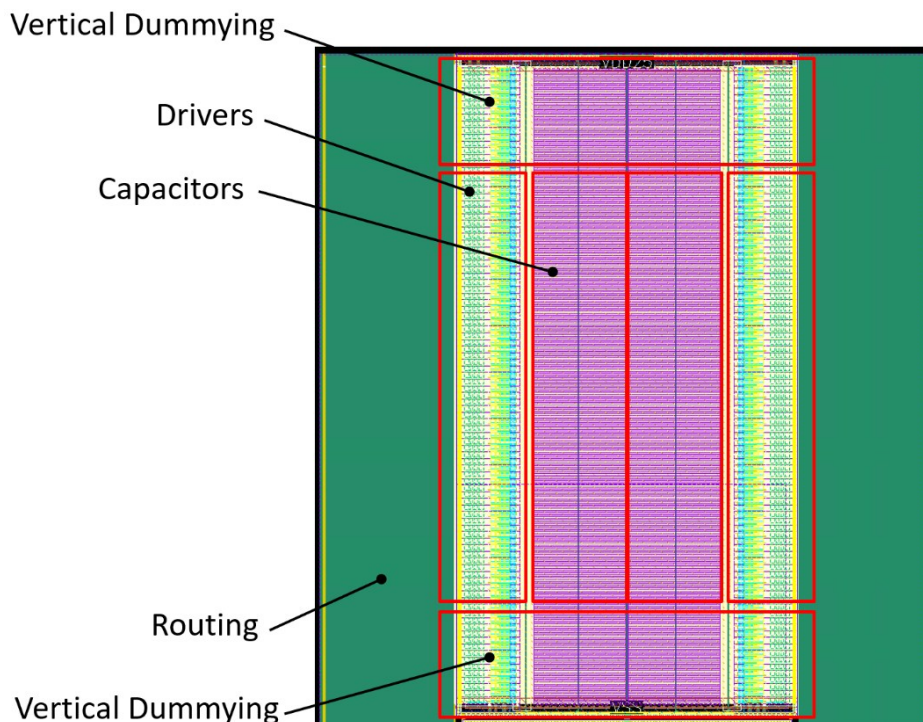


Figure 3.9. This arrangement of capacitor DAC structure is used for all optimised designs, resulting in an approximately $120 \times 100\mu\text{m}$ structure which includes the vertical routing on either side.

An expanded view of the capacitor array is shown in Figure 3.10. As the nominal capacitor is implemented in metals 2 through 5 it is necessary to locally join them with vias. The drivers connect to the capacitor bottom plate in metal 3, and the stitching via is located outside the main capacitor structure as shown and enclosed by the lateral dummy structure where possible. The intention is to relieve the stresses that may be introduced to the die from the presence of the via structure so locally to the coupling area of the capacitor as previously seen in Figure 3.5, whilst maintaining an even distribution of metal as is possible through the use of the lateral dummy structures. This is not always possible due to the proximity of the adjacent metals in the lateral dummy structures to the stitching via in the variations close to minimum geometry due to the foundry design rules. In these cases, the lateral coupling is scaled back to the outside of the stitching via.




Figure 3.10. A magnified section of a final capacitor array layout showing the positioning of the drivers, lateral dummying and the actual capacitor geometry with top plate connection. Metal 3 connects the drivers to the capacitor bottom plates but is not shown for clarity.

3.5.5 Sub-DAC Implementations

When implementing DACs with dynamic range in excess of 8-bits, it becomes necessary to re-evaluate the choice of DAC architecture, as this easily leads to arrays that are in excess of 100 μm long due to the number of total unit capacitor elements. Possibly the most common technique to alleviate this issue is to implement a sub-DAC architecture [43], [44], [45], [46], [47] as shown in Figure 3.11. The version shown allows an example implementation of the same 5-bit DAC using only 11 unit capacitors instead of the usually required 31 to demonstrate the scaling capability. Note this also significantly reduces the size of the MSB capacitor. This is achieved through the addition of a series capacitor, creating ‘Main DAC’ and ‘Sub-DAC’ sections to the whole DAC. The ‘top plate’ connection is still connected to the main DAC top plate, whilst the sub-DAC top plate connects through a series coupling capacitor to the ‘top plate’. Simply put, the series coupling capacitor effectively reduces the charge moved by operation of the sub-DAC capacitors to the ‘top plate’, allowing a higher dynamic range in capacitor bit weights without high capacitor ratios. Therefore the MSB capacitor may be significantly smaller than a topology without a sub-DAC.

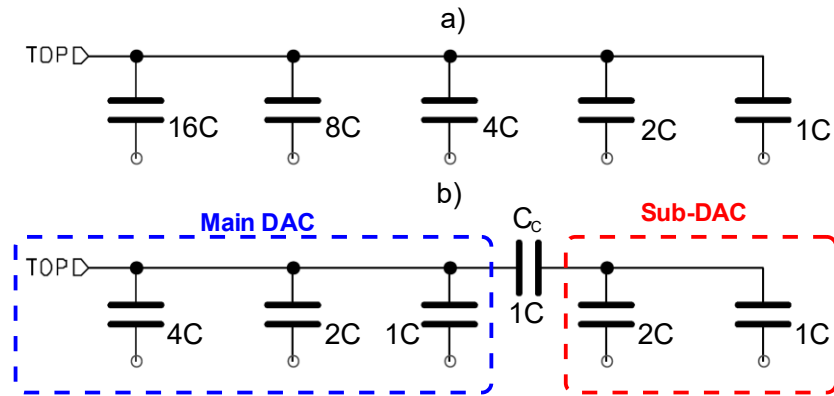


Figure 3.11. Typical sub-DAC implementation to reduce array area. a) shows a stereotypical straight 5-bit CDAC architecture with 31 unit elements, b) shows a functionally equivalent DAC with 2-bit sub-DAC implementation using only 11 unit elements.

In normal operation of the DAC, the top plate receives an injected charge determined by the bit trail capacitor switched. For example, in Figure 3.11a, when the bottom plate of the LSB 1C capacitor is switched from 0V to V_{ref} : CV_{ref} of charge is injected to the top plate and redistributed throughout all capacitors. This results in a net change in top plate voltage of $\frac{V_{ref}}{31}$ V. Through this arrangement, the smallest amount of injectable charge to the top plate is determined directly by the LSB, or unit-capacitor.

By comparison, the sub-DAC allows injection of fractional amounts of charge than the unit capacitor, and therefore allows more bit trails for a given unit and MSB capacitor size. The same principle applies. In the case of the Sub-DAC MSB (2C) of Figure 3.11b, $2CV_{ref}$ of charge is injected to the sub-DAC top plate on the right-hand side of the coupling capacitor. $V_{ref} * (2C || 1C)$ of charge is injected to the top plate and equally distributed across the total effective capacitance of the top plate. In this case $7C + (1C || 3C)$. When considered proportionally, the LSB of the sub-DAC shown gives the same voltage shift as that of the LSB of Figure 3.11. Overall, this technique allows achieving high ratios of charge sharing within the DAC structure, without creating exceptionally large arrays. This architecture can be extended to include a second coupling capacitor from the sub-DAC to a second sub-DAC. This significantly extends the dynamic range of the circuit further whilst maintain the same benefits, in principle.

This technique is far from perfect, however. While in a standard DAC the top plate parasitics generally result in a reduction in speed of the converter, any sub-DAC top plate parasitics here also significantly impact the sub-DAC bit weightings. This manifests as a gain error between the main and sub-DAC bits, with errors often in the region of found to be 2-5%. This is evidenced later in section 5.5.2 of this work. These errors are extremely difficult to accurately predict in simulation and through use of layout extraction tools as they are related to an absolute value in parasitic capacitance. This is made worse by the small values of the coupling capacitor.

Typically, a designer will scale the entire sub-DAC by a factor of 2 or more in order to use a larger coupling capacitor to alleviate this issue and reduce the error due to limitations of the parasitic field solver which is generally relied upon for tuning of these parasitics. It is also desirable to increase the capacitor size to reduce individual capacitor mismatch, although this is not generally the highest concern with this topology due to the dominance of the gain error.

Issues also arise with respect to the layout of the sub-DAC structure. General experience suggests that a layout that is disorderly and irregular will likely yield worse matching errors than one who's layout is perfectly regular and uniform. The sub-DAC necessitates a discontinuity in the top plate of the capacitor layout due to the series coupling capacitor, giving surrounding capacitors an irregular boundary. The impact of this will be discussed later when the localised matching of capacitors is considered but is shown in Figure 3.12 below. This will be shown to lead to an undesirable localised increase in capacitor error. Any technique to eliminate the need for a sub-DAC, or simply to reduce the necessary number of bits it must be used for, would be preferable from several perspectives if it removes the need for or reduces the currently required level of calibration.




Figure 3.12. Excerpt from real layout of the coupling capacitor used for a sub-DAC implementation. The sub-DAC and main-DAC top plates are highlighted green and red respectively. A large discontinuity in the vias and metal connections in the spine (centre) are visible.

3.6 Design of an Ultra-Precise Capacitor Measurement Architecture for Large Arrays

3.6.1 Requirements and Overview

To extract the statistical matching performance of large capacitor arrays, it is inherently necessary to measure the weightings of the capacitors within a large array. Whilst there exist indirect measurement methods in which a SAR ADC is developed directly around the DAC, this then requires analysis to be undertaken on the ADC output code, to indirectly infer the matching performance and errors within the DAC. This indirect method easily obscures secondary or localised errors, as well as the true nature of the errors, substantially limiting the dimensions available for analysis. As it is the goal of this thesis to better understand the specific nature of the errors, individual capacitor error is required.

A summary of the geometry variations that were implemented are shown in Table 3.2. As there is a large amount of overlap between the sweeps and possible values, the taped-out designs were a subset of these and are summarised along with their matching statistics for later reference in Appendix A.

Table 3.2. A summary of swept capacitor variables and their ranges.

Geometry Parameter	Variations
Finger Length	14.5 μ m – 72.5 μ m
Finger Spacing	70nm – 140nm
Finger Width	70nm – 140nm
Number of Metals	2 – 6 (+higher and lower metal combinations)
Top/Bottom Dummying	5 μ m – 30 μ m
Lateral Dummying	0 – 5 μ m
Optimised Sub-DACs	2 Variations

The space procured for each test-chip was a 2x2mm area, allowing for a maximum of 32 unique arrays per test-chip without the risk of placing structures too close to the pad ring, as bond-pads are a known source of error[48]. A total of 3 test chips were possible with the first only containing only a structure for measurement of the foundry capacitor. This structure was replicated on the second test-chip to ensure consistency and allow direct comparison with the custom designed capacitor cells. As much logic and amplification as was feasible was

implemented on-chip, using a common 2.5V supply for simplicity. This enabled higher speed operation of the entire test chip as, due to the methods that will be shortly discussed, automation of the measurement process was paramount to ensuring repeatability and a large number of dies may be measured. Approximately 300+ samples were requested for each test-chip design, and as many were measured as the time allowed for.

The following presented system allows individual extraction of each unit capacitor value, fully decorrelated with any measurement derived errors, providing high-precision measurement data that is later used for full analysis of the capacitor errors in Chapter 5.

3.6.2 Charge Based Measurement Technique

Other investigations [42], [49] have favoured a more direct method of measurement whereby a large number of capacitor array structures are laid out across the die, and a select few capacitors are individually connected out to pads to enable probing and measurement with an external LCR meter. Whilst this method proves sufficient for some characterisations, and is conveniently simple, the area-costly nature of the resulting design severely limits the number of measurable capacitors on the die, as well as any space for many variations in geometry, due to the requirement for probe pads. It also requires the dies be unpackaged, or at least unencapsulated. From the perspective of commercially packaged ADCs, this would mean that any packaging introduced stresses and relaxations that may occur, would not be captured in measurement. Not to mention that proper statistical analysis of capacitor errors requires a considerable number of measurements to build confidence in any statistics and trends that may arise.

Instead of this, a less direct measurement technique is chosen. The charge-based measurement technique allows the capacitors to be laid out in much the same configuration as they would be in a typical CDAC design, and then switched such that their captured and released charge may be measured [50]. The released charge is proportional to the capacitor weighting inclusive of their mismatch error. The charge is then amplified and provides a voltage at the output of the die which is then sampled externally, without loading of the capacitor itself as a factor. As will be shown, this method is largely insensitive to the parasitics of the circuit, and lends itself well to implementation in the 40nm process, whilst being compact enough to allow for up to 32 different capacitor geometry variations on a single 2x2mm die. The three stages of operation of the basic circuit are shown in Figure 3.13 as the Reset, Float, and Measurement phases.

During the reset phase the source capacitor C_a , and reset capacitors are both reset to ground by shorting them. During the Float phase, the reset switches are released. Timing of this phase

is not critical but there must be no overlap between the float and measurement phases. During the measurement phase the source capacitor is toggled. As the amplifier keeps a virtual ground at its input the injected charge is forced onto the feedback capacitor, providing a proportional output voltage.




Figure 3.13. Stages of operation of the charge amplifier.

To implement this for measurement of the large capacitor arrays, a unary DAC is constructed of 256 capacitors of nominal value C , with each capacitor having its own independent bottom plate switches implemented as inverters. This is the same as the physical structure previously illustrated in Figure 3.9, and follows the same phases as Figure 3.13. The arrangement is shown in Figure 3.14 for an array of $(n + m)$ total unit capacitors. During the reset phase all capacitors are shorted to ground and discharged. In the float phase, the top plate shorting switch is released, and allowed to float. In the third phase, the bottom plates of n capacitors are switched from ground to reference voltage V_{ref} , redistributing an amount of charge proportional to the weighting of the n switched capacitors to the top plate. This charge would normally be distributed across all capacitors and the top plate parasitics C_{po} if the charge amplifier was not used. This would result a top plate voltage V_{TP} as in equation 3.1.

$$V_{TP} = V_{ref} \frac{nC}{(n + m)C + C_{po}} \quad 3.1$$

The top plate parasitic attenuates the voltage developed on the top plate in this configuration as it would in a regular CDAC, hence the charge amplifier is the preferred solution to avoid this issue. With implementation of a charge amplifier connected to the top plate, the ideal voltage developed at the output V_{out} is given in equation 3.2 and is independent of the top plate parasitics.

$$V_{out} = V_{ref} \frac{nC}{C_{fb}} \quad 3.2$$




Figure 3.14. The three steps involved in toggling the specific capacitors for measurement of their weighting.

The switched charge is proportional to the capacitor weightings.

It's worth noting that depending on the operation frequency there may be dielectric absorption effects at play which would affect the ultimate capacitor measured value. This effect is on the same order of magnitude as the expected capacitor error ($\sim 1000\text{ppm}$). However, unlike the mismatch errors, it will manifest with some frequency dependence. The assumption is therefore made that the percentage contribution from any dielectric absorption effects that may be at play is equal in all capacitors, so provided that all measurements are done using the same times for each phase of the process, the contribution from the dielectric absorption is also constant, and serves as an offset proportional to the capacitor weighting [51], [52], [53]. The pattern generator chosen uses a clock of 200 MHz which earlier work as shown to be of the same order of magnitude as the dipolar time constant associated with this process. This leads later to largely increased reset times to ensure this effect is minimised.

3.6.3 Die-Level Measurement Architecture

The extraction of the weighting of every single capacitor in each array is desirable as it allows information of any identified systematic errors to be directly observed with fine spacial pitch. The basic die-level architecture for the measurement of a single array of capacitors is shown in Figure 3.15. Each of the 256 capacitors on each side of the arrays are individually programmable to toggle between 2 pre-defined states through a 1024-bit shift register. The states are programmed into the digital block which provides a parallel output bus running to multiplexer. which receives the two programmed settings for each capacitor, and using an external toggle signal, is able to simultaneously switch all capacitors between the two states. The multiplexer feeds the bottom plate drivers which are contained within the DAC structure as previously described.

All capacitors in the array are identical, and the patterns for the N and P array toggling are under most circumstances complementary. The output of the DAC is fed into a differential charge

amplifier of bespoke design. This charge amplifier provides an output voltage which is, by design of the switching scheme, proportional to the mismatch of the capacitors.

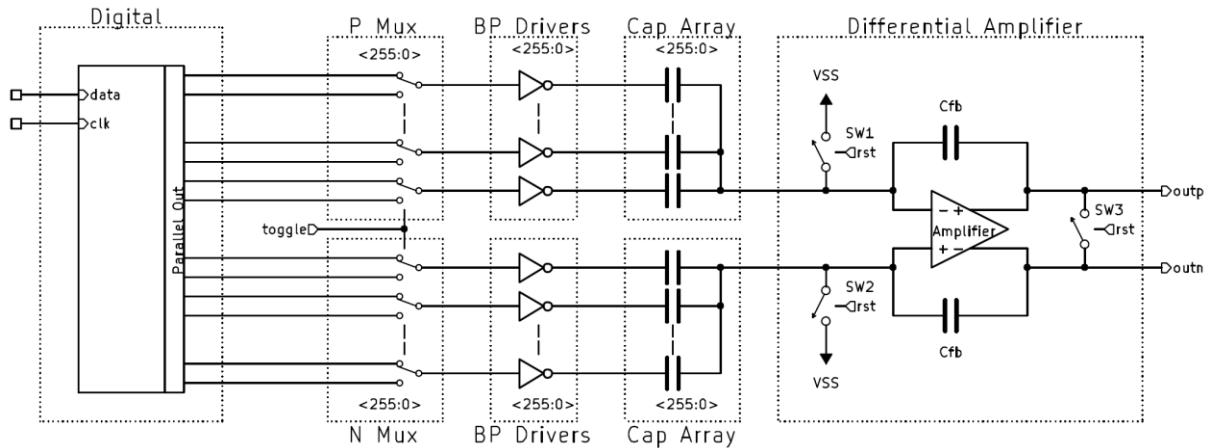


Figure 3.15. Architecture for measurement of a single differential array of unary capacitors, with integrated charge amplifier. The digital and P/N Mux elements shown are shared between multiple DACs.

At the wider die level of the design, 32 individual arrays are instantiated. All arrays share the same programmed pattern at any given time. This programmed pattern also contains the enable bits to ensure that only a single amplifier is powered and connected to the output pins at a given time. This is achieved through the multiplexer arrangement shown in Figure 3.15, whereby there are 8 channels of 4 arrays, with a single set of multiplexers from the digital for each channel. The output multiplexers ensure that for each channel, there is only one active and connected charge amplifier at any one time, to prevent competition between arrays. Through the introduction of similar logic in the toggle path, only the toggle multiplexer for the respective enabled array actually performs the toggle, and all reset and disable signals are held high. This minimises the switching activity on the die and reduces any noise from the supply.

As far as timing and durations of the three described phases of Figure 3.14 and Figure 3.16 are concerned, the reset phase is dedicated a significant amount of time to ensure that all capacitors are fully discharged, and all nodes are settled. Realistically this only required $1\mu\text{s}$ for 10τ settling, however as it would change for different sized capacitors, this remains fixed at $10\mu\text{s}$ for all measurements. The second float phase is made as short as possible whilst ensuring that the switch does not overlap with the settling phase. In practice 1 clock cycle was used to ensure sufficient margin, giving a minimum non-overlap of 50ns.

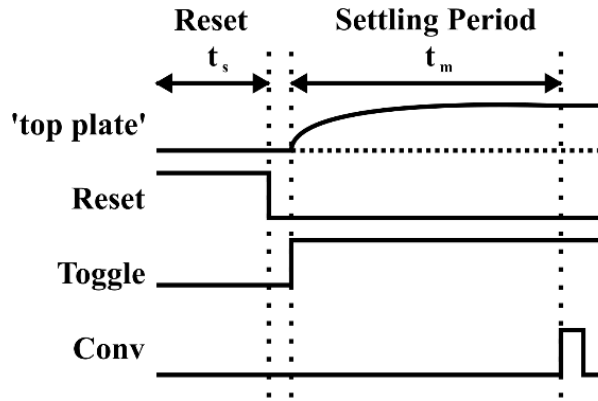


Figure 3.16. Switch sequencing timing diagram for measurement of capacitor mismatch, illustrating the three phases of the measurement.

3.6.4 Reference Capacitor

The feedback capacitor of the charge amplifier discussed was chosen to be identical for all array implementations, irrespective of capacitor value. This was done such that the absolute weighting of each capacitor may be extracted, without the localised effects that were expected influencing the value of the feedback capacitor. Of course, there will still be die gradient and localised processing effects, however the base assumption is that the feedback capacitor is located close enough to the array being measured that the effect of these gradients will affect both arrays equally. This ensures that at minimum the relative value between the capacitor geometries may be extracted.

Further to this, a reference capacitor was implemented on the die with the same parameters in Table 3.3, with the top and bottom plate terminals connected to the pad-ring, and bonded out to the package. This enables it to be directly measured in the evaluation setup using a bench LCR meter. While the feedback capacitors had only 31 fingers, giving a total value of 301fF, the reference capacitor was a 19.9pF capacitor with 2048 fingers. It included 10 μ m of top and bottom array dummyming, as well as 5 μ m of lateral dummyming, although the sheer size of the array makes it unlikely that it would be affected by the presence of the localised errors which will otherwise be discussed. The reference capacitor measured a total of 370x53 μ m.

Table 3.3. Parameters of the die reference capacitor to enable measurement of the capacitor values across all measured geometries.

Parameter	Value
S	100nm
W	70nm
L	14.5 μ m
N_m	4
N_f	2048
Width	53 μ m
Height	370 μ m
Value	19.9pF

3.6.5 Capacitor Value Measurement

As is now established by prior section 3.6.2, the output of the die is a voltage proportional to the sum of the switched capacitor values, scaled by a gain determined by its pairing with the feedback capacitor. A given capacitor value C_i for index i , within a single array, may be considered to be comprised as the sum of the nominal designed capacitor value C_{nom} , plus some additional, unwanted error E_i which this study intends to measure. Shown by equation 3.3.

$$C_i = C_{nom} + E_i \quad 3.3$$

Each measurement cycle switches n capacitors positively. The charge is then transferred to the output. This means the output voltage may be expressed as in equation 3.4.

$$V_{out} = V_{ref} \frac{nC + \sum_{i=1}^n E_i}{C_{fb}} \quad 3.4$$

As this the design is differential, and intended to be operated as such to maintain the input common mode, toggling of the capacitors in the above way is undesirable, as it severely limits the number of capacitors that may be measured at any one time due to the large output swing required from the nC_i term. This also introduces susceptibility to any non-linearities introduced by the amplifier design. To avoid this, instead of only resetting and toggling n capacitor bottom plates from ground to V_{ref} , n capacitors are also reset to V_{ref} and toggled to ground, hereby removing the nC_i term and leaving only the error term in V_{out} as seen in equation 3.5. The

relative weighting of the capacitors is determined by switching a small group of capacitors positively without this cancellation. This enables normalisation of the extracted capacitor errors with respect to the absolute capacitor weighting. It is not possible to extract the capacitor errors through individual toggling with good resolution due to their size (<100ppm). This is due to the aforementioned susceptibility to the amplifier non-linearities when using a large output swing.

$$V_{out} = V_{ref} \frac{\sum_{i=1}^n E_i}{C_{fb}} \quad 3.5$$

The feedback capacitor was selected to be nominally 300fF. This value was chosen to allow measurement of at least 2 capacitors without complementary switching for normalisation for all capacitors variations, within the accepted linear region of the amplifier's output. Normalisation of the capacitor errors is then achieved with equation 3.6. In all possible cases, the same number of capacitors were toggled positively as there were toggled negatively. This was to improve the power supply rejection of the signal chain, as balanced switching ensures that any coupling of the power supply ringing or settling errors will be only seen as common mode errors in the system, and should not significantly affect the measurement.

3.6.6 Capacitor Error Extraction

To facilitate the extraction of individual capacitor weightings using the described methodology an assignment matrix is utilised, referred to hereby as the shuffle matrix, in which each row defines the combination of capacitors that are to be switched for a single measurement, and in which direction (positive transition, or negative transition). Each element takes its value of -1, 0 or 1 which represent these physical transitions, as well as their mathematical result. This representation is used in the following matrix operations.

Letting C be the capacitor weightings within the array, and M the shuffle matrix, the voltage step (or other output signal) observed at the output of the signal chain may be represented as in equation 3.6. Equation 3.7 visualises of the operation result.

$$MC = V_o \quad 3.6$$

$$\begin{pmatrix} 1 & 0 & \dots & -1 \\ 0 & -1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & 0 & \dots & 0 \end{pmatrix} \begin{pmatrix} C_0 \\ C_1 \\ \vdots \\ C_{n-1} \end{pmatrix} = \begin{pmatrix} V_0 \\ V_1 \\ \vdots \\ V_{n-1} \end{pmatrix} \quad 3.7$$

It naturally follows that given knowledge of M and V_o from the measurements, the original capacitor weightings may be derived using the inverse of M through equation 3.8.

$$C = M^{-1}MV_o = M^{-1}V_o. \quad 3.8$$

This of course implies that there exists an inverse to the shuffle matrix M . Due to the preferred use of large over-defined, rectangular matrices in place of M to reduce measurement noise and perform further averaging of the weightings, a linear regression must be utilised to extract the individual capacitor weightings of C . This regression utilises a least-mean-squared (LMS) algorithm to minimise the error between MC and V_o .

The prime benefit to this method of measurement of groups of capacitors with a post-processing extraction of the individual contributions, is the enhanced effective measurement resolution. Assuming first that each capacitor has gaussian error distribution $N(\mu, \sigma^2)$, and that each measurement result is independent. It follows that the distribution of the variable that is the sum of two selected capacitor weights is given by equation 3.9.

$$C_T \approx N(2\mu, 2\sigma^2) \quad 3.9$$

It naturally follows that if a series of sums of β capacitors are measured, the resulting distribution will approximate to equation 3.10 with a standard distribution that has been scaled by $\sqrt{\beta}$ on the assumption that the errors under measurement are decorrelated.

$$C_\beta \approx N(\beta\mu, \beta\sigma^2) \quad 3.10$$

As the output voltage from the system is proportional to the switched capacitor values, this also increases the distribution of the output signal by a factor of β , but does not amplify the noise. This increases the SNR of the measurement.

Prior research has shown that the capacitor errors have some level of correlation at short distances [49], so we may expect up to a 50% reduction in effective resolution gain from this method. However, the simplicity of implementation, coupled with the aggressive averaging that it provides while providing some level of signal gain, this makes the application of the shuffle matrix measurement plan described here still the best option for developing switching trials and sequences for accurate capacitor measurement.

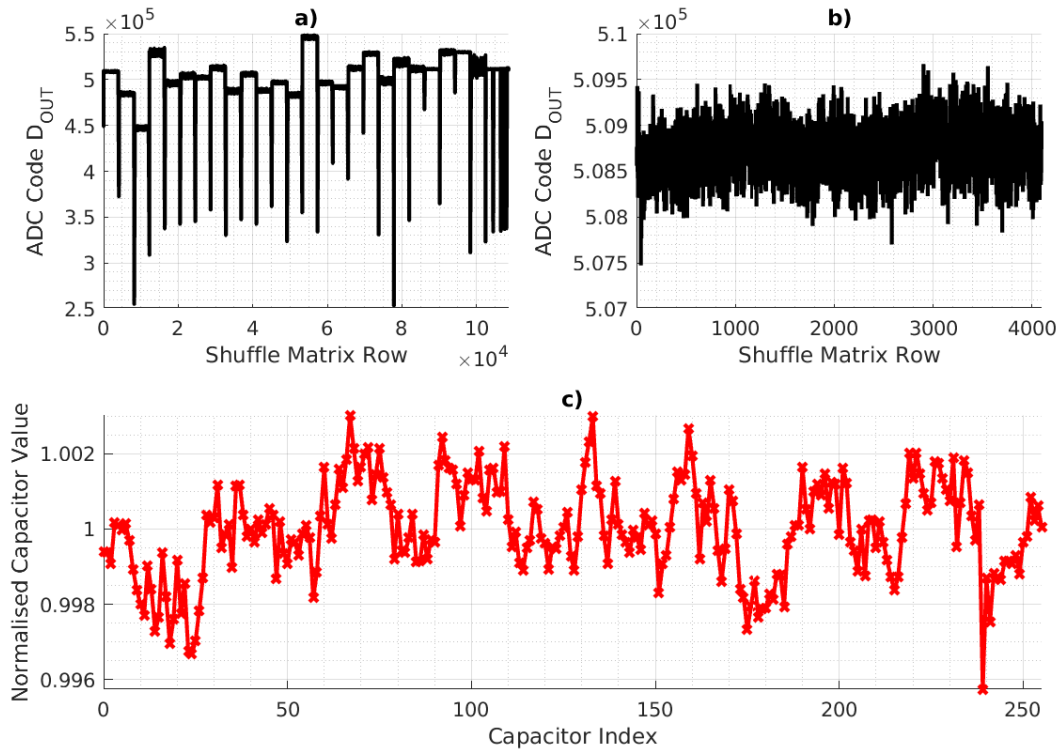


Figure 3.17. a) Raw measured ADC data from the evaluation board for all arrays on a single die. b) Magnified view of the data measured for each shuffle matrix row for a single array of capacitors. c) The extracted capacitor values from the array through linear regression.

3.6.7 Shuffle Matrix Generation

Of course, the shuffle matrix itself still has some bearing on the quality of measurement. The absolute minimum requirement of the shuffle matrix is that it has full rank such that each capacitor weighting may be individually extracted from the results. Whilst an orthogonal matrix with a directly derivable inverse is desirable for simplicity, this benefit is lost due to the presence of measurement noise in the system which would contribute directly to the capacitor weightings unless an overdetermined matrix or aggressive averaging is used to reduce.

The simplest matrix would be a diagonal one that indicates to switch only a single capacitor each trial. This directly resolves each capacitor weighting, but does not benefit from the switching of multiple capacitors at once to increase the measurement SNR, reducing noise contributions. The Hadamard matrix is often used for similar applications in the wireless transmission domain, as all rows are orthogonal, have full rank, and can be generated easily. It is a square matrix so would provide 1 trial combination for each capacitor in the array. However, the first row of the matrix would require switching of all capacitors, which is expected to be not possible due to the dynamic range of the measurement system, whilst still enabling sufficient details for obtaining mismatch error data. Also this requires switching of all capacitors for every trial, which demands careful sizing of the inverter to minimise any disturbance on the top plate which may introduce further errors. As a result, it was deemed necessary to devise a method for

generation of matrices that have the same qualities and requirements as the Hadamard matrices, but without these restrictions in order to avoid these issues, and introduce further random decorrelation of the capacitor measurements. A method for the creation of such matrices is described.

The shuffle matrix is dimensioned $N_t \times N_c$ where N_t is the total number of rows/trials, and N_c is the total number of capacitors. To construct an appropriate matrix the does not give unfair weighting towards particular capacitors, or incorporate any bias from unfairly weighted switches in the circuitry, some basic requirements are defined:

- Each trial must switch the same number of capacitors.
- Equal number of capacitors must experience a positive and negative transition.
- All capacitors must be switched an equal number of times in both directions.

The basis for the matrix construction is to construct a series of blocks of height $\frac{2N_c}{\beta}$, which satisfy the above conditions independently, and combine them into a complete matrix containing N_t rows. Construction of a single block to meet the requirements is illustrated in Figure 3.18 for $N_c = 16, \beta = 4$. To ensure each trial/row contains a unique combination, the order in which the columns are shuffled in Figure 3.18a must be unique to each block. This is easily and efficiently verified.

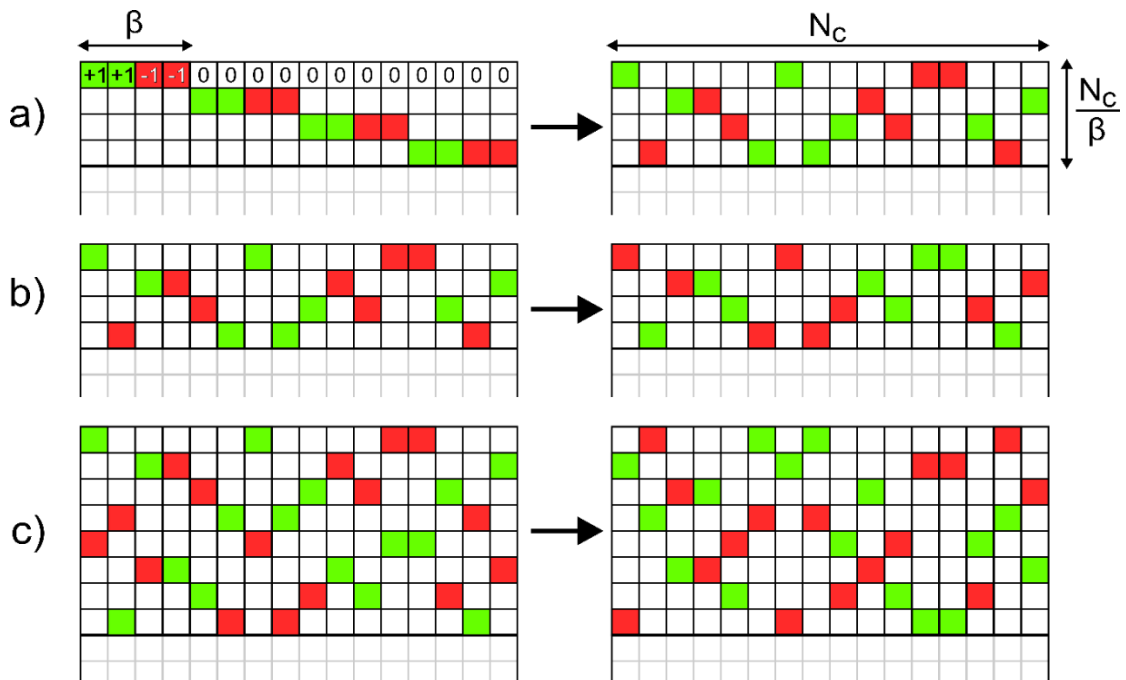


Figure 3.18. Illustration of formation of a single block of the shuffle matrix. a) The source matrix columns are randomised. b) The resulting matrix is inverted. c) The two matrices are stacked on top of each other, and the row order is randomised. This process is repeated for larger matrices to improve measurement consistency and resolution.

The above-described method prescribes some inferred requirements with respect to the sizing of the matrix which must be satisfied to ensure its construction is possible:

- N_c must be divisible by β
- N_t must be divisible by $\frac{2N_c}{\beta}$

Aside from the above-described construction method, generating purely random permutations would be the only alternative to ensure truly random trials. Due to the number required for precise measurement ($>10,000$) and ensuring proper equal capacitor weighting, this becomes a computationally costly exercise.

3.7 Setup and Evaluation of the Capacitor Measurement Die

3.7.1 Fabrication and Evaluation Architecture

To remain in line with typical commercially packaged SAR ADCs, and to ensure capture of any stresses or other dependencies that are introduced through encapsulation, all dies were packaged in a plastic encapsulated 5x5mm LFCSP. The vast majority of parts were evaluated in a socket as shown in Figure 3.21, however selected soldered devices mounted to daughterboards were also assembled to ensure all effects were considered, and that there were no other stresses introduced through the soldering process. No discernible difference in capacitor weighting was observed for pre- and post-soldered parts, so all presented measurements are for socketed parts. A photo of a fabricated die prior to packaging is shown in Figure 3.19.




Figure 3.19. Photo of a bare die from a sample of the test chip 2. The positioning of key circuit blocks is highlighted.

The packaged devices under test (DUTs) were measured in the evaluation board comprising of the basic signal chain of Figure 3.20. The 8 outputs from the DUT are multiplexed down to a single differential signal to use the same ADC for all measurements. Some amplification was provided on the board to ensure maximised use of the ADC input range, as the ADC linearity is negligible compared with that of the integrated amplifier. Some filtering is also employed to prevent kickback from the ADC during sampling, and limit the noise bandwidth. The evaluation setup is automated and controlled through LabView, interfacing with an FPGA based pattern generator.




Figure 3.20. Block diagram of the evaluation setup. The DUT outputs are multiplexed on the evaluation board for a single measurement chain comprising of amplifier, ADC, and pattern generator.

The 4-layer evaluation board was fabricated and hand assembled. The design uses near exclusively surface mount devices for the entire signal chain to reduce parasitics and ensure a compact design. Care was given to the critical paths which may be susceptible to induced noise or pickup from the digital switching of the circuitry during layout. The same was done within the test chip layout in silicon. The pattern generator mounts to the board directly as shown in Figure 3.21, whilst the rest of the supply connections to the board are delivered through SMA connectors and coax cable as a good compromise between signal integrity and size. LCR probes were not a permanent feature of the board and were connected only for initial measurement before full characterisation of each die.




Figure 3.21. Fully-assembled evaluation board for measurement of packaged die (right) on a daughter-board with LCR probes connected, and FPGA pattern generator attached (left).

3.7.2 Timing and Optimisation

Whilst the expected dynamic performance of the chip, amplifier and evaluation board design were verified ahead of fabrication to ensure measurement stability, a sweep of the settling duration t_m was performed. The period was swept from the minimum period of 20ns to 5 μ s in Figure 3.22 and the optimum value of 1 μ s for t_m was chosen for all further measurements. This ensures complete settling of the integrated amplifier and signal chain under the worst case (largest) capacitor measurements.




Figure 3.22. The duration of the measurement settling period t_m , swept from 20ns to 5 μ s to obtain optimum settling without leakage.

A 1 μ s settling time enables full settling of the circuit (including a 25% margin for the settling of the slow skewed dies), without absorption of leakage into the measurement. If the measurement period was too long the output would tend to droop linearly as the feedback capacitor of the integrated amplifier discharged due to either leakage through the dielectric or the amplifier input bias current. This discharge would be found to have a temperature dependency; however this should not affect measurements at temperature as the leakage is also proportional to the switched capacitance, and the measurements taken are always relative to the observed switched capacitance. Therefore, the leakage would only introduce a gain error that is factored for during normalisation of the results. This stands true for the comparison across geometries also.

Some optimisation sweeps were also undertaken to identify possible sources of distortion. It was noticed when using a matrix that operated 64 total capacitors in each side for the larger arrays, with a fully complementary matrix an occasional glitch occurred in the measurements. This led to the use of a maximum of 32 capacitor being utilised on each side of the array, and was believed to be due to combinations of capacitor giving an excessively negative spike at the input of the amplifier, causing the input to saturate, and resulting in a large offset in the measurement. This is likely due to the non-simultaneous switching of the P and N drivers on the bottom plate.

It was accepted that some charge injection and initial settling of the amplifier would take place during the first few moments after releasing the amplifier from the reset state. This is captured in Figure 3.23. The charge injection from release of the reset switches holds a large temperature dependency which would affect any verification over temperature. This is combatted by introducing a zero-switching trial in-between each shuffle matrix row measurement. This entails enabling the appropriate array and amplifier being measured but programming the die such that no capacitors are toggled. The result taken for each matrix row is then the difference between the switched and zero-switched measurements, which both contain the charge injection contributions, and thereby cancel this source of offset.




Figure 3.23. Output voltage at the ADC input as the measurement period is swept, with no capacitors toggled to capture the charge injection from releasing the system from the reset state.

3.7.3 Evaluation of System Noise

The noise of the evaluation system was determined through analysis of the capacitor measurements. The setup was programmed to measure the same capacitor indefinitely, and the resulting ADC codes extracted. The standard deviation of these measurements was extracted as $297\mu\text{Vrms}$. If the DUT is not operated during these measurements the noise was extracted as $148\mu\text{Vrms}$ which accounts for purely the signal chain and amplifier in reset. This is without any averaging present use which is aggressively utilised in the real measurement scheme. In practice 8192 samples are taken per measurement, and the capacitors experience an equivalent 256x oversampling due to the shuffle matrix measurement strategy.

To quantify the impact of the measurement system noise on the extracted capacitor values, and thereby the extracted capacitor statistics, the same shuffle matrix is run on the same array repeatedly (128 times). The standard deviation of the extracted capacitor errors for a single capacitor is then performed to give the effective noise of the extracted capacitor values. For the case of 32 differentially actuated capacitors per trial, 4096 trials/combinations, and 8192 averaged samples per trial, a σ of 30ppm was observed. These were the conditions under which mass measurement of the test-chips was performed. As will be seen when presenting the matching statistics of the capacitor arrays, this value does not meaningfully impact the extracted statistics. In this context, a nominal 12ppm in the matching statistics was measured and is discussed later. This is much smaller than the discrepancies that will be analysed for identification of matching trends and also will be further reduced once multiple dies are measured. This is smaller in magnitude than the variations in measurements due to the die fabrication so is considered acceptable. It is accepted that there will be some variation in error based on positioning of the array across the die due to packaging and fabrication [48], [49]. In one instance, an additional 100ppm in systematic error for identical arrays placed on opposite corners of the die (~ 1 mm separation) was seen. The discrepancy in random errors was only 17ppm however. The effect was severely reduced (< 10 ppm) for arrays located adjacent to each other, which is the case for the majority of the variations that are discussed. This can be seen in Figure 3.24.




Figure 3.24. Illustration of repeatability of 8 identical runs of extraction of the values of a single array of capacitors, achieving 30ppm σ in extracted value error.

3.8 Conclusion

An overview of the 40nm CMOS process has been provided, and the many challenges that come with designing in a low-geometry process designed primarily for digital circuit performance are discussed. The subsequent issues with the implementation of the common lateral coupling MOM capacitor are discussed, and the irregularity of the structure identified as a likely source of matching errors contributing to the capacitor structure. Their larger scale impact on large capacitor arrays is also considered. A charge-based measurement system is developed to enable measurement of an optimised capacitor geometry to ppm level using a decorrelated measurement strategy, enabling extraction of the capacitor values and their relative weightings. This design is implemented and refined, and the system performance evaluated with the extraction errors being sufficiently low for the identification of trends and relationships between the capacitor geometry and the subsequent matching performance. This enables the further chapters to address the research question's intention to single out the main contributors and quantify what improvements may be achieved.

Chapter 4 Measurement and Analysis of Single Finger Lateral Coupling Capacitors

4.1 Introduction

Within this chapter the relationships between the capacitor finger spacing, width, length, number of fingers, and the resulting capacitance are established and verified, using the test structures and methodology developed in Chapter 3. Their relation to the parallel plate capacitor equation is considered, and a significant fringing field contribution to the capacitor is identified due to the small geometry and high aspect ratios of the structures. Process-related elements which complicate the relationships observed will be discussed and used as a foundation in further chapters regarding the likely formation of observed errors and other relationships.

4.2 Capacitor Weighting Analysis

To recap for the purpose of analysis, this implementation of the laterally coupling, single finger MOM capacitor utilises the standard extra-low-k (ELK) intermetal dielectric of the foundry's BEOL process, with the precise dielectric constant k value unknown. The capacitor implementation relies on the sidewalls of the metal interconnects to produce the primary capacitance, with the pattern duplicated across N_m metal layers to increase capacitance. This design does not directly rely on any vertical coupling between the metals. In addition to this direct coupling, a fringe field is also formed. The top-view in Figure 4.1 also shows that the coupling at the finger ends must also be considered, as it contributes to the overall capacitance. Additional contributions may be split into two sections, the end and corner capacitance, as one end of the capacitor finger only consists of one contained corner contribution so must be stated separately.




Figure 4.1. Contributors from the capacitor's geometry to the total capacitance.

4.2.1 Capacitance Against Finger Length and Number of Metals

The basic understanding of how the capacitor value scales with the various geometry is described by Equation 4.1. It assumes that the capacitance is approximately a parallel plate capacitor, though with a significant level of fringing contribution α , which to first order is proportional to the parallel plate capacitance. This equation collates all coupling at the end and corner of the fingers, and the coupling between metal layers in a single term C_D .

$$C = C_U(1 + \alpha) + C_D = \epsilon_0 \frac{tL}{S} N_m k(1 + \alpha) + C_D \quad 4.1$$

Figure 4.2 and Figure 4.3 show the measured capacitor values versus the finger length and number of metal layers, respectively. All data points are measured multiple times and averaged over 100 identical capacitors on different locations on the die and on different dies. The error introduced by this variation is less than 0.1% and the size of the data points is larger than the error. Data from a linear fit as per Equation 4.1 are extracted and displayed in Table 4.1. In both cases, the quality of fit R^2 is larger than 0.99 indicating an excellent linear fit to the data.

The intercept in the case of finger length variation gives $C_D=1.0$ fF, a value which is reasonable considering the ratio of length to width for the fingers. In case of the number of metal layers, consideration of the diagonal coupling on its own would suggest an above unity increase in capacitance for each additional metal used, and therefore a negative theoretical intercept to the plot. Instead, a positive intercept was observed. This is best explained by the large fringe capacitance contributions. The additional metal layers constrict the space in which the fringe fields act (per layer) and therefore decrease the net contribution of the fringing fields. The positive intercept corresponds to a situation in which the relative decrease in fringing is larger than the contribution from the diagonal coupling.

The linear fit in Figure 4.2 and Figure 4.3 was used to extract the contribution of the fringing factor α . For this, an assumption for the metal thickness of $t=144\text{nm}$ is taken [54]. The precise value for intermetal dielectric constant in the process is also unknown, it is believed that since the process is advertised as an ELK process that the value is approximately $k = 2.5$, which is low compared to the range of classical SiO_2 dielectrics of 3.9 [29], [34].

This results in the estimated value of α of around 1.0 ± 0.1 . This is a significant factor relative to the classical large parallel plate capacitor ($\alpha \ll 1$), and implies strong fringe field effects which contribute half the total measured capacitance value. Although a parallel plate capacitor has an exact solution and some formulas exist for small fringing fields [55], these approximations are not appropriate for such strong fringing field. In extremis, the capacitor value would be identical to two parallel wires, but we are not in that regime either.




Figure 4.2. Relationship between the measured capacitance and the capacitor finger length. $S=90\text{nm}$, $W=70\text{nm}$, $N_m = 4$.

Table 4.1. Parameters for the linear fit illustrated in Figure 4.2.

Parameter	L Sweep	N_m Sweep	Unit
R^2	0.9994	0.9969	-
Gradient	0.5912 fF/ μm	5.2232 fF	-
Intercept	0.9635	1.96	fF
k	2.5 [34]		-
t	144 [54]		nm
S	90	100	nm
W	70	90	nm
L	var	43.5	μm
N_m	4	var	-
$k(1 + \alpha)$	5.2	4.7	-
α	1.1	0.9	-

COMSOL finite element simulations were performed in order to verify a reasonable order of contributions from the fringe field capacitance for a variety of spacing-to-thickness ratios. Figure 4.4 shows these simulation results plotted against the ratio of the finger spacing to the metal thickness. The capacitors have spacing of 70-90nm and an estimated thickness of 144nm which gives an approximate spacing to thickness ratio of around 0.5-0.6 which corresponds indeed to a p value of α of around 1.0. This shows that the design process for these capacitors is well understood.

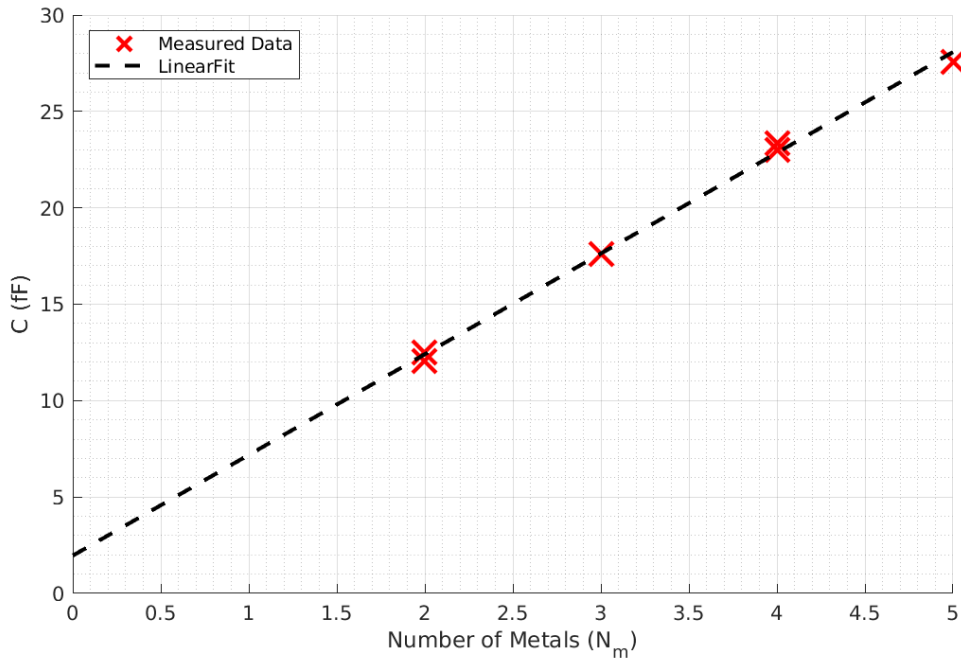


Figure 4.3. Relationship between the measured capacitance value and the number of metal layers.
 $S=100\text{nm}$, $W=90\text{nm}$, $L=43.5\mu\text{m}$




Figure 4.4. Extracted α contributions from COMSOL simulations of a parallel plate capacitor.

4.2.2 Capacitance Against Finger Spacing and Width

The basic assumption from equation 4.1 leads to an inverse relationship between the finger spacing and the resultant capacitance. A plot of this nature is shown in the data of Figure 4.5, with a significant intercept offset found and documented in Table 4.2. The increased finger

spacing leads to a decrease in overall capacitance, although the quality of the overall fit shown makes it clear this is not the true relationship between these parameters.




Figure 4.5. The relationship between the measured capacitance and the capacitor finger spacing. $W=70\text{nm}$, $L=14.5\mu\text{m}$, $N_m=4$.

Table 4.2. Extracted parameters for the linear fit observed in Figure 4.5.

Parameter	Value	Unit
R^2	0.984	-
Gradient	0.257	pF·nm
Intercept	6.5	fF

Further analysis and fitting of a log-log plot to the same data in Figure 4.6 illustrates that the $1/S$ relationship does not realistically correlate with the measured data. The trendline illustrates a $\frac{1}{S^{0.3}}$ relationship suggesting a significantly smaller effect on capacitance than expected due to the change in finger spacing. This is assumed to be due to the total capacitance being comprised of significant fringing effects, which complicate the interaction between the usual capacitor geometry parameters such as the spacing, and the total measured capacitance.




Figure 4.6. A log-log plot of the capacitance against capacitor finger spacing. $W=70\text{nm}$, $L=14.5\mu\text{m}$, $N_m=4$

Now the impact of the finger width is considered. The unit value for a sweep of the finger width of the capacitor, for 70 and 90nm finger spacing, was measured and is shown plotted in Figure 4.7. A linear decrease in measured capacitance is shown with an increase in finger width, with a larger degradation in capacitance observed for capacitors of smaller finger spacing. The direct capacitance between the two parallel plates should intuitively have no dependency on the finger width. It is probable that the reduction in capacitance here with increasing width is due to a larger scale effect dependent on the localised density of the metals. This is a known effect in processes below 65nm derived from their common BEOL processing, and specifically for their use of the copper damascene process and CMP. This effect also contributes to the results observed with an increasing of the finger spacing, and hence the respective reduction in localised metal density that comes with it. Without further characterisation of CMP dishing across similar sized metal tracks that are used in these capacitors, it is difficult to confirm this is the case, however such investigation is beyond the scope of this thesis.




Figure 4.7. Relationship between the measured capacitance and the capacitor finger width. $L=14.5\mu\text{m}$, $N_m=4$.

Table 4.3. Parameters for the linear fit in Figure 4.7.

Parameter	Value (S=70nm)	Value (S=90nm)	Unit
R^2	0.999	0.990	-
Gradient	33.3	19.7	aF/nm
Intercept	11.15	11.59	fF

These results were verified through further COMSOL simulations. The results of a sweep of the capacitor finger width are shown in Figure 4.8. The effect of the concentration of field lines is shown in Figure 4.9. The trend observed in simulation shows an increase in effective capacitance with an increase in the capacitor finger width. This agrees with theory suggesting that the increase in finger width leads to significant increases in fringing of the capacitor. The increase in fringing in this case leads to the multi-layer capacitor becoming a better approximation of the parallel plate capacitor. This increase in field line density can be seen in Figure 4.9 where the field lines are found to extend up as far as to meet those of the adjacent layer. There is notable fringing field visible above and below the overall capacitor structure. This supports the previously mentioned theory that the capacitor layers will interfere with each other

and could account for the offset previously mentioned for capacitors not having >2 fingers, with the inter-layer capacitance providing this offset.




Figure 4.8. A COMSOL simulation sweep of the capacitor finger width. This shows an inverse correlation to the trend of the test-chip measurements.

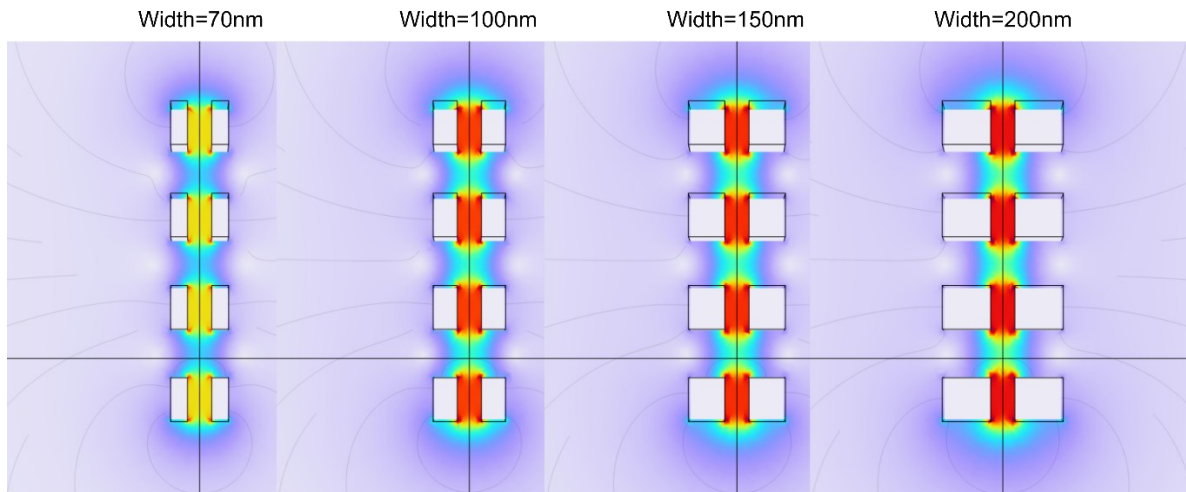


Figure 4.9. Simulation of electric field lines for a range of capacitor finger widths from 70nm to 200nm. An increase in fringing is observed for larger metal widths.

As the simulation data trends not in agreement with the measurement data, this points to another currently uncharacterised effect in this process, that is leading to this discrepancy. The effect must be significant enough to overcome the original trend twofold based on the data available. It is well known that the smaller geometry processes suffer from poor vertical

definition, leading to variations in thickness and etch depth across the wafer. It is quite possible that the localised metal thickness is varying with the localised metal density due to two of the steps used in its formation as this is a frequently documented observation [56], [57], [58], [59], [60], [61]. Firstly, aspect ratio dependent etching (ARDE) effects may be at play, producing a shallower etch depth with increased finger width due to excessive precursor which is introduced to prevent the natural trapezoidal etch. CMP effects are also commonly documented in the copper damascene process [61], [62], although they would be more likely to introduce a positive dependency with the finger width if the copper is assumed to be harder than the surrounding intermetal dielectric, although this assumption is in hindsight baseless in the context of this particular process. This is a variation on the typical description of ‘dishing’, where metals exhibit over-polishing compared with the surrounding geometry [60], [61]. While it may be expected that the foundry controls the thickness to a high degree, their control across a variety of densities cannot be expected.

Clearly, as the spacing is increased the fringing contribution also increases, so whilst the direct coupling field may decrease, the $1/S$ reduction in this is decreased by a proportional increase in the lateral coupling. A similar principle is applied to the relation with the finger width. It is expected that moving away from the minimum geometry in these parameters would lead to a less dependent fringing field, and there a capacitor that behaves more conventionally.

4.3 Conclusion

Relationships between the key capacitor geometry parameters have been analysed and have revealed significant fringing contribution within the lateral coupling capacitor in a 40nm process, approximately equal to the parallel plate contribution. This contribution yields a less than unity increase in capacitance value with an increased number of metal layers. The capacitance has proportional dependency on the capacitor finger length. A non-obvious relationship is identified regarding the width and spacing of the metal fingers to the capacitor value which is not explained by the usual capacitor theory, nor proved by basic simulation. This points to processing density and linewidth dependent effects in the etching and polishing stages, yielding a variation in metal thickness. This contributes to the knowledge of likely contributors to the matching variations that will be discussed in the following chapters.

Chapter 5 Detailed Analysis of Capacitor Errors and Trends

5.1 Introduction

Full calibration of an ADC leads to significant production-test cost during manufacture of a high-precision ADC. By providing a better understanding of the very nature of these errors, either test cost can be reduced through a more limited calibration scheme, or improved matching can be attained through informed design such that some tests become unnecessary. This understanding is achieved through independent analysis of the systematic (repeatable) and randomly sourced errors in the capacitor arrays, and their dependencies on the capacitor geometry: Finger spacing, width, end spacing, metal combinations, and number of fingers. Understanding the nature of the matching errors on top of quantifying them is a key goal of this project.

In this chapter, a theory is constructed for the analysis of both dominant mismatch error types: systematic and random. This theory is then applied to the previously described capacitor array measurements. Several key and novel relationships between the capacitor geometry and other layout parameters are identified. Discussions focus predominantly on the presence of large systematic elements observed in errors across 174 measured dies, as well as unexpected relationships between the matching, and seemingly inconsequential layout parameters such as the finger width. Significant portions of the capacitor errors are found to be repeatable across multiple wafers and devices, allowing the possibility of implementation of blind trims, or other more limited, calibration schemes.

5.2 Derivation of Matching Statistics from Array Measurements

5.2.1 First Presentation of Array Capacitor Weighting Measurements

A representative example of the normalised extracted capacitor weightings for four different capacitor designs is shown in Figure 5.1. Each design is measured on many different dies but for clarity only 3 randomly selected die are shown. It is clear in all four instances, that a large component of the errors is consistent across the measured dies. This is also observed to be the case when looking at arrays from the other measured dies. This systematic error has been previously discussed in various studies although in those cases only large scale linear variations were observed [42], [49], [50], [63], [64]. None of those studies provided similar resolution of the

measurement error. It is evident that at the resolution of our measurement, many different systematic phenomena can be observed which require more sophisticated analysis. To produce a single figure and thereby allow comparison of the errors of this nature across geometries, the standard distribution of the average error is considered.

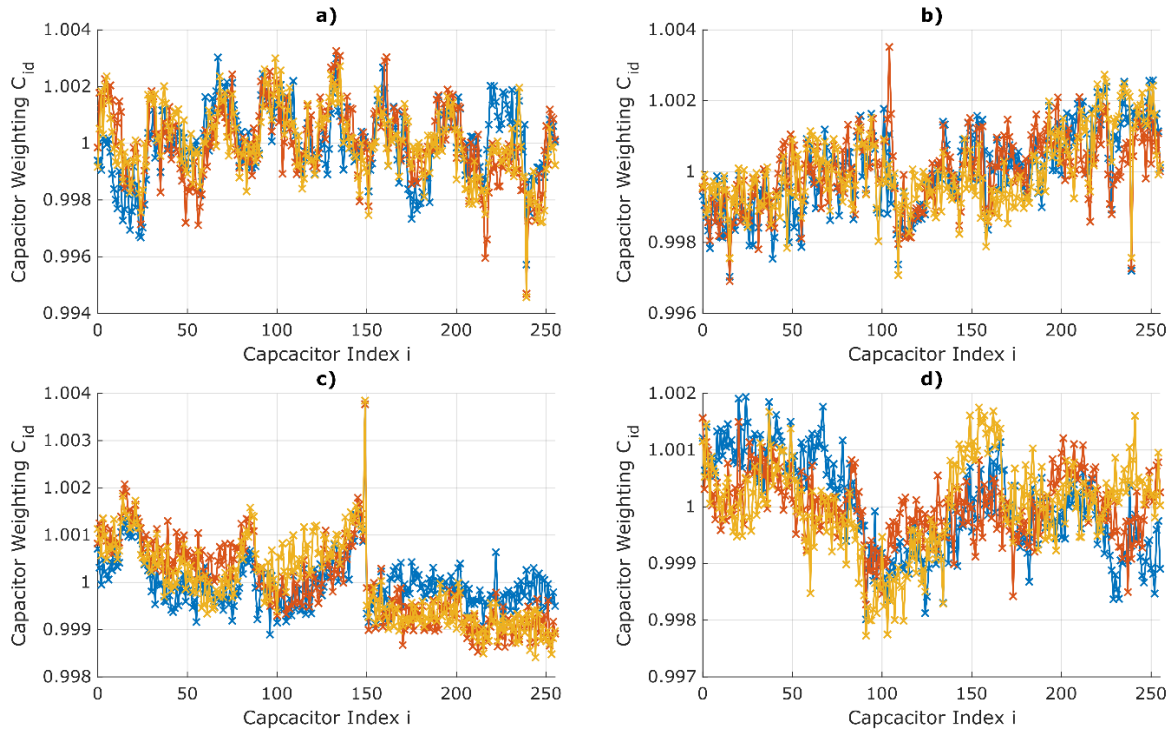


Figure 5.1. Examples of three measured arrays of four different capacitor designs, illustrating common errors for a single array with respect to position within the array. Yellow, blue and red represent data for three different measured dies.

Averaging of the normalised capacitor weighting is performed to extract the average capacitor values by the index seen in Figure 5.2 (top right). This data allows the analysis of the systematic components of various nature, such as sinusoidal variation, step profiles, single spikes, and linear gradient. These average capacitor errors are then subtracted from the measured weightings for each die. The remaining errors, also shown in Figure 5.2 (bottom right), appear random in nature. Note though that since capacitor allocation commonly involves some level of folding or interdigitation, the indexation of the capacitors results in some lost knowledge on the nature of the systematic component. Also, no wafer map information is present. This means that, whilst there may be some contributing systematic sources to these errors, there are insufficient dimensions in the recorded data to analyse them further and they are treated as random gaussian errors for the purposes of this analysis.

It is by independent analysis of the systematic and randomly contributing errors in this chapter that the value of this research may be realised. It then becomes possible to understand the impact of blind calibrations of the predictable errors, such that test time, and other calibration

efforts are thereby reduced, and the product of other linearisation methods may be further enhanced.

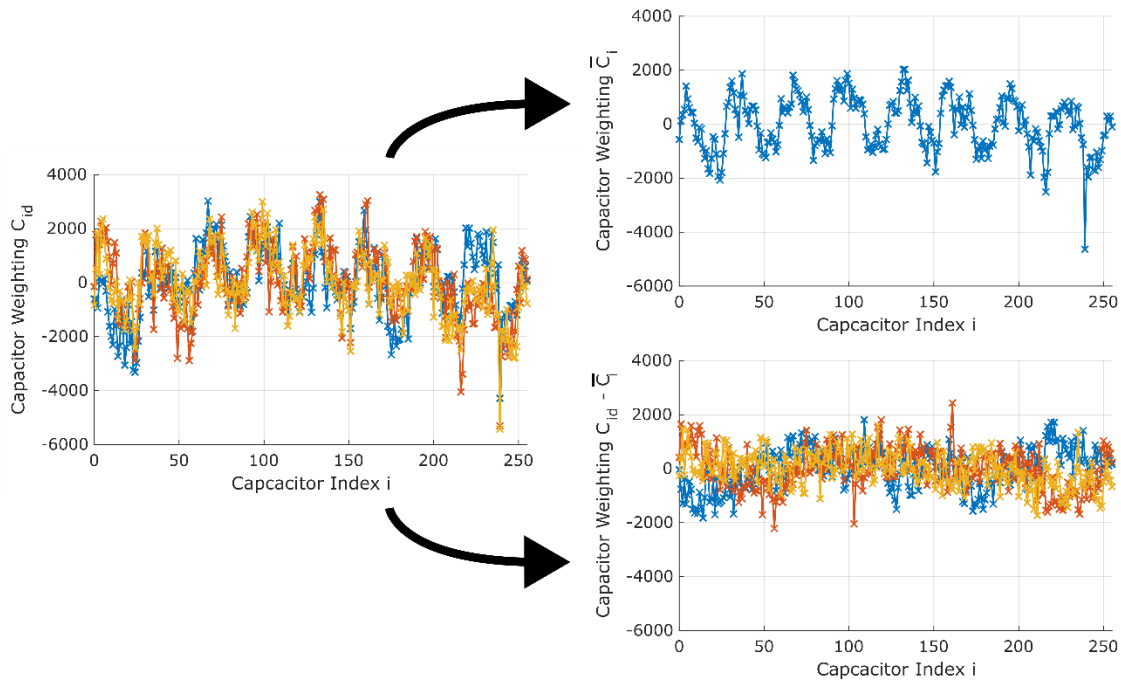


Figure 5.2. The measured capacitor errors (left) may be extracted into two components: A systematic component common to all measured dies (top right), and a random component (bottom right).

5.2.2 Systematic and Random Matching Definitions

Before any statistic is derived, the measured weightings for each capacitor are normalised to the die mean. This enables comparison of the errors with each variation against one-another.

Firstly, a statistic to represent the systematic errors is determined. $C_{i,d}$ represents the measured value of the capacitor at location i of N_c capacitors, on die d , where there are a total of N_d measured dies. The average weighting at each array index \bar{C}_i is then calculated in equation 5.1.

$$\bar{C}_i = \frac{\sum_{d=1}^{N_d} C_{i,d}}{N_d} \quad 5.1$$

The subsequent array contains the errors which contribute to all die, including those resulting from specific design geometry and deterministic stresses and die gradients. To facilitate comparison between geometries, the standard deviation of \bar{C}_i is taken, and defined as the systematic mismatch σ_S for the specific geometry in equation 5.3, using the normalised array mean \bar{C} in equation 5.2. This statistic is used as an approximation at this stage to represent the

systematic errors present on all dies, although the validity of this approximation by standard deviation is later verified.

$$\bar{C} = \frac{\sum_{i=1}^{N_c} \sum_{d=1}^{N_d} C_{i,d}}{N_c} = 1 \quad 5.2$$

$$\sigma_S = \sqrt{\frac{\sum_{i=1}^{N_c} (\bar{C}_i - \bar{C})^2}{N_c - 1}} \quad 5.3$$

Next the random error contributions must be analysed. These errors, in all cases, are determined through removal of the previously identified systematic errors \bar{C}_i from the measured errors $C_{i,d}$ to give the array random errors $R_{i,d}$ in equation 5.4.

$$R_{i,d} = C_{i,d} - \bar{C}_i \quad 5.4$$

The array random error standard distribution σ_{Rd} is calculated in equation 5.6 using 5.5, although the mean random error for each array \bar{R}_d can be assumed to be zero as it is removed through earlier normalisation.

$$\bar{R}_d = \frac{\sum_{i=1}^{N_c} R_{i,d}}{N_c} \quad 5.5$$

$$\sigma_{Rd} = \sqrt{\frac{\sum_{i=1}^{N_c} (R_{i,d} - \bar{R}_d)^2}{N_c - 1}} \quad 5.6$$

Taking the root-mean-square (RMS) of the array random errors σ_{Rd} over all dies yields the final figure for representing the random error distribution σ_R in equation 5.7. Whilst these errors cannot be confirmed to be random in nature, they are for all intents and purposes in the context of this analysis random.

$$\sigma_R = \sqrt{\frac{\sum_{d=1}^{N_d} \sigma_{R,d}^2}{N_d}} \quad 5.7$$

Further to these primary statistics, it is also valuable to consider the random errors experienced at different positions within the array σ_i . Again, the standard deviation is calculated in equation 5.8.

$$\sigma_i = \sqrt{\frac{\sum_{d=1}^{N_d} (C_{i,d} - \bar{C}_i)^2}{N_d - 1}} \quad 5.8$$

A final generalised statistic must be considered, and is believed to be the usual choice for statistical representation of capacitor errors in the literature. This is the total array error standard deviation σ_T and does not distinguish between the random and systematic error contributions. Similarly to the calculation for the random error, the standard deviation for a single die σ_d is first calculated in equation 5.10 with respect to the array mean C_d in equation 5.9. The RMS is then taken over all measured arrays to build the total error standard deviation σ_T in equation 5.11.

$$C_d = \frac{\sum_{i=1}^{N_c} C_{i,d}}{N_c} \quad 5.9$$

$$\sigma_d = \sqrt{\frac{\sum_{i=1}^{N_c} (C_{i,d} - C_d)^2}{N_c - 1}} \quad 5.10$$

$$\sigma_T = \sqrt{\frac{\sum_{d=1}^{N_d} \sigma_d^2}{N_d}} \quad 5.11$$

Note that all presented statistics are normalised to the array mean and are all tabulated in Table 5.1. This is because the DAC designer has interest only in the proportional matching of the capacitors, and there is no benefit in analysis to be gained from the absolute errors in value for which there is no control to be gained over, and which typical variations are on the order of 25%. The systematic and random contributions part in the capacitor array values are illustrated in Figure 5.3.




Figure 5.3. Illustration of the two independent types of mismatch errors that contribute to a capacitor's total value.

Table 5.1. Variable definitions relating to the description of capacitor matching and respective errors.

Variable	Description
$C_{i,d}$	Measured capacitor weighting at index i on die d
\overline{C}_i	Mean capacitor weighting at index i (normalised to be 1)
\overline{C}	Mean array weighting over all arrays (normalised to be 1)
\overline{C}_d	Mean array weighting (normalised to be 1)
$R_{i,d}$	Random component of capacitor weighting at index i on die d
\overline{R}_d	Mean capacitor random error weighting on die d
σ_S	Capacitor systematic error standard deviation
$\sigma_{R,d}$	Capacitor random die-to-die error standard deviation on die d
σ_R	Capacitor random die-to-die error standard deviation
σ_i	Capacitor die-to-die error standard deviation at index i
σ_T	Capacitor total error standard deviation

5.2.3 Relative Matching Statistic Definitions

In addition to the systematic, random, and total error statistics that are now defined, the relative matching between capacitors must also be considered. This is previously discussed in [49] where an unusual relationship in a similar process is presented, and attributed to the presence of correlated errors in the fabricated die. These errors act over a finite area, and so the variation of the relative errors between capacitors with respect to their separation must also be considered in tandem with the existing matching statistics. The insights gained through analysis of the relative matching are primarily relevant when considering folding and grouping of the capacitors within large DACs or arrays, in order to reduce the observed matching.

For each index of capacitor separation x and each die d , the error between the capacitor $\epsilon_{i,d,x}$ at starting index i , and index $(i + x)$ is taken in equation 5.12. This is performed on all available dies, and for all starting indices up to $i = (N_c - x)$. For clarity this is visualised in Figure 5.4.

$$\epsilon_{i,d,x} = c_{(i+1),d} - c_{i,d} \quad 5.12$$




Figure 5.4. Visualisation of the relative errors of capacitors within the same array separated by x fingers, on die d .

The standard deviation for a given starting index over all dies $\sigma_{i,x}$ is calculated in equation 5.14, with the systematic error $\overline{\epsilon_{i,x}}$ contained in equation 5.13 rejected, so as to only consider the random errors between the capacitors. It is well established that the systematic offset between capacitors may be lower if they are co-located, so this statistic is constructed to establish if this is also the case for the apparently randomly sourced errors.

$$\overline{\epsilon_{i,x}} = \frac{\sum_{d=1}^{N_d} \epsilon_{i,d,x}}{N_d} \quad 5.13$$

$$\sigma_{i,x} = \sqrt{\frac{\sum_{d=1}^{N_d} (\epsilon_{i,d,x} - \overline{\epsilon_{i,x}})^2}{N_d - 1}} \quad 5.14$$

The RMS is computed across $\sigma_{i,x}$ for all starting positions to obtain the final statistic for the relative matching of the capacitors at a given separation σ_x in equation 5.15. Therefore σ_x represents a relationship between the random matching between two capacitors, and the distance between them.

$$\sigma_x = \sqrt{\frac{\sum_{i=1}^{N_c-x} \sigma_{i,x}^2}{N_c - x}} \quad 5.15$$

It may be useful to also consider the relative matching including the systematic offset between capacitors. This is more relevant for larger separation distances and can allow direct comparison between this work and the relevant literature. Equation 5.16 captures the systematic errors as well as the random variations previously considered into the total relative error mean $\overline{\epsilon_x}$.

$$\overline{\epsilon_x} = \sqrt{\frac{\sum_{i=1}^{N_c-x} \overline{\epsilon_{i,x}}^2}{N_c - x}} \quad 5.16$$

5.3 Verification of the Gaussian Statistical Model for Matching Errors

With the variables for statistical analysis now defined, they may now be applied to the measured capacitor data. Before any conclusions are drawn however, the basic assumptions made in the development of the matching error statistics are tested. Specifically, this refers to the assumption that the errors, of all types, may be suitably represented by a normal distribution. It was noted that the simulator used for extraction of the capacitor statistics from the foundry model artificially limited the variations to a maximum $\pm 3\sigma$ range (99.7% coverage) during monte-carlo simulations. Whether this is a result of the PDK or the simulator itself is unclear, but whether this is a fair result is discussed within the following.

5.3.1 Systematic Error Distribution

A Q-Q plot is firstly produced in Figure 5.5 a), showing the quartiles of a standard normal distribution, against that of the extracted systematic capacitor errors. A close alignment is observed between the two distributions, illustrated with reference to the ideal normal trendline in red. It is also observed in the histogram plot in Figure 5.5 b) that all systematic capacitor errors are within the 3σ boundaries previously assumed.

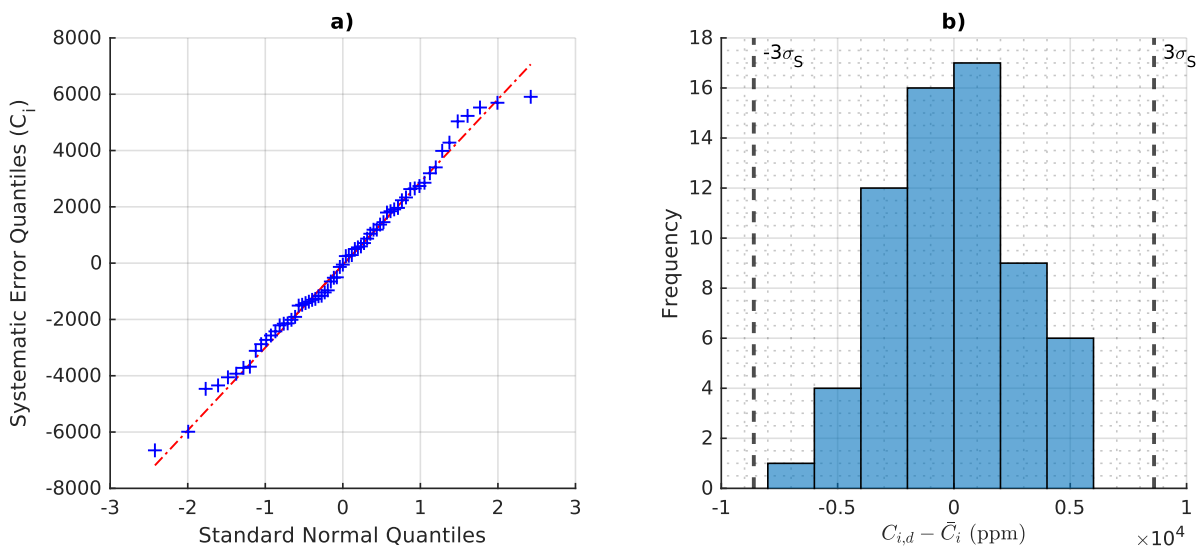


Figure 5.5. a) Q-Q and b) Histogram plots of the distribution of systematic errors of foundry cell capacitor measurements from test chip 1.

5.3.2 Random Error Distribution

Next the random errors are considered in isolation. The Q-Q plot in Figure 5.6 a), whilst showing reasonable alignment with the central portion of the distribution ($\sim \pm 1.5\sigma$), deviates significantly from the gaussian distribution outside this region. The resulting distribution observed in the histogram of Figure 5.6 b) is symmetric, but with larger content at the ‘tails’

of the distribution than a gaussian or normal distribution would otherwise predict. The result of representing this observed error distribution with an ideal gaussian distribution, is the under-representation of the larger errors that may occur during fabrication, and the over-representation of the errors closer to nominal ($< 1.5\sigma$). The limit of the foundry distribution model of 3σ on the distribution is reasonable as the spread of the distribution does not meaningfully extend beyond the 3σ range, with $<0.5\%$ of errors beyond this range as expected.

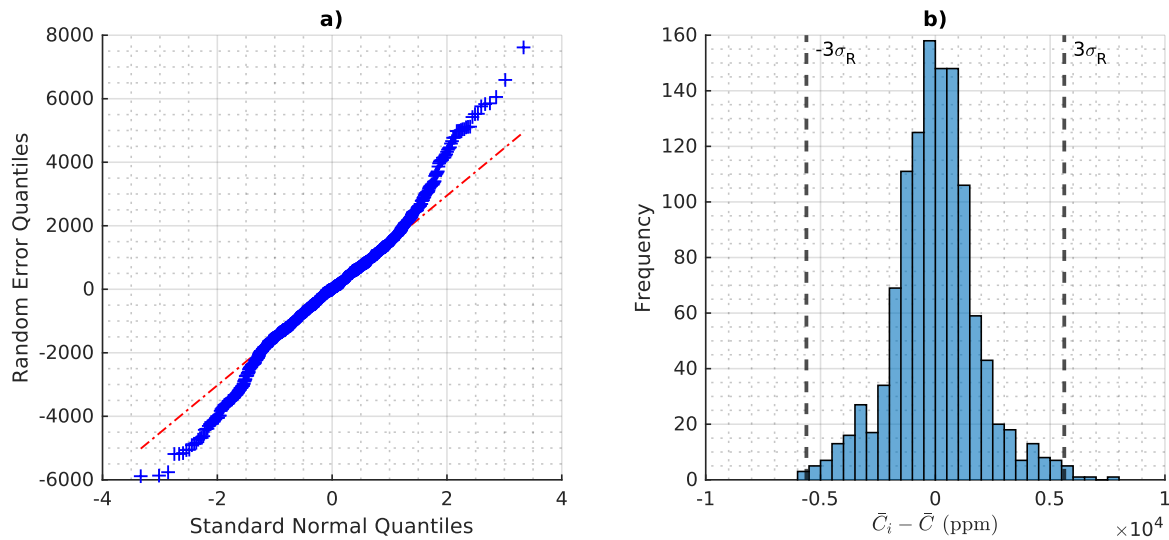


Figure 5.6. a) Q-Q and b) Histogram plots of the random error distribution of foundry cell capacitor measurements obtained from test chip 1.

5.3.3 Total Error Distribution

We now finally consider the total error contributions, such as are considered in foundry models for the simulation models previously used. Figure 5.7 a) shows a similar trend to the random error distribution data. A good alignment with the gaussian distribution trendline between $\pm 1.5\sigma$, but a reasonable deviation beyond this range. Some level of asymmetry of the distribution is seen in the Q-Q plot that is also observable in the histogram of Figure 5.7 b). The maximum deviation is equivalent to 3.8σ , but this accounts for less than 1% of errors. Approximately 85% of the sampled errors are represented well by the distribution.

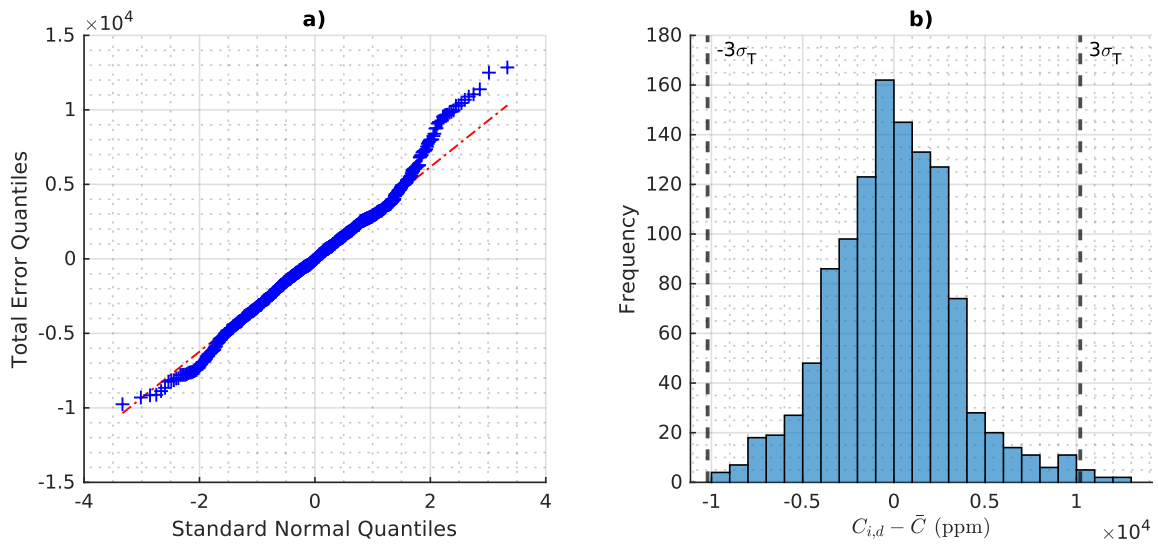


Figure 5.7. a) Q-Q and b) Histogram plots for the total capacitor errors for the foundry cell capacitors measurements from test chip 1.

Overall, representation of mismatch errors for the foundry capacitors are suitably performed by a gaussian distribution. The errors extracted for both the systematic and random based errors fit the distribution well, although an underestimation of the errors is expected with respect to the random matching. Considering the typical use case of this error model, corner and monte-carlo analysis, the fit is quite sufficient and does not warrant further analysis.

5.4 Analysis of Observed Systematic Errors

Focus is now brought to the systematic errors on the die. Referring back to section 5.2, these errors are created through averaging of the normalised errors for each capacitor index over all 96 measured dies, and the 6 measured wafers. There are some exceptions to this relating to dependency on skew, however these are highlighted where appropriate.

5.4.1 Verification of Large Periodic Error Content

As previously discussed, there exists a set of large systematic errors dependent on the specific capacitor geometry. The most pronounced of these is for a capacitor of minimum finger spacing and width. This error is illustrated in measurements of the foundry cell capacitors in Figure 5.8 on 5 measured dies. An overall gradient is observed, correlating to 60 ppm/ μm , leading to an approximate offset of 4000 ppm between capacitors at each end of the array. This gradient is typically referred to by the literature as process gradient and is normally thought to be the dominant contributor to the mismatch [31], [36], [42], [48], [49], [50], [64], [65], [66]. On top of this linear gradient, a significant periodic component is observed. This pattern approximates to a sinusoid of 2300ppm with 8.95 μm period, and is unrelated to any multiple of the capacitor

pitch of $1.54\mu\text{m}$. A large deviation is also observed at around $10\mu\text{m}$ along the array. This was consistent across all measured devices.

The magnitude of the systematic component observed in the arrays, is much larger than could be expected, yielding a figure of 2849 ppm where all capacitors are considered. If the error at the last $15\mu\text{m}$ of the array are removed as they are presumed due to a lack of dummy capacitors at the end of the array, then a reduced error figure of 2517 ppm for the systematic matching error is calculated. This however only equates to a 10% improvement. A periodic error is previously undocumented beyond this work, so further measures were taken to decorrelate its presence from any layout or measurement derived sources.




Figure 5.8. Systematic error capacitor by index, showing a strong systematic profile of a sample of 5 measured dies, from arrays using only foundry capacitor geometry.

While the measurement method has already been evaluated in previous sections, a brief revisit to verify this content is required. The periodic component is observed on all four arrays of each die in Figure 5.9, for both vertically and horizontally orientated arrays, also for the custom-designed minimum geometry arrays. The period in all cases is extracted as $5.84\mu\text{m}$ through linear regression, and is observed to align for each pair of capacitor arrays as in Figure 5.9. The discrepancy between the $5.84\mu\text{m}$ extracted period for the custom structures compared with the foundry structures $8.95\mu\text{m}$ period is down to the foundry capacitor consisting of 5 capacitor fingers per measured unit capacitor. This error is effectively the averaging of those 5 capacitors fingers. The deviations between the horizontal and vertical orientated arrays at $10\mu\text{m}$ are due to an irregularity that will be discussed further in section 1.1.1. This irregularity is later determined independent of the periodicity observed here so may be ignored for this discussion.




Figure 5.9. Mean capacitor errors by capacitor position, for each measured minimum geometry array. The errors are consistent in period across the four arrays.

Possible sources of the observed capacitor errors are leakage of the capacitors or drivers during the measurement period, or a localised effect stemming from an artifact of their layout. If this error is due to leakage, this would yield variations in measured error with respect to changes in supply voltage, settling time, and most significantly temperature. This is verified first through variation of the clock speed and reference supply voltage to the DUT in Figure 5.10. No noteworthy differences in the relative extracted errors were observed. The magnitude of any variations were below the measurement noise (30ppm).




Figure 5.10. Measured capacitor errors for a single array with varied settling times through changing of the clock speed. Supply voltage variations are also introduced to decorrelate the matching errors from any settling errors.

A temperature forcer is employed for operation of the chip under different temperature conditions as leakage has a strong correlation with temperature. The results from measurement of the array at -40 , 0 , 25 , 65 , and 90°C in Figure 5.11 showed no net movement in the normalised capacitor error due to temperature. A drift of $0.015\%/^\circ\text{C}$ was observed in the nominal capacitor value however, although the setup was not calibrated for the drift of the amplifier and surrounding circuitry, hence this value is an order of magnitude higher than would be typically expected. This confirms that the observed systematic profile is highly unlikely to be related to the device leakage, whether this be the bottom-plate drivers, or the capacitors themselves. The errors are also shown decorrelated with the permutation of shuffle matrix used in Figure 5.12. The anomalies at capacitor indices 15 and 63 are attributed to the use of the Hadamard matrix which was established to introduce large artefacts at certain capacitor indices and are provided here only for illustration.




Figure 5.11. Extracted capacitor mismatch errors for the minimum geometry capacitor under different temperature conditions, resulting in limited variation.




Figure 5.12. Extracted capacitor errors from 7 difference shuffle matrices to establish decorrelation of the observed periodicity from the measurement order and matrix.

The specific layout of the capacitors and their drivers were investigated to ensure that there was no surrounding geometry which may contribute as an error source to the capacitors in a periodic manner. To assist in decorrelating the layout of the foundry capacitors, the heatmaps of errors are overlaid on the layout of the foundry capacitors in Figure 5.13, and their drivers in Figure 5.14, and illustrate no clear discrete pattern in error in relation to the layout as a function of a whole number of unit capacitor or driver cells.




Figure 5.13. A heatmap of systematic capacitor errors against the capacitor position in the fabricated array.

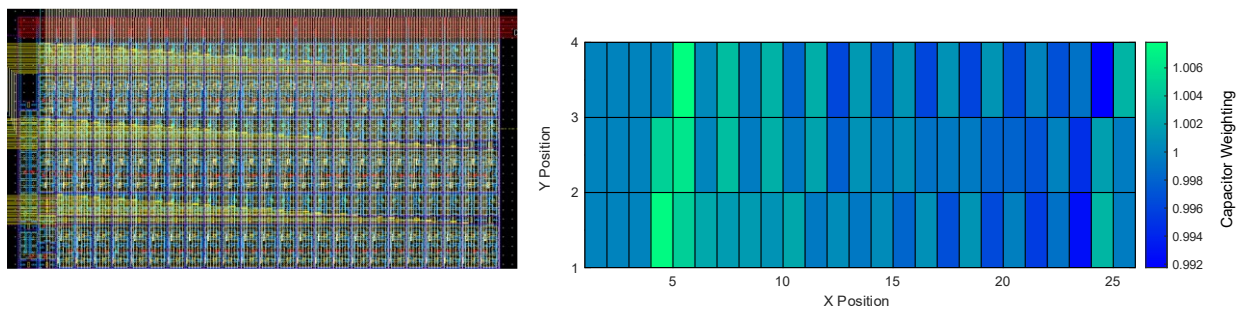


Figure 5.14. A heatmap of capacitor error in relation to the respective driver of the capacitor, showing no clear pattern relating to the layout of the drivers.

An FFT was taken of systematic errors for all measured capacitor geometries with respect to the capacitor positioning within the array, and normalised to the same highest contributor. This plot in Figure 5.15 illustrates a number of key periodic components that are consistent amongst all other measured capacitor geometries. Only the capacitor with geometry using the minimum foundry-permitted finger spacing and finger width displayed any isolated contributions at a pitch of $5.84\mu\text{m}$. The periodic components of Figure 5.15 were not visually identifiable in the plots of the capacitor systematic error for the other capacitor geometries, although the FFT confirms a strong presence.




Figure 5.15. Normalised FFTs of all mean capacitor arrays highlighting several periodic components in the errors. The larger periodic element discussed is masked by the larger low spatial-frequency contributors of the other arrays.

The above data was used to confirm that the observed periodicity in errors were not derived from the design of either the measurement system, shuffle matrices, or some artefact of the capacitor layout. This is believed primarily as the capacitor error period is not a multiple of the pitch of the unit capacitor, and the same period is correlated between both the foundry and custom laid out capacitors of a different geometry. This overall implies that the source of this error is likely process derived. The likely processing stages are the lithography, or some pre-processing stages performed by the foundry before fabrication which result in raster. Without further information held by the foundry on the specific parameters of these procedures, it is difficult to speculate as to the most likely cause.

5.4.2 Step Error in Capacitor Value Across Die

As previously observed, the systematic capacitor error varies massively with the drawn geometry. However, some systematic profiles appear not just with respect to the positioning within the array, but related to the positioning of the capacitor on the die. Suspicions were raised when the majority, but not all, of the arrays on one design of chip presented large steps in capacitor weighting of the order of 2000ppm, in positions that appear unrelated to any drawn geometry in layout, and are independent of the specific shuffle pattern that was used in measurement.

It was observed as shown in that when the mean capacitor errors values were plotted against their positioning from the bottom edge of the die, in Figure 5.16 and Figure 5.17, in the fabricated layout the steps for example at $472\mu\text{m}$ in row 1, all align across the arrays horizontally with the same approximate magnitude of error. The pitch of the step was measured reliably at $111\mu\text{m}$. This does not correlate with any known figure relating to the process density checks, design rule checks (DRC), optical-process-correction (OPC), or metal fill generation windows. Neither does the pattern align with any known geometry at the die edge. Any documented reference point would be either $28\mu\text{m}$ inboard of the die drawn edge, or $83\mu\text{m}$ outside. There is no available documented geometry placed at this location during either top-level die finishing by the foundry to the author's or any consulted party's knowledge.

The natural question arises as to why not all die exhibited this step. While most array designs were under $111\mu\text{m}$ in length, it was conceivable that the arrays might be placed, luckily for the designer, between the position that these steps would occur. The plot of confirms this theory for the test-chip 2 chip arrays. Whilst there are still steps in the capacitor weightings of the expected magnitude at the extrapolated positions from the other measured arrays, most arrays were positioned in-between these step locations. From this data it is deduced that this step is created through some stage inherent in the process, and is not derived from the design.

Lithography or a pre-process raster of some description is the favoured source, however due to the foundry holding the exact process parameters and recipe a closely guarded secret, there is little more that can be done to deduce more about the error's source. Other possible origins include the non-linearity of the stepper during mask alignment, or quite possible something related to the optical-proximity-correction (OPC) algorithm window. This is a process run by the foundry before fabrication of the masks to correct for the etching profile and ensure the fabricated geometry best matches the drawn geometry [67], [68].

Chapter 5

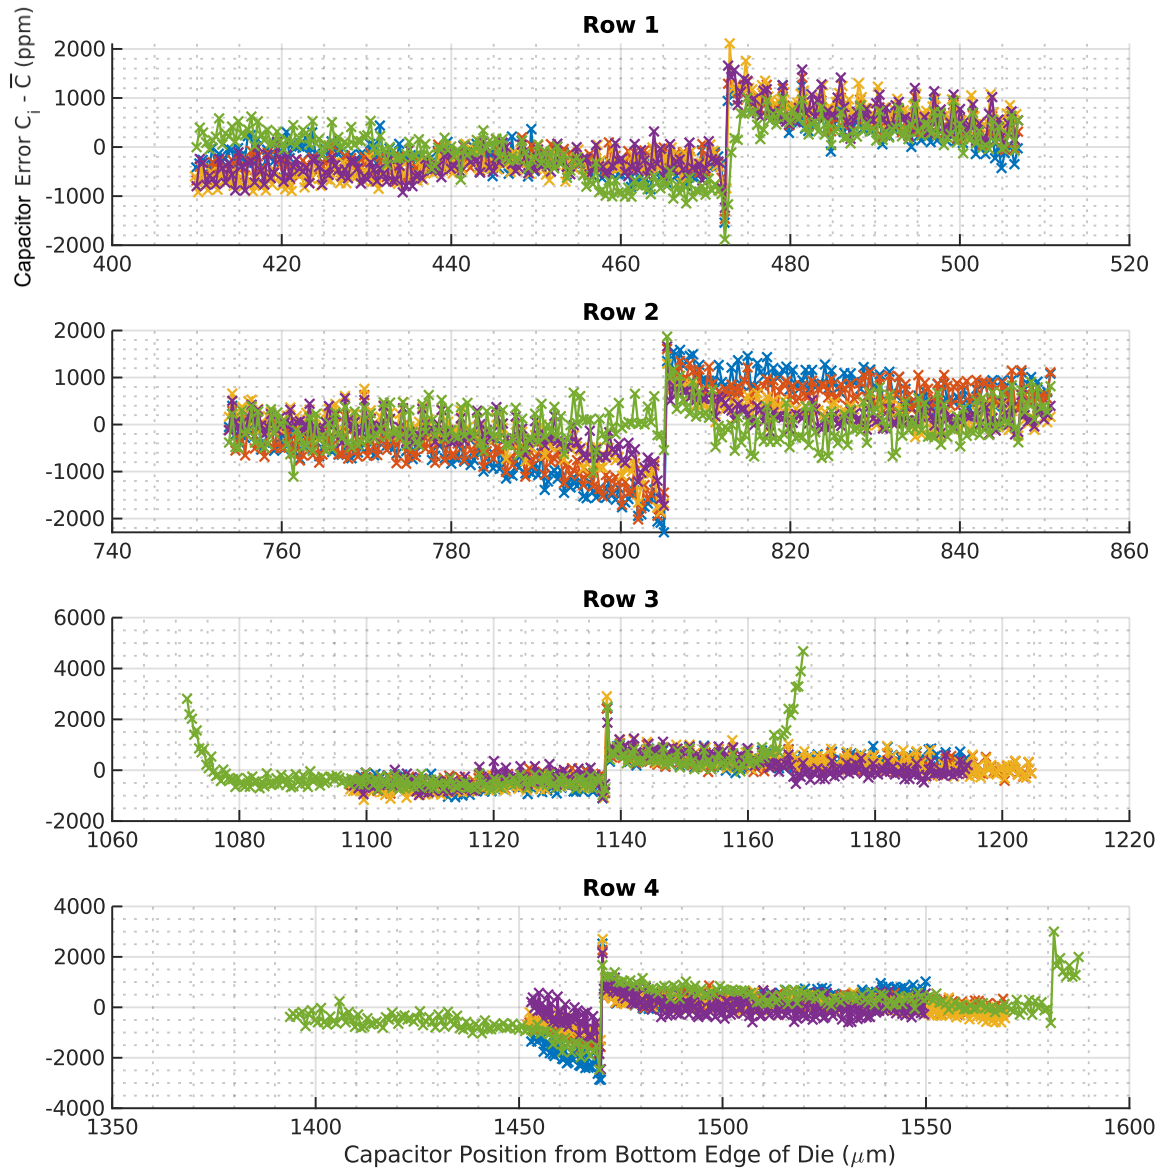


Figure 5.16. Normalised capacitor values from test-chip 3 plotted against their positioning from the die bottom edge revealing patterning errors ('steps') independent of capacitor specific geometry.

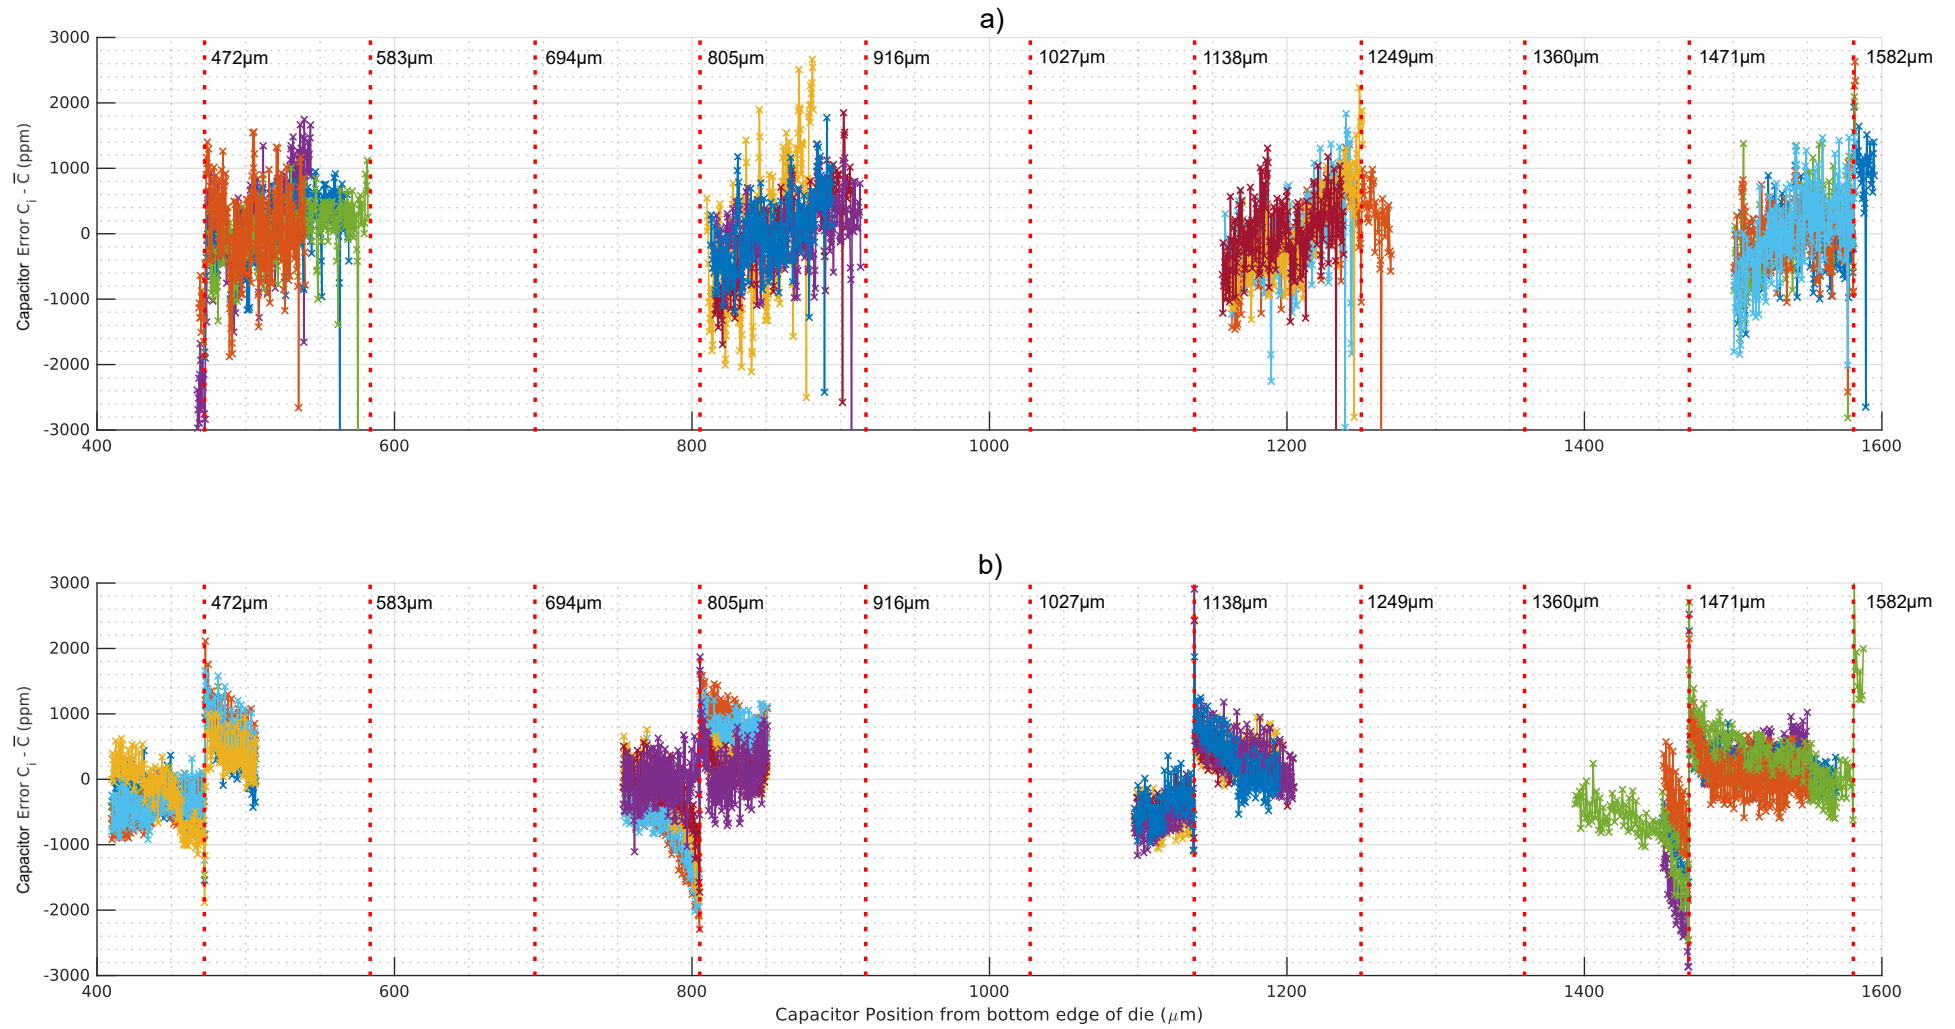


Figure 5.17. All capacitor errors plotted with respect to the die bottom edge to illustrate the step patterning. The theoretical positioning of the steps is indicated by the dashed red line at 111nm intervals. a) test chip 2 arrays, b) test chip 3 arrays.

The steps were consistent in size and location on all 174 measured dies over 6 wafers, including skewed material. The step positions are further visualised across the two chips in Figure 5.18. This was consistent across all BEOL metals with equal contribution to capacitance error. This clearly illustrates that a currently unknown sensitivity exists within the fabrication process of the BEOL metal layers. The sources of these errors are later discussed further in section 6.6.

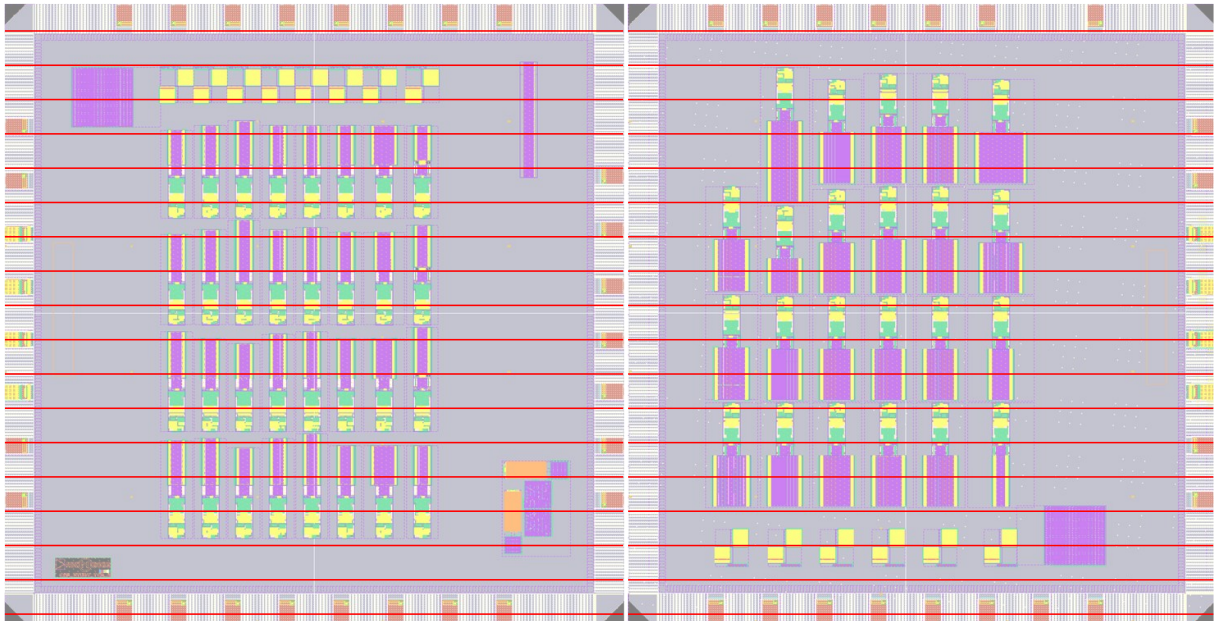


Figure 5.18. Visualisation of the positioning of the steps on the two chip designs. The steps align predominantly with the ends of arrays on the test chip 2 die (left), and in the centre of most arrays on the test chip 3 die (right). The positions of the measured and predicted steps are shown as red horizontal lines.

5.4.3 Step Error Removal for Further Analysis

Before the statistics may be extracted from the measurement data, the systematic step in the capacitor's weightings is first removed. As the step is relative to the die edge, and not geometry specific as has been established, it provides uneven contribution to each array and therefore inflates the systematic errors in a way uncorrelated with the geometry. Without removal, this would lead to corruption of any extracted systematic error trends. This applies only to the systematic matching error σ_S , and the total matching error σ_T . For calculation of random matching error σ_R , the method of calculation is agnostic of systematic errors so data inclusive of the capacitors around the step is still used. No localised increase in random error was observed in the region of the step, so the random error statistics remain unaffected.

In the case of the relative matching statistics and trends, the original measured data must also be used. The resulting trend relies on knowledge of the distance between capacitors, so no capacitors may be discarded. The calculation is independent of systematic errors, so the

results remain unaffected by the presence of the step. This case also applies to FFT analysis of the periodic systematic content in the following sections to avoid corruption. Furthermore, removal of the step in the data thereby allows valid comparison between the data from test chips 2 and 3 without these location dependent matching effects.

Where required, the step is extracted by removal of several capacitors surrounding the step, and renormalisation of the array. The capacitors indices for removal and the scaling factors for the two remaining halves of the array are determined by analysis of the mean capacitor profile \overline{C}_i which allows easy identification of the step location n in Figure 5.19. Indices $(n - 3)$ through to $(n + 4)$ are removed from the dataset. Capacitor weights for indices greater than $(n + 4)$ are then offset such that $C_{n+5} = C_{n-3}$, then the two halves of the array are rejoined, and re-normalised. This offset correction is of the order of 2000ppm, and the impact on σ_S is proportional to the number of capacitors the offset is applied to. The absolute worst case applies to a 2000ppm offset to 246 capacitors, resulting in a 10ppm error in the final statistic. In the majority of cases in , approximately half of the array experiences this offset, and as such the impact is anticipated to be closer to 5ppm which is similar the measurement error.




Figure 5.19. Mean capacitor profile step removal process. A red cross indicates a removed capacitor, whilst a blue cross indicates a retained capacitor.

5.5 Trends in Matching Error with Capacitor Geometry Variations

5.5.1 Capacitor Error Variations with Core Capacitor Geometry

The systematic and random matching error statistics for two capacitor geometries with varied finger length is shown in Figure 5.20. The two capacitor geometries are sized to assess the

consistency of the relationship across other geometry parameters. The systematic matching error, while consistent for each finger length measured, does not show a clear trend. A number of anomalies were observed for finger lengths of $19\mu\text{m}$, and $29\mu\text{m}$ where significant increases in matching errors above what is expected is seen. This may be attributed to the dependency of the systematic error exhibited on the array's location on the die. To this note, three arrays were instantiated twice on the same die, located at opposing edges of the matrix of arrays (the far left and far right columns). Of these arrays, the matching statistics for the first two showed a 38ppm discrepancy in systematic error. The third array geometry however showed a significant 132ppm discrepancy (of opposite polarity). This not been able to be attributed to any boundary or proximity related effects based on the layout, and the absolute values of the capacitors were within 1% of each other. This effect may account for the discrepancy noted here, however all arrays for the finger length sweep were located within the same column so the link to this effect is limited. It is still noted however that moving away from the smaller geometry capacitors, does yield a general decrease in systematic error.

The randomly sourced errors are illustrated for the same capacitors in Figure 5.20b. For both of the trialled capacitor geometries, an increase in finger length yields with little exception a decrease in capacitor random errors. An improvement in the random matching error with an increase in finger length is expected as the random defects that accumulate in the capacitor are averaged out further as the size of the capacitor is increased. As previously established, the finger length contributes directly to the coupling area of the capacitor, so this trend is expected. Features such as line edge roughness (LER) from lithography and etch processes are generally accepted as the main contributors to these mismatch errors [37], [63], [69], [70], [71].

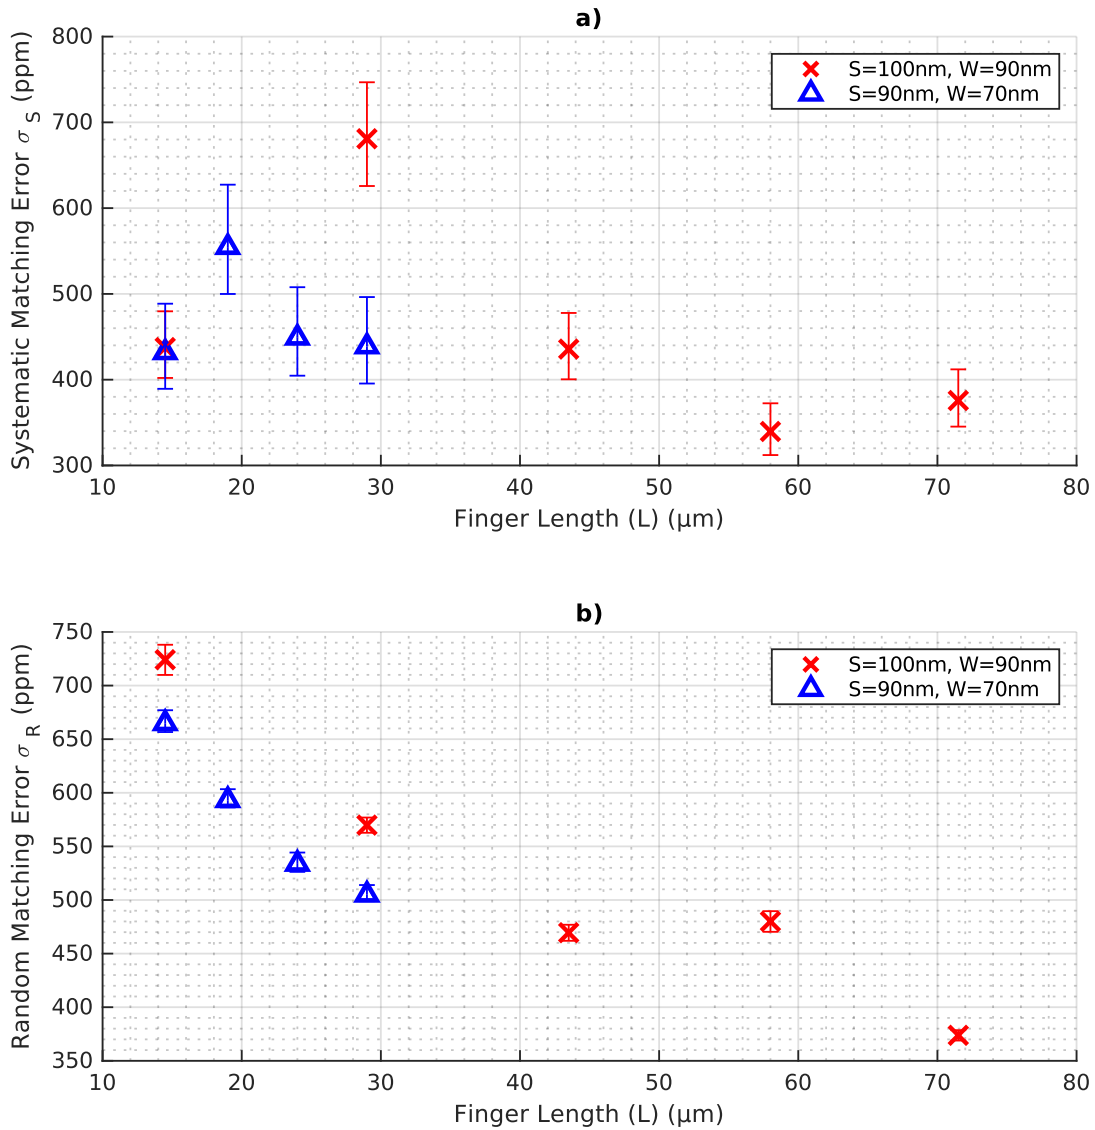


Figure 5.20. Capacitor matching error shown with respect to capacitor finger length (L) for both the a) systematic matching error σ_S , and b) random matching error σ_R .

The resulting capacitance is proportional to the number of metals used to implement the capacitor; therefore, it would be expected that a similar trend to varied finger length would exist for the number of metals N_m . Figure 5.21 shows that for an increased number of metals used, a large decrease in the random matching error for the capacitor may be attained. Intuitively this improvement may be attributed to the same reasons that were previously described for the finger length, and the total coupling area of the capacitor will increase proportional to the number of capacitor metals, and as such, provide more averaging of the LER errors resulting in a reduced total error for larger capacitors.




Figure 5.21. The relationship between the systematic (a) and random (b) smatching error and the number of metals used in the capacitor.

In earlier discussions it was proposed that by implementing the capacitor in the higher core metals, it may achieve better matching performance than those of the lower ‘dirty’ metals which are more susceptible to errors and stresses from the substrate below them. This is also the theory that drives investigation into the use of metal shielding below and above the capacitor in later discussions. Table 5.2 gives the matching statistics for capacitors implemented in upper and lower metals. For a 2-layer capacitor, the two tested metal combinations were M4-5 and M5-6. 49ppm better random matching error is seen on the lower metal implementation using M4-5 while significant improvements in systematic error are seen by utilising the higher metal combinations. The improvements for a 4-layer capacitor are less conclusive with the variations between matching errors being on the same level as the derivation of the statistics (<20ppm). The takeaway from these results implies that a capacitors random matching error may be reduced through implementation in multiple metal layers, with negligible impact resulting from metal choice between metals 2 through 6.

Table 5.2 Matching error statistics for 2-metal and 4-metal capacitors implemented in upper or lower metals.

Metals Used	N_m	σ_S (ppm)	σ_R (ppm)
M4-5	2	936	978
M5-6		673	1027
M2-5	4	436	545
M3-6		461	568

The results for the matching errors across varied finger spacing are shown in Figure 5.22. The systematic errors in Figure 5.22a show a sharp decline when geometry is increased beyond minimum spacing. This is true only in the case of also having minimum finger width (70nm). This represents a 38% improvement in systematic matching for only a 7% increase in top view area when moving from 70 to 80nm spacing. This is a significant improvement that is well above the usually assumed $1/\sqrt{A}$ trend which should comparatively yield only a 3.4% improvement were it to apply here. The decrease in systematic errors appears to continue for increased finger spacing, but appear to plateau above 110nm. The discontinuity for 100nm finger spacing appears an outlier. Previously discussed data regarding variations in matching based on position on the die showed significantly smaller variations than this point 220ppm from the expected position. This increase in error is marked as being linked to either the lithography or OPC processing windows that are discussed towards the end of this chapter.

The random matching error shown in Figure 5.22b illustrates a similar trend. The matching error for the smaller trialled capacitor consistently decreases with an increase in finger spacing. For the larger capacitor with a more limited spread of geometry, the trend is inconclusive due to the small data set. The larger capacitor's errors are generally smaller due to the increased finger length as previously established. The data shows that at least until 120nm finger spacing, and improvement in random matching errors may be achieved using minimum finger width.

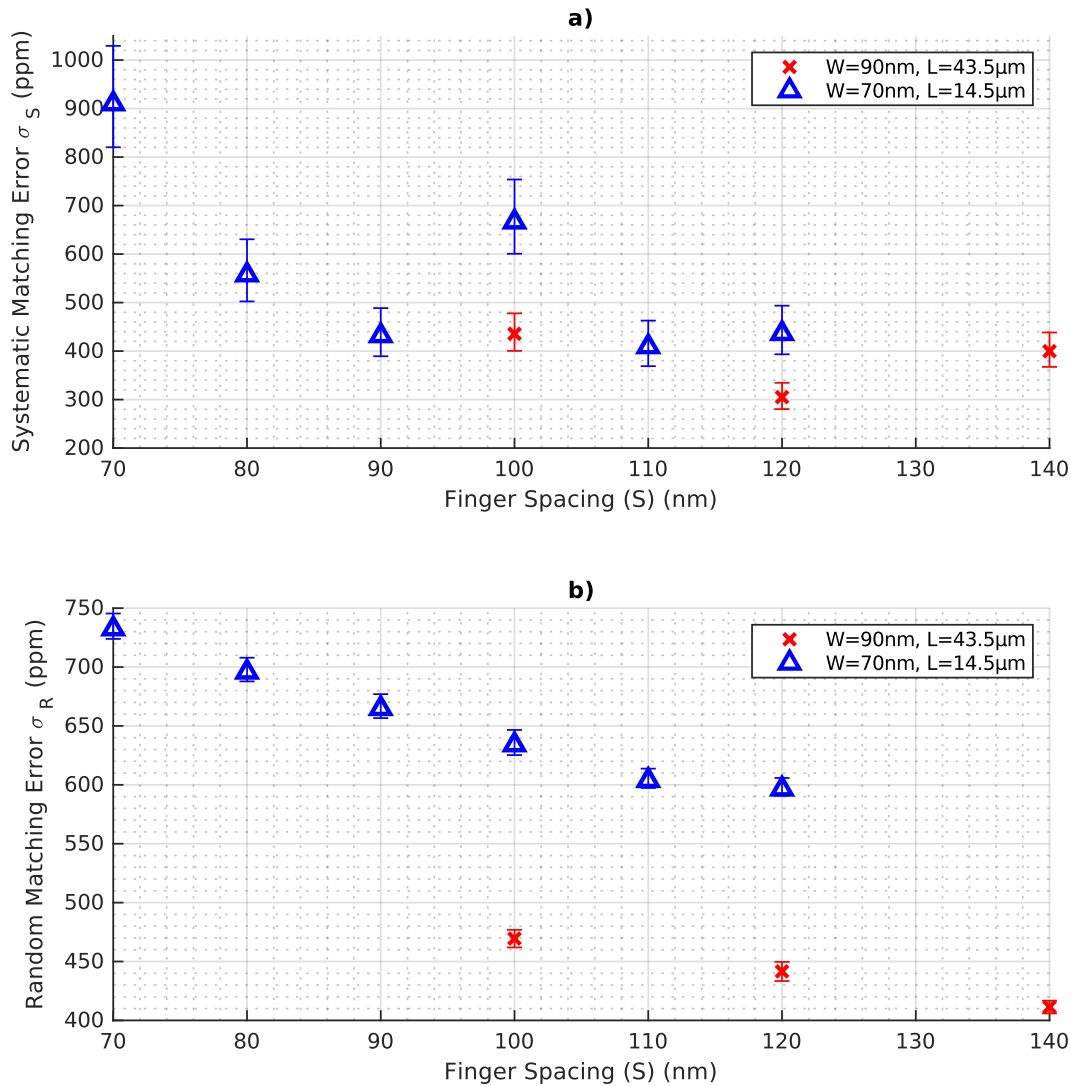


Figure 5.22. The relationship between the matching error and the lateral capacitor finger spacing (S). a) Systematic matching error σ_S and b) random matching error σ_R .

A similar trend of improved matching is also observed for the finger width parameter in Figure 5.23a, where a 44% improvement in systematic error is made for a 28% increase in finger width for a capacitor of minimum finger spacing, and a 38% increase for a capacitor of 100nm finger spacing. A minima, which would be expected in the data, is located at around 110nm finger width however the range of the data is insufficient to determine this conclusively across all geometries. The minima is expected in the data as it is related to the classical matching theory regarding the die gradient. If the spacing between capacitors is increased, they therefore should achieve better matching due to the larger devices having more variations averaging out. However, if a device becomes large enough, the gradient across the die becomes the dominant contributor to the matching error. Therefore there must be a sweet spot where the maximum amount of averaging is achieved without the die gradient being dominant.

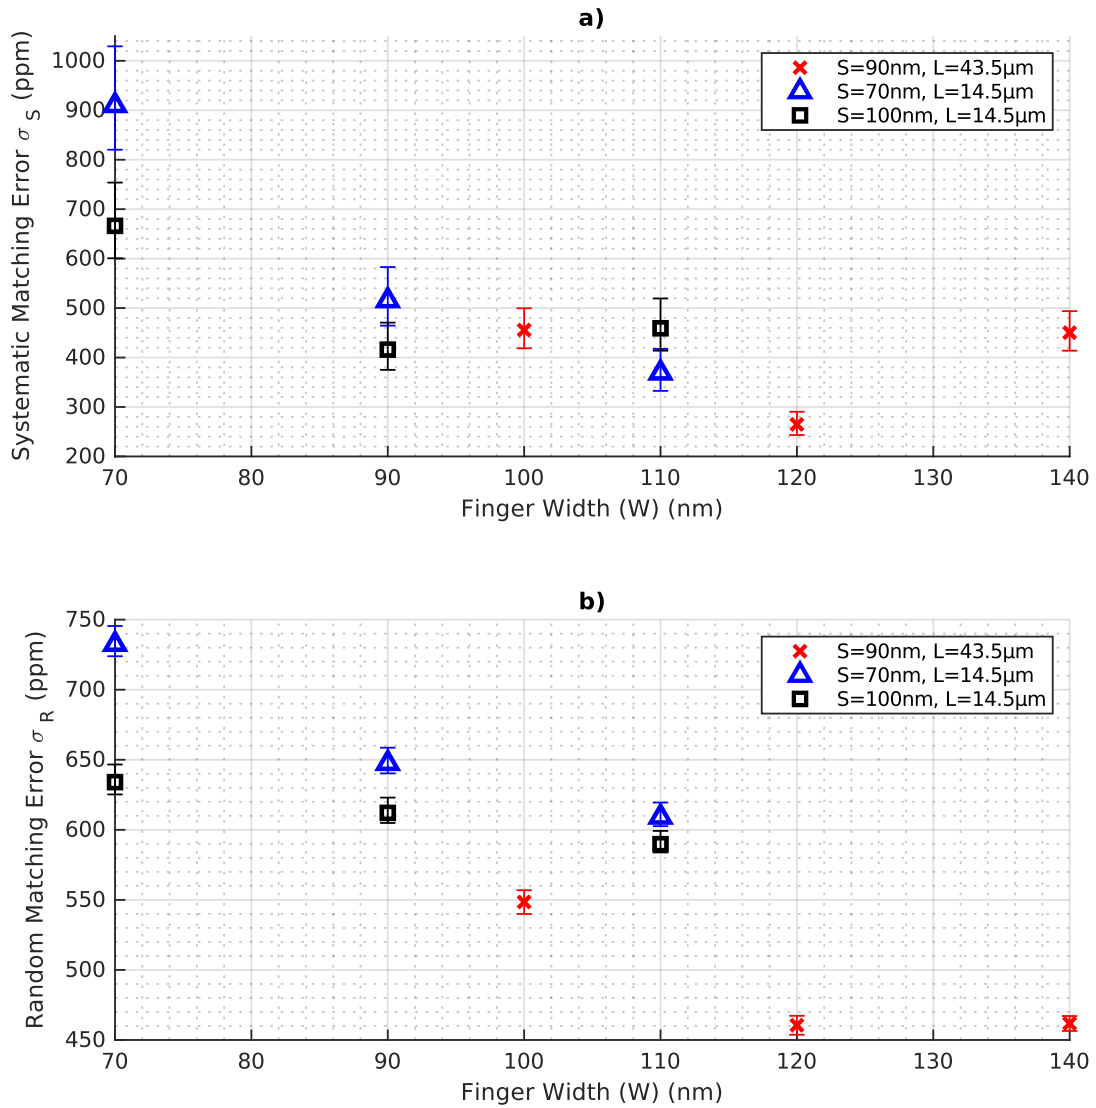


Figure 5.23. The relationship between the matching error and the lateral capacitor finger width W . a) systematic Matching σ_S , b) Random Matching σ_R .

It can easily be reasoned, particularly with the evidence from Figure 5.15 that the absorbed systematic error of a capacitor array will increase with the size of the array due to the large presence of errors active over large distances. This was illustrated in the FFT of capacitors in Figure 5.15. For example, a linear gradient would have a larger delta from one end of the array to the other, and subsequently introduce larger errors. This is not an effect that was observed in the presented data. For the increased finger length, it may be expected that the more complex gradient components are averaged out as the array dimensions are increased in a direction that maintains the maximum height of the array. Any increase in this dimension only serves to average the process gradient and other systematic components further, and thus reduce the total absorbed systematic error.

With the finger spacing and width however, increasing these parameters increases the maximum distance between capacitors, and would intuitively lead to an increase in capacitor

systematic error. This was not observed to be the case for the results discussed, which were all sized very close to the minimum geometry possible in this process node.

The reduction in random errors with increase finger spacing and width however, poses a more intriguing puzzle. Neither of these parameters has previously been linked to a decrease in capacitor matching errors of any description before. In the case of finger width, nor does it contribute directly to the capacitor value. However, as has been previously established in the previous chapter, the fabricated capacitor geometry in the 40nm process appears highly dependent on the specific patterning and metal densities. It has already been identified that there are a number of second order dependencies on the capacitor value due to both the high level of fringing, as well as some metal density dependencies. If LER is to be decided as the main contributor, then this would demonstrate an LER dependency on metal trace width and spacing [56]. Alternatively, this error may come from formation of the porous intermetal dielectric which may not as reliably form, or be damaged by the complex metal deposition process.

The finger end spacing S_e is often attributed as a source of error in the capacitor value. While never formally discussed or analysed due to the complexity of the manufacturing process, the definition of the lateral capacitor finger would intuitively be the least consistent at the end due to the difficulty in defining such a small shape in lithography. This is for the same reasons that 45° traces and curves are not permitted for regular metal wires in the 40nm process. This will of course introduce some offset into the capacitor value from the idealised version and equations discussed in earlier chapters, although it is expected that this error would occur equally on every instance of the capacitor. Even if the geometry varies more significantly than the rest of the finger, it has already been established that this portion of the structure contributes very minorly to the overall capacitance. Hence, its impact is minimised.

The measurements of the random matching in Figure 5.24 reflect these intuitions. A particularly minor decline in random mismatch error is seen for an increase in the finger end spacing. A 4% decline for a 150% increase in end spacing and a 2% increase in area. A sweet spot is noted at 140nm spacing for the random matching errors, whilst the systematic errors increase with increased finger end spacing due to the discontinuity in metal density it introduces.




Figure 5.24. Random matching errors for varied capacitor finger end spacing. $S=100\text{nm}$, $W=90\text{nm}$, $L=43.5\mu\text{m}$, $N_m=4$.

It has already been established that in general an increase in the parameters contributing to the top view area of the capacitor, as well as the number of metal layers, leads to a general decrease in the random matching error. The top view area is defined in equation 5.17. The random matching error is plotted against this defined top view area inclusive of the number of metals N_m in Figure 5.25 with the results grouped by the two test chips that they were derived from, with a summary of parameters presented in Table 5.3. The reason for this is evident in the three trendlines that are drawn for the same relationship considering only the test chip 2 data, test chip 3 data, and all test chip data. Analysing only the data from the second test chip, a strong fit for the relationship of equation 5.18 is seen with an error scaling parameter γ of $3827 \text{ ppm}/\sqrt{\mu\text{m}}$. The fit for the test chip 3 data is even stronger, and the slightly reduced value of γ of 2879 implies a reduced improvement in random matching error was seen on these arrays compared with that of the second test chip. Combining all data to generalise the fit yields for this model gives a γ of 3062 . The literature frequently cites a $1/\sqrt{C}$ relationship as an approximation for the total matching for both vertical and lateral coupling capacitors [5], [31], [36], [41], [49], [69], [70], [71], [72], [73], [74], [75], [76], [77], [78], [79], [80], [81], however due to the contributions from increased finger spacing and width being identified, it is apparent that these models are unsuitable for description of this capacitor geometry.

Given the quality of the fit over the range of capacitor sizes and values that have been measured in this work, this is presented as a reasonable first order model for the random matching error of

a well-designed single finger capacitor of this precise geometry. Expansion beyond this to alternative geometries cannot be made. The foundry geometry capacitor is a prime example of this, achieving a σ_R of 706ppm for a value of 40fF consuming an equivalent area of $14.4\mu\text{m}^2$. Scaling this for a capacitor of equivalent value to the reference capacitor of the optimised geometry using the proposed model, this produces a 23fF capacitor with a σ_R of 931ppm consuming an area of $8.3\mu\text{m}^2$. This is nearly twice the random matching error of the optimised capacitor geometry, which represents almost a 2x improvement. For these reasons also this model could only confidently be applied to the optimised capacitor structure within this thesis given the collected data. If the total matching statistics are considered, then this improvement is extended to 4x.

The discrepancy between the two test chips was consistent across the designs they each contained. The duplicated designs across both test chip designs (54 capacitor arrays, 174 measured dies) provided values for σ_R that were consistent between the two instances to within 30ppm, with the value for test chip 3 consistently higher. This test chip contained arrays that were generally larger in value than the others. The first test chip not considered here only contained designs of the foundry cell capacitors, which do not align with either of the trends shown here, and achieve considerably worse matching for their capacitance due to their use of a large number of fingers, and poor metal density and other inconsistencies in layout.

$$A_{Top} = 2L(S + W) \quad 5.17$$

$$\sigma_R = \frac{\gamma}{\sqrt{A_{Top}N_m}} \quad 5.18$$




Figure 5.25. Random matching error of all capacitor arrays plotted against the most common model used by the design community [31], [65]; $1/\sqrt{N_m A_{top}}$. Three fits for equation 5.18 are shown and detailed in Table 5.3.

Table 5.3. Parameters and measures of fit for the extracted trendlines in Figure 5.25 for use with equation 5.18.

Data	γ (ppm · μm)	R^2
Test Chip 2	3827	0.925
Test Chip 3	2879	0.984
All Available Data	3062	0.435

As was noted before, the systematic and therefore the total considered error of the capacitor structures did not show any clear trend with the capacitor geometry aside from the significant improvement observed when the finger width or spacing is increased away from minimum sizing. Comparison against the foundry's mismatch model is not feasible here due to the lack of separation it provides between random and systematic contributing errors.

Figure 5.26 illustrates the systematic matching components plotted against the capacitor pitch. It was previously identified that the finger length relates differently to the systematic error so results here are grouped by length. For the third test chip designs of shorter $43.5\mu m$ length the minimum in systematic error is seen at $0.42\mu m$ pitch. Significant improvements are seen in the second test chip designs of $14.5\mu m$ finger length, for pitches moving away from the minimum

geometry. The improvements continue until a pitch of $0.36\mu\text{m}$ after which there is no improvement seen.




Figure 5.26. Systematic matching errors from an assortment of custom capacitor geometries plotted against the resultant capacitor pitch. The overall trend shows improvement until approximately $0.38\mu\text{m}$, with no consistent improvement beyond this.

5.5.2 Sub-DAC Matching Errors

The sub-DAC is the primary method through which to achieve a high-resolution single-stage DAC, without requiring very large numbers of capacitor elements to achieve high dynamic range or large number of bits. The two custom designed DACs utilise the same optimised capacitor structures that have been previously discussed. The critical section of the DAC where the coupling capacitor is positioned is shown in Figure 5.27 and Figure 5.28. The two primary considerations in the performance of a sub-DAC are the gain error and the localised errors of the capacitors directly adjacent to the coupling capacitor. These are the quantities that the analysis is focused on. The gain error is considered relative to the error that the layout extraction tool produces. A gain error will always exist between the main and sub-DACs due to the presence of parasitics in the sub-DAC and coupling capacitors, and designers expect to have to trim the layout to account for these effects relying on layout extraction tools, typically 2.5D field solvers, to do so. A typical error expected between the schematic and extracted layout for the gain error is in the region of 1-5%. For the purposes of this discussion the gain error is defined here as in equation 5.19.

$$\text{Gain Error} = \frac{\left(\frac{LSB_{Ideal}}{MSB_{Ideal}}\right)}{\left(\frac{LSB_{Measured}}{MSB_{Measured}}\right)} \quad 5.19$$

Each sub-DAC consisted of 224 unit capacitors connected as standard to the top plate, with 32 connected through an equivalent of a 23 fF series coupling capacitor. This latter section formed the equivalent of a 5-bit sub-DAC. A substantial amount more averaging was required in order to produce accurate and low noise measurements of the same quality as the main DAC capacitors for the sub-DAC capacitors due to their comparatively small output range. The sub-DAC structures also only featured 1.5 μm of capacitor dummy structures between the active DAC capacitors and the coupling capacitor to give visibility of any possible systematic errors.




Figure 5.27. Sub-DAC variation-1. Blue shows the main DAC 'top plate'. Pink shows the sub-DAC 'top plate'.




Figure 5.28. Sub-DAC variation-2. Red shows the main DAC 'top plate'. Yellow shows the sub-DAC 'top plate'.

The gain error variations of the two custom sub-DAC variations are shown in Table 5.4. The magnitude of the gain error itself was quite large compared with the predicted value ($\sim 3x$), so is not considered usefully here. The standard distribution is considered reflective of the useful metric of the sub-DAC design, as offsets are easily trimmed out blindly and contribute the same as a systematic matching error, whereas the variation of the gain error between die manifests as a randomly sourced error, and would require part-by-part calibration to remove. Neither of the

sub-DAC coupling capacitor variations showed particular merit in their performance and their gain errors are in line with expectations from the foundry cell capacitors. The equivalent σ_R for an individual capacitor of equivalent value scaled by finger length equates to approximately 3000ppm. Clearly, the main source of the sub-DAC error is derived from the coupling capacitor, incurring a large penalty due to the disruption of the patterning of the array through its presence.

Table 5.4. Summarised gain errors of the two custom sub-DAC variations measured, and the foundry cell equivalent implementation.

Sub-DAC Variation	Gain Error Standard Distribution σ
Custom Variation 1	0.47%
Custom Variation 2	0.48 %

This disruption is highlighted in Figure 5.29 where the systematic errors of the capacitors across the main-sub-DAC boundaries are presented. In both cases, the capacitors located directly adjacent to the series coupling capacitor (with only 1.5 μm of dummyming) incur substantial errors due to disruption of the patterning. The effect is clear in the sub-DAC portion of the arrays with up to a 1.7% error observed on the closest capacitors. These errors indicate a further 3 μm of dummyming between the used DAC capacitors and the series coupling capacitor are required for elimination of the gross systematic error observed. Overall, variation 1 labelled 'DAC120' would be preferred for proper implementation in a SAR sub-DAC. This is due to the substantially lower systematic errors observed within the sub-DAC whilst still maintaining only a 0.47% gain error variation, a figure that is substantially lower than the 4-5% which is typical of the foundry sub-DAC implementations.




Figure 5.29. Sub-DAC mean capacitor error around the main to sub-DAC boundary for the two trialled geometries shown.

5.5.3 Improvements Using Multiple Capacitor Fingers

In a CDAC it is expected to have the same finger length for each unit capacitor, and then to scale the number of capacitor fingers to achieve the required value for each binary weighted capacitor. As discussed, it is expected that the capacitor errors will scale with the number of fingers used, due to the effectively increased finger coupling length allowing for more averaging of the coupling interface surface roughness. To test this, an array with a unit capacitor having two fingers was designed and measured. For comparison, the matching statistics of the capacitor designs with single, and double fingers are shown in Table 5.5. The capacitor matching statistics for a capacitor with similar increased finger lengths is also shown for comparison of the trade-offs.

From the presented data, it can plainly be seen that the increased vertical size of the array from use of 2 fingers per unit capacitor has incurred a significant near-two-fold increase in the systematic error contribution over the single finger design. This statistic is calculated after the systematic step in section 5.4 was removed which would otherwise significantly inflate this systematic error figure, as the double-fingered array was large enough to incur two steps as opposed to the majority of the arrays which incurred only one. The longer pitch of the double-fingered array has led to the further inclusion of longer periodic components, including any die gradient that is present, unlike the smaller arrays. This is also apparent when comparing the

array with extended finger length which has the same array vertical length, but near double the array width on the die. This array actually displays a reduced systematic error, as the longer fingers expand the array in a dimension that allows more averaging of the gradients. This leads to the conclusion that when implementing large arrays, it is significantly beneficial to implement the LSB capacitors as single, longer fingers as opposed to two shorter ones due to the significant increase in absorbed systematic error due to the increased height of the array.

The lack of any reduction in the random matching statistic is unexpected. It has previously been noted that an increase in coupling area through increasing the finger length yields improvements, however, no notable improvement is seen here with the increase in the number of capacitor fingers for a similarly dimensioned capacitor. There is no current reasoning for this relation, but it may be explained if the ends of the capacitor finger are somehow contributing significantly to random error. This is unlikely to be the case as it was previously seen that the finger end spacing has no meaningful impact on the final capacitor random matching error. This is ultimately somewhat, but not completely, explained by the reasoning that the double finger capacitor still incurs the same metal definition based errors as the single finger capacitor, although they would be expected to average out further than was observed with the increased capacitance.

Table 5.5. Capacitor matching statistics for a single, and double fingered capacitor, as well as a capacitor with increased finger length for comparison of effects.

Variable	Single Finger Capacitor	Double Finger Capacitor	Single finger Capacitor with Increased Finger Length
S	100nm		
W	90nm		
L	43.5 μ m		71.5 μ m
N_m	4		
N_f	1	2	1
σ_S	436 ppm	882 ppm	376 ppm
σ_R	545 ppm	572 ppm	438 ppm
σ_T	687 ppm	1027 ppm	567 ppm

5.6 Conclusion

This chapter presents a method of analysis for independent extraction of the systematic and random matching error contributions for the lateral coupling MOM capacitor in large arrays. Significant periodic and systematic errors are identified, and the core capacitor geometry analysed. This revealed several previously unidentified relationships between the finger spacing and width, and the systematic and random matching errors of the capacitors. Independent analysis of the random and systematic errors allows validation of the usual matching models in use to be analysed, and a modified model for the random matching errors and its applicability are presented. The systematic errors are far more complex than the standard linear gradient which is typically assumed in literature. It emphasises the large benefit for aggressive folding or interdigitation schemes in capacitor groupings for DACs for effective cancellation. The collected data suggests that an optimised capacitor geometry can achieve in excess of a factor of 4x improvement in total matching errors, with potential further improvement possible with the use of blind calibration due to the consistent nature of the significant systematic contribution across many independent dies, and the large content of errors being attributed to being systematic where smaller geometries are used. This chapter is the basis for this thesis' demonstration that there are further advances still to be made within the confines of the 40nm process, and that proper understanding of the mismatch's nature can lead to higher performing capacitors.

Chapter 6 Capacitor Errors and Trends in Large Arrays

6.1 Introduction

In the previous chapter the individual capacitor geometry is studied for optimisation of the collective matching performance of the capacitors. There are however further parameters that may be considered at the array level that can improve the overall performance when integrated into a full capacitive DAC. In this chapter, the effects of the relative matching between arrays, and between capacitors is presented through analysis of the test chip data. The advantages from a matching perspective of using a fully differential DAC as opposed to single ended are shown, and the effectiveness of both lateral and top and bottom array dummyming structures are quantified, illustrating that a significant amount of dummyming is required for full removal of the localised errors at the ends of arrays. The data and relationships gathered are then reflected on to summarise the identified avenues for improvement of the later coupling capacitor's matching in the process, as well as where such techniques are transferrable between processes.

6.2 Matching Within Differential DAC Arrays

Throughout the rest of this thesis the capacitor statistics presented are for a pair of capacitors, assuming that in the majority of high-speed and precision applications a differential DAC is preferred. It has also been assumed that the matching statistics presented may be taken as approximately $\sqrt{2}$ larger for a single ended array, as is the case for other devices. The mean capacitor values of the P and N array capacitors are shown in Figure 6.1. The strong periodic component that was previously discussed is evident in both mean capacitor values across the array, and is in alignment between the two arrays as expected. A confident measure for the systematic contributions was not able to be calculated due to the limited data collected for independently measured P and N arrays, and the substantial time required for performing further low-noise measurements. The improvement in random mismatch error experienced through actuating the capacitors as a fully differential array was in line with the predicted $\sqrt{2}$ improvement to within 5%.




Figure 6.1. Mean capacitor value for the P and N arrays of a single differential DAC array measured independently. The large periodic systematic error is seen to be correlated between both arrays and is aligned by position.

6.3 Impact of Dummy Structures on Capacitor Matching

6.3.1 Dummy Structures at the Array Ends

Figure 6.2 presents the systematic error of two capacitor arrays with $5\mu\text{m}$ and $30\mu\text{m}$ of dummy structures at each end. As mentioned before these dummy structures are all replicas of the main DAC capacitors, and their respective drivers. It is clear from the figure that the presence of the additional $25\mu\text{m}$ of end dummifying has removed the large systematic error at the ends of the array. This is clear as a 4000ppm error is seen to act over the first $15\mu\text{m}$ of the array when only $5\mu\text{m}$ of dummy structure is used. These errors are only due to the systematic effects which are consistent during processing and are likely a result of changes in density in the metal patterning. Whether this is due to density sensitivity in the plasma etch process stages of fabrication, or an artefact of the CMP process, which has already been shown to product dishing over larger structures, is unclear from this data. However, the data is conclusive that at least $20\mu\text{m}$ of dummy structures total is required for assured removal of these errors. Further dummy structures may be required depending on the adjacent cell metal densities. However this was not considered in these designs.




Figure 6.2. Impact of capacitor array end dummy structures for the minimal 5 μm , as well as the predetermined 30 μm of capacitor and driver dummying.

Consider now the effect of the dummying on each capacitor's random variations. In Figure 6.3 a clear difference in random variations is observed between the array with 5 μm , and 30 μm of dummying. The array with 5 μm of dummying displays a >1500ppm increase in random matching at the end of the array, with the end 10-15 μm capacitors affected. The inconsistency in this value is due to the difference in neighbouring circuitry each end of the array. A small increase in random errors are observed at the end of the array with further dummying, although at a notably smaller level: approximately 200ppm increase in σ_i for the outer 20 μm of the array compared with the center. The baseline magnitude of error for the central portion of the array is similar in both cases, excepting for a localized increase in errors at 15-50 μm with only 5 μm of dummying.




Figure 6.3. Capacitor random matching error σ_i for arrays with 5 μm and 30 μm of top and bottom (end) dummy structures.

The difference in baseline σ_R for the two arrays may be attributed to the scale of the area that they occupy. Fundamentally, a more even polish from CMP would be expected over larger areas, such as larger arrays. Increasing the dummyming at each end from $5\mu\text{m}$ to $30\mu\text{m}$, results in a 47% increase in occupied area. It is therefore be expected to achieve more predictable results due to this significant increase in area due to the reduction of localised effects. If one were aiming to completely eliminate all increases of errors at the ends of the capacitor array, then a further $10\text{-}20\mu\text{m}$ may be required. However, the level of improvement that would be observed does not justify this massive increase in area, as with this the useful area would occupy the same area as the dummy structures.

6.3.2 Lateral Dummy Structures

Further to the discussion of top and bottom array dummy structures, the effect of lateral dummyming between the capacitors and their drivers is also considered. The matching statistics are tabulated in Table 6.1, and detail more than 10% variation through presence of the lateral dummyming structures.

For the capacitor geometry utilising minimum finger spacing and width, and $14.5\mu\text{m}$ long fingers to achieve a nominally 8fF capacitor, significant detriment is found through use of lateral dummyming. Systematic errors are increased by nearly 30% with $5\mu\text{m}$ of dummyming. The average capacitor error profile in Figure 6.4 shows that the previously discussed periodic content is upset by the lack of dummyming, likely due to the reduced regularity around the array structure. This is further analysed through use of the FFT of Figure 6.5 in which the two dominant periods at $8.96\mu\text{m}$ and $2.24\mu\text{m}$ are only seen where lateral dummyming is utilised. In the alternative case, this periodic component is replaced with another large periodic component at a considerably higher spatial frequency, a pitch of $0.84\mu\text{m}$. No other notable periodic content is identified. The FFTs of Figure 6.5 are normalised to the assumed-constant gradient component to allow for relative comparison.

Table 6.1. Matching error statistics for identical capacitor arrays with and without lateral dummyming structures in layout.

	S=70nm, W=70nm, L=14.5μm			S=90nm, W=70nm, L=29μm		
Lateral Dummyming (μm)	σ_S (ppm)	σ_R (ppm)	σ_T (ppm)	σ_S (ppm)	σ_R (ppm)	σ_T (ppm)
0.0	638	659	919	629	510	733
5.0	910	733	1186	439	505	673
Improvement (%)	-42.6	-11.2	-22.5	+30.2	+1.0	+8.2




Figure 6.4. Mean capacitor profile for array with and without lateral dummying structures. ($S=70\text{nm}$, $W=70\text{nm}$, $L=14.5\mu\text{m}$) showing disruption of the longer periodic error.




Figure 6.5. FFT analysis of the mean capacitor profiles with and without lateral dummying present. The dominant periodic components are labelled.

For the case of the longer geometry finger, which avoids the use of minimum sized geometry, a limited improvement in all matching statistics is observed. This variation does not exhibit periodic components as strong as the minimum geometry capacitors discussed, and no notable degree of change is observed in the periodic content that is detectable.

When the random matching is considered on a capacitor-by-capacitor basis, the profile remains predominantly the same. For minimum geometry capacitors in Figure 6.6a, the general decline in σ_R with lateral dummying is observed, with approximately equal degradation across the array σ_R for non-minimum geometry

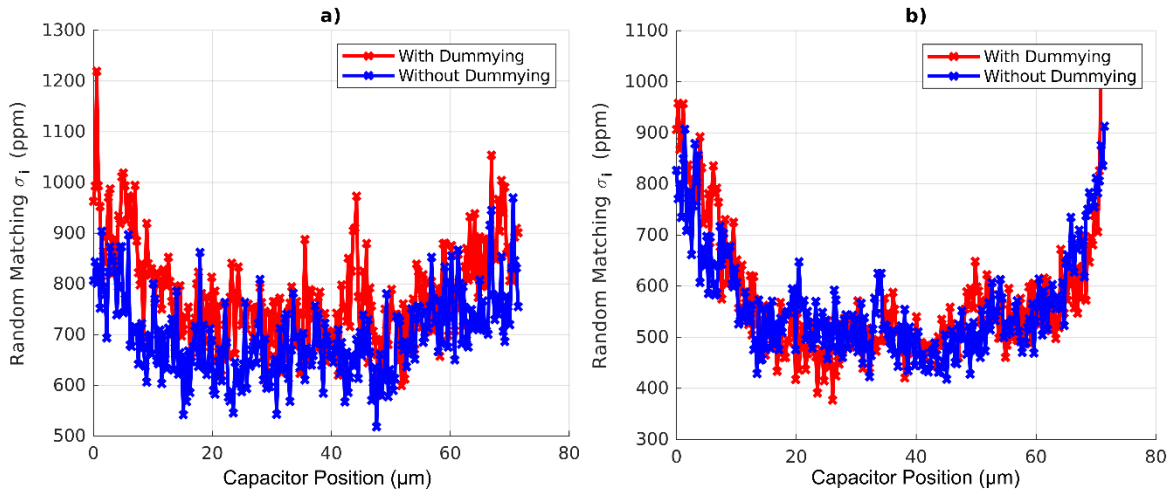


Figure 6.6. Random matching σ_i across the array with and without dummyming for a) short 14.5 μm finger lengths, and b) long 29 μm finger lengths.

Shorter period systematic components are arguably preferable to longer periods of oscillation as this ensures multiple repetitions within each grouping when folding and other layout cancellation techniques are employed. Of course, the extreme case of a linear gradient is easily cancelled, but provided the periodic component is far removed either way from the minimum grouping size in the folded layout, then the averaging between the multiple capacitor groups cancels the systematic components as intended.

The discrepancy in improvement may indicate that the use of minimum geometry, as a rule, leaves the capacitors most susceptible to the presence of periodic components. This is owing to the lack of consistent periodicity in the latter measured array of non-minimum geometry. These periodic components then contribute to the overall capacitor errors. The improvement in matching for the larger array is explained through typical theory of improved consistency and regularity in layout through presence of the dummyming, leading to reduced dishing and other inconsistencies resulting from the CMP and etching processes.

6.4 Impact of Array Shielding on Capacitor Matching

The capacitor array may be shielded from the potential impact of upper metals or the circuits below by introducing another layer of metal above or below the array. This is implemented in the polysilicon layers below the array as well to minimise the substrate capacitance and reduce the coupled noise. The same principle applies to the top metal shielding where power is generally distributed across the die in the top-most metal layers.

The impact of using such shielding layers is seen for a single capacitor geometry in Figure 6.7 with the systematic and random matching for the combinations of top metal (AP) shielding, and polysilicon (poly) shielding shown. The lowest random matching errors are attained with only

top metal shielding, although the margin by which this is the case is only small. It appears that poly shielding does not meaningfully affect the random errors. Increased systematic matching errors are seen with the inclusion of either top metal or poly shielding. It can be concluded that the best matching performance is achieved for an array utilising only AP shielding if a blind trim or calibration is to be used. If no shielding is to be used a 30ppm improvement in total matching errors may be attained, although this improvement (4%) is unlikely to yield measurable performance improvements in a full ADC. Realistically a design may choose to include these layers irrespective of their matching performance impact for electrical shielding purposes. At the very least, the presence of these layers does appear to meaningfully affect the overall matching properties of the array, provided a blind trim can be used.




Figure 6.7. Matching error performance of arrays using a varied combination of vertical shielding above and below the capacitor array. S=90nm, W=70nm, L=14.5 μ m, M2-M5.

6.5 Localised Errors in Large BEOL Capacitor Arrays

6.5.1 Relative Matching of Capacitors

Since capacitors are generally grouped when used in charge redistribution DACs to form larger value capacitors of relative weightings, the actual relative matching between the capacitors is of great value in understanding the applicability of layout techniques such as interdigitation and layout folding. The work presented previously throughout this chapter assumes that the errors categorised are random and decorrelated from each other. However, by undertaking the analysis through the methods detailed in section 5.2.3, correlated errors are revealed within the array.

Figure 6.8 illustrates the result of this analysis performed on a large sample of arrays, showing the relationship between the separation of any two selected capacitors within the array, and the distribution of the relative error between them. This method results in no contributions from systematic offsets between the capacitors such as die gradient or other periodic or stepped errors that have previously been discussed. There are two distinct regions observable in the curves of Figure 6.8, with a transition between them at around $10\mu\text{m}$ separation. Firstly, above the transition point, a linear relationship is observed, with the relative matching error of capacitors increasing as they are spaced further apart. No sensible relation between the capacitor geometry and the gradient of this region of the statistic has been identified.

Below the transition region of $10\mu\text{m}$, a sharp decrease in relative error is observed. Put plainly this means that the relative error between capacitors is substantially smaller when the distance between them is less than approximately $10\mu\text{m}$. The lower limit of the relative errors at minimum separation correlates somewhat with the top view area of the capacitor. This is due to the presence of correlated errors in the manufacturing of the die. Correlated errors in this context are those that affect multiple fingers of the capacitor array over a finite area. The nature of these correlated errors and their formation is discussed in the following section, however it is evident that these are due to localised affects, such as a limited deficit in etching, or scratches from the CMP process [58], [61], [82]. These are effects that are previously documented, and both have dependencies on the localised metal density in the magnitude of their effect, hence they may affect different capacitor geometries differently.

It is typically accepted that devices closer together have improved matching, although generally it is presented in the context of die gradient and other larger area effects. In work by Tuinhout looking at a similar statistic [49] the gradient for large capacitor separations is attributed to the systematic errors. However, even with the systematic errors artificially removed the same phenomena is observed. The method used here also allows analysis without the presence of systematic errors.




Figure 6.8. Capacitor relative matching error σ_x against the separation of the capacitors. Calculated over all available die, wafers, and arrays.

If the systematic contributions are also considered the results of this study are very different. This is achieved through taking of the rms value of the errors as opposed to the standard deviation, although as is shown, this complicates the relationships for examination. Figure 6.9 illustrates a wide variety of shaped profiles for the errors inclusive of the systematic matching. Whilst some data retains the shape previously discussed with a sharp increase at low separations, and a shallower linear increase above a threshold, this is not observed for the data from all geometries here. The geometries here which appear to be flatter are also those whose typical errors are dominated by the systematic component. Realistically this is not a useful quantity for further analysis, and it is unclear how the measurements were prepared for analysis in the literature, so focus here remains on the relative matching.




Figure 6.9. Capacitor relative matching considering both random and systematic errors for a variety of array geometries.

6.5.2 Periodicity Relative Matching Statistics

The development of the sinusoidal pattern in the relative matching of capacitors in Figure 6.8 is a very unexpected one, particularly as the random errors that remain after removal of the sinusoid of the original data did not show the same periodicity under analysis. Performing the same analysis with the periodic systematic errors removed still yielded the same results. Yet another unusual aspect of this observation is shown in Figure 6.10 where the analysis is performed on the designs giving the sinusoidal pattern previously, but isolated by wafer. Of the three wafers, for which there was no requested skew or other fabricated difference, wafer B shows an extreme case of the sinusoidal effect, whilst A shows none. Wafer C shows a very weak periodic component. It was also verified that the periodic component was present in approximately the same quantities consistently on the three wafers with respect to the aggregated data shown.

The only conclusion that may be drawn from this discrepancy is that given the wafers did not contain any differently specified BEOL processing (skewed caps, poly, or resistor sets), and that they were all packaged and otherwise assembled identically, this periodicity is yet another component arising from the variations in the fabrication process. To illustrate that this effect is not a periodic component in the random error after the systematic has been removed, such data is shown in Figure 6.11, and the FFT analysis thereof in Figure 6.12 with data isolated by wafer. The periodic component of the same pitch of around $9\mu\text{m}$ in Figure 6.10 is observed for

wafer C. The same cannot be said for wafer D, although the low frequency gradient of the data dominates. The increased noise floor for wafers C and D is explained by the reduced number of capacitors sampled from this wafer compared with that of wafer B. The period of this component is notably the same as the extracted systematic error previously discussed, so is therefore likely to be derived from a similar processing step.

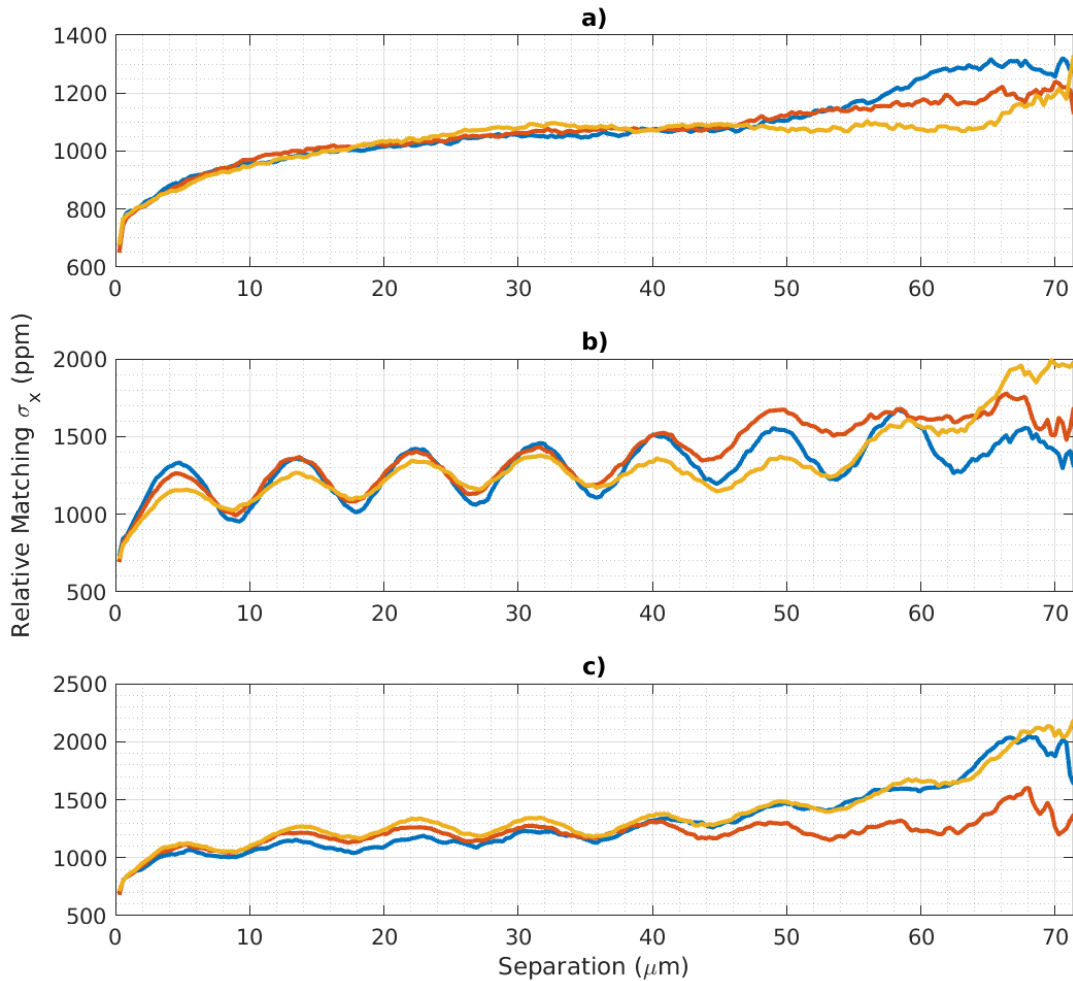


Figure 6.10. Capacitor relative matching against capacitor separation for the minimum spacing geometry arrays, calculated over single wafers. a) Wafer B, b) Wafer C, c) Wafer D. The colours represent each of three samples of an array of the same geometry within the design on each wafer.




Figure 6.11. The random matching of each capacitor within the array by position for the three wafers. a) Wafer B, b) Wafer C, c) Wafer D. No repeatable patterning is observable.




Figure 6.12. FFT analysis of the random matching of each capacitor within the array for the same array over the three wafers. The same periodic component is not observed across all three wafers.

A similar investigation that identified the presence of short range correlated errors in BEOL capacitors is presented in Ref [49]. The scale of separation of capacitors was an order of magnitude larger in the presented measurements than here, but the structures were still fabricated in a 40nm node and a similar relationship to this work identified. As a result of the coarseness of the variations, the inflection points of the short-range correlated errors are difficult to compare. Here it was extracted at around $9\mu\text{m}$, whereas Ref [49] places this point

within the 10-100 μm region. The presented reasoning for the improvements in matching at short distances are related to the CMP metal processing deficiencies, and localised dishing and erosion effects. It is unclear from this work whether the random and systematic contributions to the error are independently considered in the presented analysis, as the precise methodology is unclear. However, the observed trends are still consistent with this work.

6.5.3 Simulations of a Scratch Based Source of Correlated Errors

To better understand the source of the errors presented, simulations were performed through MATLAB to replicate the theoretical impact of types of correlated errors on the capacitor fingers. The primary suspect for a source of correlated errors in the capacitor finger formation is that of CMP scratches [30], [58], [82]. These defects come as the inherent result of a mechanical process. Their effect is limited to the metals and intermetal dielectric surfaces, and whilst intuitively it is hard to conceive of a simple scratch introducing a sizeable error into the capacitor value, the level of error considered is in the parts-per-million region where the slightest defect has a measurable impact on the capacitor value. These scratches can effectively be considered localised reductions in the capacitor finger (or top plate finger) thickness.

A visual representation of the CMP ‘scratches’ considered is shown in Figure 6.13 where many scratches are combined to produce errors in the capacitor shown by their colouring. The scratches are considered to have a defined statistically nominal length S_i , and occur at an angle θ to the orientation of the array. A further representation is shown in Figure 6.14. This results in an effective scratch length of S_{eff} . A single capacitor’s error in this context is the superposition of multiple scratch events of a given weighting. The simulation randomises the positioning of the scratch across a constructed array, and introduces an error to each capacitor value that it touches. As shown, some errors can conceivably extend beyond the occupied area of the array, and some will of course terminate mid-capacitor. These effects are also modelled, with the weighting of the error impacting a capacitor value being proportional to the length of the scratch in its predefined area.




Figure 6.13. Illustration of how CMP scratches may fall on the capacitor array. As shown, they may affect multiple capacitors, giving rise to correlated errors. Each scratch however may be decorrelated.




Figure 6.14. Determination of the effective length of the CMP defect with dependency on the angle of incidence with the array. Scratches may occur at any angle to the array.

The results of the simulation of the impact of the periodic scratches on the relative matching are shown in Figure 6.15. The finite maximum length of the scratches causes a strong decrease in relative matching as the elements are brought closer than the maximum scratch length. This is used to explain the significant improvements in error below $\sim 10\mu\text{m}$, achieved by simulating an error source with nominal scratch length of $9\mu\text{m}$. The larger scale errors are effectively represented by very large-scale scratches of lengths $>100\mu\text{m}$, significantly larger than the distances considered. The impact of varying the capacitor pitch is also observed in Figure 6.15, whereby capacitors on a longer pitch exhibit lower error contributions than those of smaller pitch. Capacitors of longer finger length are also statistically more likely to both be hit by a scratch together rather than individually. Both of these observations are consistent with the measurement data.




Figure 6.15. Simulated recreation of the observed dependency of the relative matching statistic on the capacitor pitch represented by each coloured line. The result is a combination of $9\mu\text{m}$ long randomly arranged scratches, and $100\mu\text{m}$ long scratches.

The periodicity previously presented was also considered in the context of this scratch model. The simulation in Figure 6.16 considers that the scratches do not occur singularly, but that they occur as pairs of scratches separated by a predefined distance, at the same angle and weighting. The result is a periodicity in the extracted relative matching versus separation in Figure 6.10, bearing a similar relation to the real measurements. By analogy, this may imply that there is some scratch source which on some of the measured wafers was producing a correlated series of scratches. For clarity, this may be described analogously to the grit on a piece of sandpaper, where there is a nominal spacing between the grains on the paper which give a fine set of scratches when used of a nominal length, at a nominal pitch.




Figure 6.16. Simulated results for the presence of periodic correlated errors at a pitch of $9\mu\text{m}$. Each error introduced occurs at a randomised position within the array.

Other types of correlated error sources were also trialled in this investigation such as two-point sources of pre-defined separation. It was found that this model produced errors that were too discrete and defined, and did not effectively represent the trends and relationships to the geometry that have been discussed. Whilst the above exercise may only be theoretical, it does provide what is believed to be a reasonable model for the sources of correlated errors, given the limited information available about the process. Other sources such as interference patterns from the lithography process may be considered, but the information is too limited by the foundry to meaningfully evaluate this.

6.6 Sources of Systematic Errors

This is an opportune time for discussion of the systematic errors. In section 5.4.1, large period and systematic errors were observed. These errors were found consistent in their shape, position, and magnitude across in excess of 96 die, all measured wafer reticules, and across three separate lots of wafers. This was the first indication of periodic content in the capacitor errors. The FFT of Figure 5.15 confirms the consistency of the errors, as well as the presence of spatial harmonics to the error content.

Further systematic errors were observed with respect to the large steps in matching found in 5.4.2. By necessity these were removed for further analysis of the remaining errors, however these errors were also consistent across all measured die. Moreover, they were able to be linked to their position on the die, and most notably were consistent when the design was rotated on the reticule. This data held consistent across 3 unique designs of test chip and multiple capacitor geometries, with the die aspect ratios being varied too. This closely links with

an array level patterning effect. While the potential for this to be an OPC effect was discussed in the context of this data alone, this does not hold true due to the consistency across multiple dies, as the OPC is run in a stepped window on the exported die layout. Therefore, the precise positioning of the arrays on the die should result in a variation in the systematic patterning that was observed. A diffraction related lithography issue is far more likely given the data available, particularly the alignment of the periodic patterns to the array and its consistency in position relative to the array despite multiple orientations and positions on the die. It is believed that the pattern is the result of an interference pattern of the light from the lithography resulting from a combination of the very narrow tracks that form the openings within the mask during exposure, and the short wavelength of the light used. This is an effect known to occur on a shorter scale in similar processes, as this effect is responsible for the need for OPC pre-processing for the masks. In order to verify this assumption, more details would be needed to be known about the process, including both the wavelength of the exposure light, likely in the UV region as well as the geometry of the exposure setup, as it is not published whether the process uses a contact or immersed based lithography setup which results in a dramatically different setup for exposure.

The step in capacitor weightings that was observed, however, does appear likely to be related to the aforementioned OPC. This is likely due to the apparent independence of the error's position within the array, with respect to the arrays position on the die. This consistency in position ties it to a die level effect, such as OPC, or other pre-processing effect. As it is confirmed that the die measured come from multiple sites within a single reticule, it is assumed that this is not a reticule level effect. Whether this could be attributed to alignment or stepper issues remains to be seen, although the data to support this is low, due to the position's independence on specific instance within the reticule.

This chapter described a slightly different presentation in the systematic errors. These errors were present in what has been presented as the 'random' errors. It therefore appears that these errors are clearly only random with respect to their absolute position in the array. Being an aggregation of a large number of correlated errors implies that this is unlikely to be due to the lithography process, due to the repeatability in the lithography and etching processes with regards to the patterning, hence why the theory regarding CMP scratching is favoured. CMP scratching, unlike the lithography or stepper effects, comes unaligned at the die and reticule level, due to the polishing occurring at a wafer level. A similar result was explained in an independent study[49] however this paper did not factor the separation of random and systematic matching errors in the arrays. It is this distinction that reaffirms that the CMP process may be to blame. While this study would explain the curve of Figure 6.8, it does not explain the periodicity. Loose media in the polishing compound may be to blame for the original relationship, however scratching from a well-defined larger structure as discussed in section

6.5.3 is required to explain the formation of the periodic patterning in the relative matching relationships. It is still unclear why these periodic patterns only occur at minimum geometry, however.

To summarise: The periodicity found in the systematic errors is likely to be from interference patterns in the lithography. Steps in the matching across the die may result from OPC or other mask pre-processing windows, and the correlated errors in the relative matching performance of the capacitors is likely to arise from the hard media used during the CMP polishing of the metal surfaces.

6.7 Conclusion

The wider scale considerations for the matching of capacitors in larger arrays have been discussed. Dummying has been identified as a significant contributor to the introduction of errors into the capacitor array if it is not well utilised in layout. A typical requirement of $25\mu\text{m}+$ of top and bottom array dummying is identified for full removal of systematic errors, while in excess of $10\mu\text{m}$ more may be required to also remove the random errors that were observed to significantly increase for those capacitors that are near the ends of the array. Short-range correlated errors are also found have a strong presence, presenting another avenue for further improvement in capacitor matching. Large improvements in random matching can thereby be achieved by placing capacitors within groups of separation of less than $10\mu\text{m}$. This is in addition to the existing improvement in systematic matching which is broadly utilised in modern array layouts. This is achieved through analysis of the relative matching of capacitors and captures errors that are likely derived from the CMP steps during fabrication. The model is verified through basic simulations, and through analysis of the periodic trends linking the observed matching errors and their resultant performance in the array to the underlying process a designer has to work within. The threads of investigation regarding the systematic errors are brought together in the final section of this chapter, and mark a significant leap forward towards proper understanding of the nature of the capacitor errors, so that the errors themselves may be worked around. This improvement will now be described further in Chapter 7.

Chapter 7 Simulated Performance Impacts of Improved Matching in a Behavioural 18-bit SAR ADC

7.1 Introduction

This chapter presents a behavioural test-bed for evaluation of the achievable performance improvement that may be attained using the improved structures and observations that are analysed throughout this thesis. The test-bed comes in the form of a model of a single-stage, 18-bit, flash-assisted, charge-redistribution SAR ADC quite typical in design for a 40nm ADC. This reflects back to the research objectives and questions, demonstrating that when the observations of the previous chapters are considered during the array design, many circuits may achieve significant improvement in their performance above that of the foundry provided cells, and matching models. The choice of an ADC to demonstrate the performance improvement is important due to its fundamental sensitivity of to the capacitor matching. Conventional matching improvement techniques are also demonstrated and analysed in this chapter.

7.2 Behavioural ADC Architecture

The test-bench ADC of Figure 7.1 is a single-stage, flash-assisted, 18-bit SAR ADC operating at 100MSPS using a charge redistribution DAC. The architecture utilises several techniques in order to achieve the wide dynamic range. This is first done through implementation of 4 bits of shuffling on the MSB capacitors, necessitating use of segmentation which forms the majority of the DAC capacitors. The ADC is also flash assisted, with the flash ADC quantising the input signal to 4-bit level, and passing the result to random barrel, or round-robin, shuffler. This was found to be sufficient in shuffling the capacitor combinations between samples so as to linearise the mismatches between the thermometer MSB DAC. Dither is also integrated with 9 bits of range starting from the LSB capacitor. This would usually require the addition of capacitors specifically for introducing dither, however this implementation was performed by sampling the dither values on the lower 9 LSBs of the array, as opposed to the sampling of the input voltage. Whilst this does attenuate the sampled input voltage, the impact was limited to around 6% of attenuation, so was a tolerable penalty given the significant reduction in capacitors required through use of this technique.

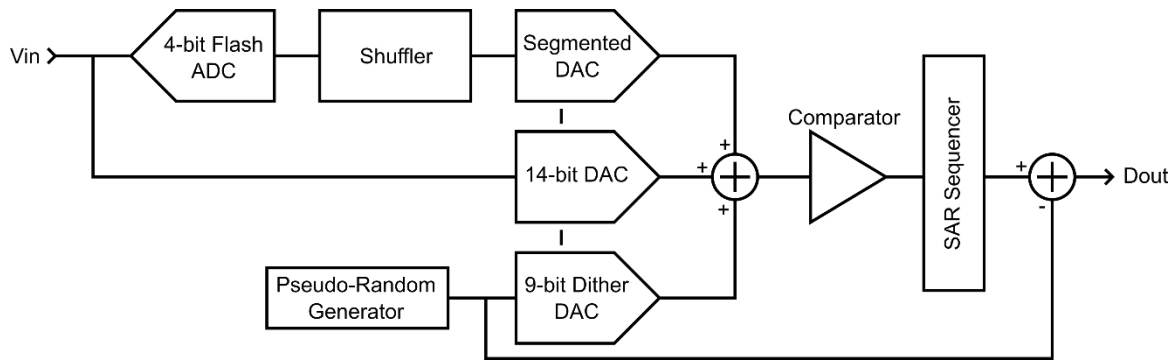


Figure 7.1. Block diagram of the behavioural 18-bit flash-assisted SAR ADC, utilising both shuffling of the 4 MSBs and 9-bit dither of the LSBs.

At the core of the design is the CDAC shown in Figure 7.2. This design of CDAC is split into four distinct sections to allow for integration of the aforementioned techniques, and to reduce the total number of unit capacitors to an implementable level as shown in Table 7.1. This is done primarily to conserve area on the die, but also to reduce the incurred systematic error, as this thesis has previously established that significant reduction in capacitor matching of all types may be achieved through reducing the physical size of the capacitor array. To achieve this the array is split into 4 distinct sections: The 4 MSBs of the DAC (B17-B14) are implemented as thermometer coded bits. This is to enable the use of shuffling on them, and enhance the linearity. The following 5 bits (B13-B9) are straight binary weighted bits, followed by a 5-bit fractional reference section, and a 4-bit sub-DAC. Both the fractional reference section and the sub-DAC operate under the fractional reference.

To reduce the required dynamic range of the sub-DAC without the DAC MSB becoming excessively large, a fractional bit scheme was implemented. Fractional bits may be achieved through a number of means. Most commonly, this is by use of a sub-DAC, or by fractional finger capacitors [23], [38], [83]. The issues with sub-DACs have been discussed, introducing a gain error of approximately 3-5% to the bits implemented and limit their usefulness to approximately 4 or 5 bits without calibration due to this. A purely binary weighted implementation of an 18-bit SAR ADC would require 262,144 unit capacitor elements. Use of even a 5-bit sub-DAC would only reduce this figure to 8224, which for the capacitors analysed in this section would give a capacitor array at least 2.3mm long. This is larger than the typical die size of a fully integrated high-precision ADC.

The physical implementation of a fractional reference scheme was chosen as a more effective alternative. The capacitors are identical to that of the straight binary weighted section as shown in Table 7.1, however the positive reference used for switching is a derived $1/32^{\text{nd}}$ of the main ADC reference. The main ADC reference used was 1.1V, and so the fractional reference was 34.375mV. While this is a rather aggressive division of the reference, given the matching between two capacitors or even resistors, a capacitive or resistive divider is easily implemented

to derive this reference with far superior matching than that is achievable using only a sub-DAC. For example, assuming the unit capacitor only achieves 1200ppm total matching error for both systematic and random components combined. This easily translates to less than a 0.02% gain error for the affected portion of the DAC. In reality, the reference buffer offset is likely to introduce more error than this and thus be the dominant error contributor, although many options for self-calibration or chopping of the amplifier are available. The sub-DAC is also operated on the fractional reference scheme in order to prevent the need for a very large sub-DAC coupling capacitor. In an ideal design, the shuffled and dithered sections of the DAC would be extended such that every single capacitor experienced some shuffling or dithering in order to maximise the linearity of the DAC. This has already been shown to be a highly effective approach to high-linearity ADCs [10]. However, implementations of this level, particularly in a single-stage SAR, would be rather extreme due to the introduced bit-propagation delay introduced by the shuffling algorithm and block, as well the respective dither coders/decoders. Whilst the impact of these blocks are not modelled in this behavioural ADC, the levels at which these techniques are implemented are kept reasonable for implementation on a practical 100MSPS ADC design, without these propagation delays affecting performance.

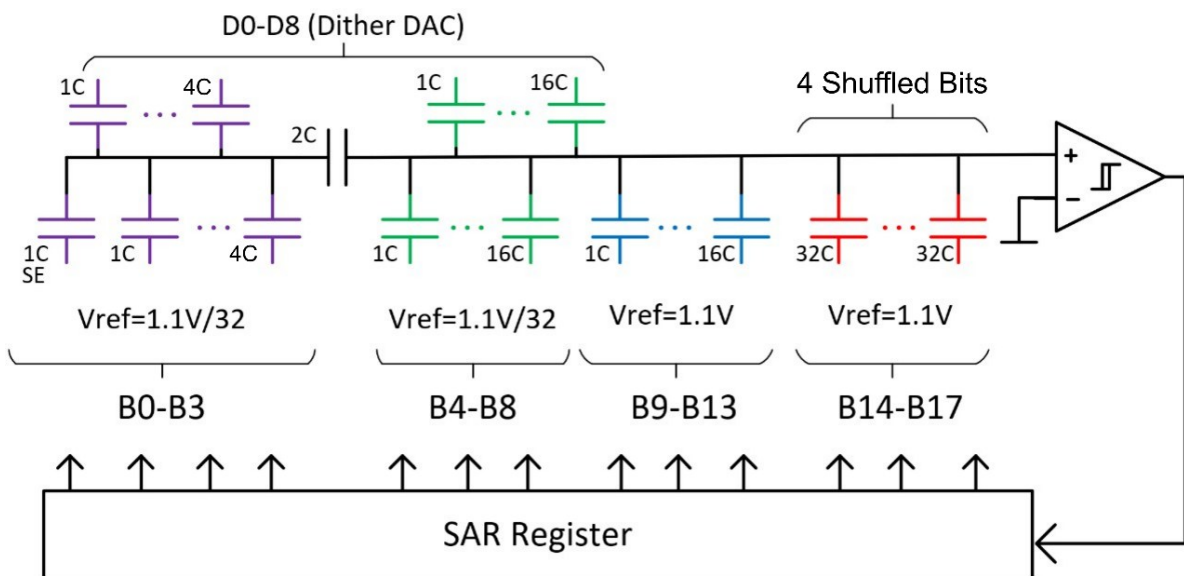


Figure 7.2. Overview of the chosen straight 18-bit SAR ADC CDAC architecture, with the four main DAC sections and their relative capacitor weightings highlighted.

Table 7.1. Each bit in the designed behavioural ADC and its respective bit weightings. This allows conversion to a total of 262,144 codes with only 550 unit capacitors without use of fractional capacitor fingers.

Section	Bit	Unit Capacitors	Reference (V)	Respective Bit Weighting
Segmented/ Thermometer Coded - Shuffled	B17	15x32 (Shuffled)	1.1	65536
	B16		1.1	32768
	B15		1.1	16384
	B14		1.1	8192
Straight Binary Weighted	B13	16	1.1	4096
	B12	8	1.1	2048
	B11	4	1.1	1024
	B10	2	1.1	512
	B9	1	1.1	256
Fractional Reference – Binary Weighted	B8	16	0.034375	128
	B7	8	0.034375	64
	B6	4	0.034375	32
	B5	2	0.034375	16
	B4	1	0.034375	8
Sub-DAC	B3	4	0.034375	4
	B2	2	0.034375	2
	B1	1	0.034375	1
	B0	1	0.034375	0.5

Redundant bit-trials were also introduced for bit trials B8 and B12 to account for some level of correction of accumulated errors from the capacitor mismatch. When capacitor errors are considered, the shuffler produces additional dither in the converter, improving the linearity further for the central 7/8th of the input voltage range. Outside this range, the segmented DAC is at maximum or minimum code, and thus there is no shuffling occurring, and no additional linearisation.

7.3 Simulation Methodology

The metrics used to evaluate the performance of the converter in the presence of capacitor mismatch are the SNR, INL, and DNL. The SNR is unaffected by the use of any of the linearisation techniques implemented in this ADC aside from the redundancy. The shuffle and dither techniques do not affect the SNR beyond a slight noise degradation, but a significant improvement is expected in the INL and DNL of the ADC as a direct result of these techniques.

For a fair comparison, both the foundry and custom capacitor evaluations are performed using the exact same ADC model, with the only difference being the capacitor values. The assigned capacitor values for the CDAC come directly from the silicon measurements of the capacitor errors, as will be described. To evaluate the SNR of the ADC for a given assigned DAC, a near-full-scale low frequency sinewave is input to the ADC. This signal is also ideally driven and therefore no kickback or settling from the ADC is considered. The shuffle and dither blocks are fed with a randomised pattern each conversion and large number of samples taken (>4096). From this, an FFT is derived from the ADC output code from which the SNR is calculated.

The linearity of the ADC is massively more complex to evaluate in simulation. Due to resource constraints, the ability to queue up over 1 million simulations, each lasting around 5-10 minutes was not feasible. As a result, the usual statistical method for evaluating the INL and DNL of high-precision ADCs is also not feasible [31], [36], [73]. Instead, the ADC errors must be evaluated at discrete, predefined input voltages with either a DC input or a slow-moving ramp, optimising the simulations to be fewer in number, but of considerable computational cost to run. The approach to extract the performance therefore is taken through two stages. First the overall linearity of the ADC is determined by running (in parallel) 8192 simulations with a slow ramping input to the ADC. This is configured to result in input conversion steps of $\frac{1}{4}$ LSB, resulting in precision in the INL/DNL that can be trusted to this level. Each simulation operates with a different dither and shuffle code combination, and the resulting ADC codes are then averaged across the 8192 runs for each input voltage to give the oversampled linearity plot. This is performed for a single measured DAC of a given design.

To extract the distribution of the INL and DNL values across multiple measured dies, the worst-case sections of the overarching INL and DNL profiles are selected, and higher resolution sweeps run in this range. These are referred to as 'short-codes' and massively reduce the time to evaluate 96 of the available arrays. These are implemented as fixed-input simulations with the dither and shuffle codes cycled between conversions to ensure full dither and shuffle code coverage, and the final result an average of the conversion results. This method reduces the time to evaluate DAC layout performance in the ADC model from around 3 weeks for a single design, to around 3 hours to evaluate all 96 permutations.

7.4 Incorporation of Measured Arrays for Simulation

The key requirement is to use the real measured capacitor errors from the test chip arrays directly in the simulation. This ensures simulation results directly reflect the capacitor matching error's influence on the potential real-world performance of the SAR ADC. To do this, the matching errors for arrays of 572 unit capacitors are required. These arrays are generated

through extrapolation of the real measured data of the arrays, as the measured arrays have only 256 capacitors. The method chosen for extrapolating the arrays, was to mate three measured arrays from three different die for the selected capacitor geometry together. This was performed by alignment of the means of the final 16 capacitors in the array, such that any low order gradient or periodic components in the systematic error would be extended, as this was expected to be the primary source of error in such a large array given the previous data. To maintain consistency in the same way as the matching statistics were calculated, the array data that was used for extrapolation has the previously addressed ‘step’ removed.

The measured foundry capacitor arrays were considerably shorter than the custom capacitor arrays, so required a large number of extrapolations to achieve a sufficiently large array. The arrays were extrapolated as such to continue the large periodic pattern that was previously noted, with simulated dummyming added to remove the large increases in systematic error at the ends of the array. This extrapolation method was chosen as to ensure that for all capacitor geometries, when layout folding was applied it would demonstrate the attainable improvement through not just the systematic error cancellation, but also the random correlated error improvements that were alluded to in section 6.5. The capacitor geometry chosen and demonstrated in the proceeding simulations was that shown in Table 7.2. Other geometries were also trialled and displayed the same relationships shown, however minimisation of the random matching performance was found to be paramount in achieving a high-performing DAC in this ADC, as opposed to minimisation of the systematic content.

Table 7.2. Details of the capacitor geometry chosen for demonstration of the achievable SNR and linearity improvement of a behavioural SAR ADC. This was found to be the optimum geometry for this size of array and 16th order folding.

ID	C (fF)	S (nm)	W (nm)	L (μm)	N_m	S_e (nm)	Number of DACs	σ_R (ppm)	σ_S (ppm)
DAC104	22.7	120	90	43.5 μm	4 (M2-5)	140	96	495	265

7.5 Impact of Capacitor Interdigitation and Folding in Layout

Conventional theory suggests that layout folding schemes applied to the DAC capacitor array are only present to cancel out the gradient and other systematic effects from contributing significantly to the total DAC errors. Each further order of folding applied as shown in Figure 7.3 creates an additional symmetry point to the array layout and further distributes the elements from each capacitor more evenly throughout the array. This is achieved by splitting each lumped set of elements in half, and then mirroring the structure. For example, a capacitor of 8 individual

fingers may be split into two capacitors each with 4 fingers, and then interdigitated at that level. The result of this is for each lumped capacitor to be more evenly affected by the large acting gradients, with higher order folding capable of averaging out higher order, more complex, gradient profiles[36], [65], [84].

To evaluate the required amount of folding necessary to produce the desired 16-bit linearity and increase the mismatch-limited SNR, folding of varying orders is applied to the segmented capacitor's weightings for all available array data. SNR simulations are then performed. The expectation was that the SNR will increase with increasing folding orders and stabilise at the SNR achievable given the random matching figure that the chosen capacitor array should achieve. This array was extracted to provide a σ_R of 495ppm which theoretically limits performance to 96 dB, with reference to Figure 2.7 assuming there is perfect systematic cancellation. As each segmented capacitor is comprised of 32 unit-capacitors, a maximum folding order of 16 was possible. The specific folding implementations used within simulation are shown in Figure 7.4. The non-segmented sections of the DAC were placed centrally to the overall layout, and simulated with a similar folding system.




Figure 7.3. Example of the common-centroid folding principle to generate more points of symmetry, to reduce the effective mismatch error of each grouped capacitor.




Figure 7.4. Further visualisation of folding schemes of increasing order. Schemes a) through f) illustrate folding order 0, 1, 2, 4, 8 and 16x respectively.

Simulations are performed using measured data from all available dies with the various folding orders applied. This gives a figure for both mean achievable SNR, as well as the spread of the achieved SNR, which is represented in Figure 7.5. Firstly, a general increase in SNR is observed

for increasing orders of folding, along with a general decline in spread of achieved SNR, σ_{SNR} , due to the improved matching attained. Due to the specific nature of the gradient, a significant increase in cancellation occurs from folding order 0 to 1, resulting in a large increase in SNR, which is not cancelled quite as well by second order folding. Beyond this however, a consistent 4dB SNR improvement is observed with increased folding. This improvement continues beyond the 96dB theoretical limit for the 495ppm random matching figure observed. This result implies that the impact of folding goes beyond complete gradient cancellation for extreme cases, and may be enabling the utilisation of localised, correlated errors, to achieve better matching performance, and hence superior SNR.




Figure 7.5. The relationship between the behavioural ADC SNR performance with varied folding orders in layout, as well as the spread of the SNR over all measured dies.

Referring to the data previously obtained in section 6.5 the impact of the folding may be matched to the relative decrease in random errors in the array. Table 7.3 details the decrease in relative matching as the cluster size (the distance containing one instance from each capacitor in Figure 7.3) is reduced through higher order folding. The source of this reduction is observed in Figure 7.6, where despite the exponential decrease in separation (a.k.a. cluster size) the magnitude in the reduction of error is maintained. This is consistent despite the sizings spanning 2 orders of magnitude due to the presence of correlated errors over small distances. This results in the improvements in SNR observed for higher folding orders as the errors become correlated and affect multiple capacitors at once, removing the net error.

Table 7.3. The distance containing an instance of every grouped capacitor at each folding order, and the resulting relative matching improvement of the capacitors from Figure 7.6.

Folding Order	'Cluster' Size (μm)	Equivalent σ_{rel} from Figure 6.8 (ppm)	$\Delta\sigma_{rel}$ (ppm)
0	188.16	-	-
1	94.08	822.4	-
2	47.04	739.1	143.3
4	23.52	620.8	118.3
8	11.76	543.5	77.3
16	5.88	459.6	83.9




Figure 7.6. Net reduction in capacitor matching errors as the maximum capacitor separation is reduced by increasing the order of folding in the array layout from 2x through 16x.

It is clear from the simulation results that there is significant benefit to be obtained through the use of higher order folding schemes where the implementation allows. For higher order folding, whilst the exact structure of the DAC may remain the same, the placement of the drivers and their routing between the capacitors can pose significant design challenges in layout. This is both due to the fine capacitor pitch and long interconnect distances. This introduces larger

bottom plate parasitics due to increased routing complexity, as well as increased driver input parasitics, and loading. For the purposes of the following simulations in this chapter a 16x folding scheme is to be used to demonstrate the potential improvements that combining of the above techniques with a more optimised capacitor array may bring to the ADC's performance. By comparison, in this instance of a capacitor array design that does have a particularly large systematic contribution, a blind trim with no folding applied only results in an increase of 2dB in the SNR, supporting that careful design may be much more effective than post-fabrication correction techniques.

7.6 Simulated ADC Performance

7.6.1 Linearity

The overall shape of the INL and DNL of the behavioural ADC for a sample array is presented in Figure 7.7, with simulations performed without thermal noise. Figure 7.7a shows the overall trend of the INL where errors from the binary weighted section of the DAC are reproduced 16 times across the profile due to the 4-bit segmentation as expected. It is clear that the shuffling is functioning due to the lack of any large lower order components visible in the INL. The dither can be seen to be working within this profile, resulting in an interpolation of the smaller DAC errors below ~9-bit level. The dither primarily results in the improvements in the DNL profile, where no systematic steps are observable. The thermal noise contribution in these linearity simulations also adds to the dithering effect, and results in the clear DNL profile seen. Without the thermal noise present, the DNL would be seen to be overall lower, and the upper and lower 1/16th sections would be larger than the central portion of the range, due to the lack of shuffling present. This particular DNL profile indicates that this array achieves 19-bit monotonicity, and 17-bit linearity. For these purposes the monotonicity is defined as commonly $|DNL| < 1\text{LSB}$ [31] and linearity as $|INL| < 1\text{LSB}$. The full INL run was not able to be repeated on the foundry capacitor measured structures due to the period of time required to simulate. Despite efforts to optimise this the full process totals around 3 weeks.




Figure 7.7. Full a) INL, and b) DNL profiles from a randomly selected array, with thermal noise present in the simulation.

The DNL results of the short-code simulations without thermal noise performed on each available measured array for the geometry described in Table 7.2 are shown in Figure 7.8. The worst case DNL observed over simulated arrays was 0.27ppm, equating to 0.07 LSBs at 18-bit level. This indicates that the ADC design if extended might be capable of achieving 21-bit linearity and that the linearisation techniques work effectively as the variation of the maximum DNL is very low at only 0.12ppm. The INL distribution in Figure 7.9 of the ADC is also incredibly promising, averaging approximately 3.67ppm. This equates to 1LSB at 18-bit level. The variation between measured arrays is also exceptional, giving a maximum INL of 15ppm, which is still suitable for 16-bit linearity. Overall, all dies were able to achieve maximum non-linearities below 16-bit level uncalibrated. This level of linearity uncalibrated is typically only achieved in the literature due to the matching requirements of the DAC which has been greatly improved here to achieve this.




Figure 7.8. Maximum DNL of the behavioural ADC using real array measurements for capacitor errors. Data is accumulated through short-code simulations of 96 measured arrays.




Figure 7.9. Maximum INL of the behavioural ADC using real array measurements for capacitor errors. Data is accumulated through short-code simulations of 96 measured arrays.

7.6.2 Dynamic Performance

The SNR gives a good indication of the dynamic performance of an ADC and is used as a point of comparison between the foundry capacitor structures and the optimised structures of this thesis. The presence of matching errors in the DAC, even with the use of linearisation techniques as has been discussed, will degrade the SNR. For these simulations no thermal noise was considered.

The results for both the optimised capacitor structure as well as the foundry capacitors are shown in Figure 7.10. The mean achieved SNR of the behavioural ADC, uncalibrated is 16.19 dB

higher for the optimised capacitors than for the foundry capacitors. The thermal noise of the 6pF DAC in this simulation would in theory limit the SNR to 96 dB if thermal noise is introduced, meaning that performance of all foundry capacitor DACs for this architecture are capacitor mismatch limited. This is not the case of the optimised capacitors, as they easily achieve a nominal 99.34 dB SNR, with a mere 7% of arrays being mismatch limited, and the rest limited by thermal noise. This demonstrates that by purely optimising the capacitor structure and geometry for their matching performance, a 16-bit linear ADC with thermally limited performance is quite achievable, and that the use of the use of dither and shuffling linearisation techniques are enhanced through the use of said capacitor arrays.




Figure 7.10. Comparison of SNR performance using this work's optimised capacitor geometry, compared with the foundry capacitor's matching model.

7.7 Conclusion

This chapter has provided an overview of an 18-bit behavioural SAR ADC, and utilised it to benchmark the performance of the measured optimised capacitor geometry against the measured foundry capacitor structures. This was achieved through direct introduction of the in-silicon measured capacitor array errors from 96 die into the simulated ADC model. The simulations demonstrated that by use of the improved capacitors, the extracted INL reliably achieved 16-bit linearity and DNL in-excess of 20-bit monotonicity. This was achieved through moderate implementations of standard dither and shuffling linearisation techniques already in

common practice. This work overall demonstrates that the matching improvements are considerable enough as to result in an SNR not limited by the mismatch of the DAC capacitors, and an improvement over the foundry cell capacitor more than 16dB. The above-expected increase in SNR with increased orders of layout interdigitation and folding also demonstrates the exploitation of correlated errors. This results in superior matching performance of capacitors at high folding orders. This chapter has also shown that the effectiveness of conventional matching theories (dummying, shuffling, dithering) are only enhanced by the use of the optimised capacitor geometry. Without the improvements and understanding presented in this thesis the gains are only minimal, and leave significant untapped potential within the process for higher performing circuits.

Chapter 8 Conclusion and Outlook

This thesis sought to quantify and examine the matching performance of laterally coupling BEOL capacitors in low geometry nodes, and whether opportunities for optimisation existed within the constraints of the process. An optimised capacitor geometry has been designed and refined to enable such characterisation and has resulted in new knowledge about the nature of the capacitor errors. This has been investigated into how they may be exploited within the boundaries of a low geometry node utilising the copper damascene process to achieve more than a 4x matching improvement. This has then been demonstrated to equate to more than 16dB potential improvement in the performance of an 18-bit SAR ADC.

The performance of CDACs in SAR ADCs are frequently limited by the errors developed during fabrication of the capacitors in the capacitor array. This is particularly true in recent developments, due to the complexities introduced by use of smaller geometry processes such as the 40nm node and below. Use of low-k dielectrics and copper interconnects for reduced parasitics, plus the requirement for the capacitors to be implemented exclusively in the interconnect metals are largely responsible. This has necessitated this investigation in order to better understand the and quantify the possible improvements that may be made.

The relationship between the key capacitor geometry and the resulting capacitor values has been analysed, showing proportionality of capacitance to both finger length and number of metal layers. The lateral coupling capacitor was determined to derive approximately half of its capacitance from the fringing fields, leading to a reduced relationship with the finger spacing, and an increased dependency on the finger width. This allowed identification of the processing steps, such as the copper dual-damascene process, which are likely to determine the value of the capacitor and therefore the matching errors of the capacitors.

These effects have been investigated further in the context of the matching statistics. A high-precision charge-based measurement method was developed for independent extraction of the systematic and random matching errors for analysis. The random matching components have been found consistent with a $1/\sqrt{A_{Top}N_m}$ relationship and are not directly correlated with the capacitor value as is usually portrayed by popular literature. Significant periodic and step components were analysed in the systematic matching errors for capacitors near the minimum permitted geometry. Providing that the designer avoids the minimum geometry, and the locations specified on the die that have been noted in this study, these two systematic effects may be avoided.

Array-level matching considerations were investigated, with shielding of the array found to have less impact on matching than initially expected. A requirement of more than $25\mu\text{m}$ of capacitor dummy structures at each end of the capacitor array has been established to remove large contributions of both systematic and random errors. The relative matching of capacitors with a variety of separation distances has been evaluated. This identified a strong presence of correlated errors that can be exploited in practical DAC designs to reduce the random errors that would otherwise require calibration to remove. Such errors are subsequently theorised to link to the copper polishing process resulting in correlated scratches in the metal surface. This is a key development in the understanding of the nature of the capacitor errors, particularly this study's identification of this error to be from what were previously thought to be purely random errors.

The performance increases due to the optimisation of the capacitor structures has been demonstrated in a test-bed simulation of a behavioural 18-bit SAR ADC utilising a fractional reference scheme to reduce the requirement for a large sub-DAC. The simulations use the results of a total of 96 die from experimental capacitor array measurements. Overall, the optimised ADC achieves more than 16dB SNR improvement compared to an ADC using the foundry capacitors. The array is also sufficiently linear to achieve the performance required for an uncalibrated 16-bit SAR ADC. This achievement implies that full implementation of the optimisation procedures described in this thesis would reduce the capacitor matching error to below the influence of thermal noise. It will instead leave the latter as the main limitation in CDAC performance for uncalibrated SAR ADCs required in high-speed and precision instrumentation applications.

This study has made significant progress towards its goals of understanding of the capacitor errors. The design of the test chips and the methods of characterisation have enabled capturing of quality data that has identified several presentations of systematic errors in the capacitors. These presentations have been linked back to the processing steps that are responsible for their creation. This allowed demonstration of the potential improvements of the capacitors in the SAR ADC. The next stages in this study, were time and resources available, would be a more focused investigation into the BEOL metal formation itself. It is clear from this wider capacitor-level study, that there are several non-linear mechanisms at play that complicate the resulting capacitance of the overall capacitor structure. These link back to dependencies of the metal and dielectric formation on the precise geometry of the capacitor. A more focused study would investigate the metal geometry and its dependencies using FIB and SEM techniques to identify correlations between the metal cross-section and the matching performance. Any investigation into the low-k dielectric would be particularly challenging as the exact nature or material of it is unknown and is likely to remain a closely guarded secret. Such secrets also affect the other

interesting avenue for investigation: The interference patterning arising from the lithography stages. Any analysis here would require more information on the exact setup using for the process' lithography, which will certainly unpublished for the lifetime of the process. The more general copper damascene process is well established and, as such, proves a more fruitful and applicable avenue for further investigation.

With the removal of the capacitor matching errors as the bottleneck of performance for uncalibrated SAR-ADCs, other avenues for improvement could also be pursued without the added burden of the previously dominant mismatch errors. Dielectric absorption is one issue which would likely become a main error source in high-speed ADCs (>100MSPS), typically manifesting as a memory type effect in the DAC. They may also be considered as frequency-dependent mismatch errors which change with sampling frequency, making them challenging to compensate for. A large portion of the work of this thesis may be applicable to this endeavour, as the structures and methods used for the measure of the capacitor errors can be reused for the measure of the dielectric absorption, subject to some optimisation for the higher switching rate that would be necessary. The level of precision would also not need to be as high as this study, due to the expected magnitude of the errors due to dielectric absorption in this process (>1000ppm). This would bring benefit and allow optimisation of the existing calibration schemes by more specific targeting of the compensation mechanisms. This is as standard DAC calibration routines are generally unable to handle the memory effect presented by dielectric absorption, and often 2nd order schemes or more are necessary.

Ultimately, the power and area previously spent on addressing mismatch may then allow two-stage or pipeline architectures with a higher dynamic range per stage, overall increasing the precision of the ADC, as well as potentially increasing the speed. The ease-of-use of the optimised structures presented act as near drop-in replacements for existing DACs to remove the matching performance limitation.

Appendix A Tables of Capacitor Variations

Table 8.1. Capacitor matching performance figures for pairs of capacitors from the TC2 Tapeout.

DAC ID	EN ID	S (nm)	W (nm)	L (μm)	Value (fF)	σ_S (ppm)	σ_R (ppm)	σ_T (ppm)	Description	Dummying	Metals
1	N/A	70	70	N/A	40.00	2868	1872	3404	Replica of tc1 DAC (40f caps)	5 μ dmy	M2-M5
3	0	70	70	14.5	10.18	910	733	1186	Spacing Sweep	15-25 μ dmy ↓	M2-M5 ↓
4	1	80	70	14.5	9.69	557	696	896	Spacing Sweep		
5	4	90	70	14.5	9.46	432	665	799	Spacing Sweep		
6	5	100	70	14.5	9.20	666	634	923	Spacing Sweep		
7	8	110	70	14.5	8.87	409	604	731	Spacing Sweep		
8	2	120	70	14.5	8.59	436	597	752	Spacing Sweep		
3	0	70	70	14.5	10.18	910	733	1186	Width Sweep		
9	6	70	90	14.5	9.86	515	648	830	Width Sweep		
10	3	70	110	14.5	9.40	369	609	715	Width Sweep		
6	5	100	70	14.5	9.20	666	634	923	Width Sweep		
11	7	100	90	14.5	8.51	416	612	744	Width Sweep		
12	9	100	110	14.5	7.87	459	590	764	Width Sweep		
23	23	70	70	14.5	9.43	638	659	919	No side Dummying		
24	22	90	70	14.5	9.39	435	647	781	AP and Poly shielding		
25	21	90	70	14.5	9.38	449	661	800	Poly Shielding		
26	20	90	70	14.5	9.42	496	642	817	AP Shielding		
31	24	90	70	29.0	18.02	439	505	673	Extended		
32	28	70	70	14.5	10.13	872	761	1161	DA Spacing Sweep		
33	29	100	70	14.5	9.17	629	631	894	DA Spacing Sweep		
34	30	120	70	14.5	8.59	569	605	876	DA Spacing Sweep		
35	25	90	70	19.0	12.24	555	593	837	Extended		

Chapter 8

36	26	90	70	24.0	15.25	449	534	709	Extended		
37	27	90	70	29.0	18.12	629	510	733	Extended (No side dummy)		
38	31	90	70	14.5	9.47	689	624	938	Retracted vias		

Table 8.2. Capacitor matching performance figures for pairs of capacitors from the TC3 Tapeout.

DAC ID	EN ID	S (nm)	W (nm)	L (um)	Value (fF)	σ_S (ppm)	σ_R (ppm)	σ_T (ppm)	Description	Dummying	Metals
100	4	100	90	43.5	23.05	436	545	687	Default cap	30 μ dmy ↓	M2-M5 ↓
100	8	100	90	43.5	23.05	252	562	614	Default cap, offset 1 finger		
101	10	120	90	43.5	21.15	305	513	596	Finger spacing sweep		
102	11	140	90	43.5	19.60	400	474	617	Finger spacing sweep		
103	5	90	100	43.5	23.67	456	610	892	Finger width sweep		
104	6	90	120	43.5	22.73	265	495	560	Finger width sweep		
105	7	90	140	43.5	22.43	450	509	678	Finger width sweep		
106	0	100	90	14.5	8.66	437	801	884	Finger length sweep		
107	1	100	90	29.0	16.35	681	616	1022	Finger length sweep		
108	2	100	90	58.0	28.05	340	511	610	Finger length sweep		
109	3	100	90	71.5	31.83	376	438	567	Finger length sweep		
110	12	100	90	43.5	23.08	402	559	635	Finger end spacing 80 nm		
111	9	100	90	43.5	23.01	728	537	984	Finger end spacing 200 nm		
113	13	100	90	43.5	27.63	573	463	796	M2, M3, M4, M5, M6 only		M2-M6
114	14	100	90	43.5	23.39	461	568	742	M3, M4, M5, M6 only		M3-M6
115	15	100	90	43.5	12.10	936	978	1200	M4 and M5 only, M3 shielding		M4-M5
116	16	100	90	43.5	17.64	726	852	1007	M4, M5 and M6 only, M3 shielding		M4-M6
117	17	100	90	43.5	12.53	673	1027	1099	M5 and M6 only, M3 and M4 shielding		M5-M6
119	18	100	90	43.5	24.30	649	788	963	Only 5u of top dummying	5 μ dmy	M2-M5 ↓
124	19	100	90	43.5	38.73	882	572	1027	All caps 2 fingered	30 μ dmy	

References

- [1] W. Kester, ‘Understand SINAD, ENOB, SNR, THD, THD + N, and SFDR so You Don’t Get Lost in the Noise Floor’, Analog Devices Technical Articles. [Online]. Available: <https://www.analog.com/media/en/training-seminars/tutorials/MT-003.pdf>
- [2] B. Murmann, ‘Introduction to ADCs/DACs: Metrics, Topologies, Trade Space, and Applications’, in *IEEE International Solid-State Circuits Conference (ISSCC)*, 2022. doi: 10.1109/ISSCC42614.2022.11005942.
- [3] Boris Murmann, ‘ADC Performance Survey 1997-2024’. Accessed: Jan. 03, 2025. [Online]. Available: <https://github.com/bmurmann/ADC-survey>
- [4] ‘Understanding SAR ADCs: Their Architecture and Comparison with Other ADCs | Analog Devices’. Accessed: Aug. 31, 2025. [Online]. Available: <https://www.analog.com/en/resources/technical-articles/successive-approximation-registers-sar-and-flash-adcs.html>
- [5] J. H. Mueller, S. Strache, L. Busch, R. Wunderlich, and S. Heinen, ‘The Impact of Noise and Mismatch on SAR ADCs and a Calibratable Capacitance Array Based Approach for High Resolutions’, *International Journal of Electronics and Telecommunications*, vol. 59, no. 2, pp. 161–167, Jun. 2013, doi: 10.2478/eletel-2013-0019.
- [6] E. Alvarez-Fontecilla, P. S. Wilkins, and S. C. Rose, ‘Understanding High-Resolution Dynamic Element Matching DACs [Feature]’, *IEEE Circuits Syst. Mag.*, vol. 23, no. 4, pp. 34–43, 2023, doi: 10.1109/MCAS.2023.3325504.
- [7] Kok Lim Chan, N. Rakuljic, and I. Galton, ‘Segmented Dynamic Element Matching for High-Resolution Digital-to-Analog Conversion’, *IEEE Trans. Circuits Syst. I*, vol. 55, no. 11, pp. 3383–3392, Dec. 2008, doi: 10.1109/TCSI.2008.2001757.
- [8] E. Alvarez-Fontecilla and P. S. Wilkins, ‘Linearity Through Democracy [Feature]’, *IEEE Circuits Syst. Mag.*, vol. 25, no. 1, pp. 58–69, 2025, doi: 10.1109/MCAS.2024.3495566.
- [9] R. T. Baird and T. S. Fiez, ‘Linearity enhancement of multibit $\Delta\Sigma$ A/D and D/A converters using data weighted averaging’, *IEEE Trans. Circuits Syst. II*, vol. 42, no. 12, pp. 753–762, Dec. 1995, doi: 10.1109/82.476173.
- [10] J. Steensgaard, R. Reay, R. Perry, D. Thomas, G. Tu, and G. Reitsma, ‘A 24b 2MS/s SAR ADC with 0.03ppm INL and 106.3dB DR in 180nm CMOS’, in *2022 IEEE International Solid-State Circuits Conference (ISSCC)*, San Francisco, CA, USA: IEEE, Feb. 2022, pp. 168–170. doi: 10.1109/ISSCC42614.2022.9731652.
- [11] I. Galton, ‘Why Dynamic-Element-Matching DACs Work’, *IEEE Trans. Circuits Syst. II*, vol. 57, no. 2, pp. 69–74, Feb. 2010, doi: 10.1109/TCSII.2010.2042131.
- [12] F. Ye, S. Li, M. Zhu, Z. Ni, and J. Ren, ‘A 13-bit 180-MS/s SAR ADC with Efficient Capacitor-Mismatch Estimation and Dither Enhancement’, in *2019 IEEE International Symposium on Circuits and Systems (ISCAS)*, Sapporo, Japan: IEEE, May 2019, pp. 1–4. doi: 10.1109/ISCAS.2019.8702487.
- [13] J. Shen *et al.*, ‘A 16-bit 16-MS/s SAR ADC With On-Chip Calibration in 55-nm CMOS’, *IEEE J. Solid-State Circuits*, vol. 53, no. 4, pp. 1149–1160, Apr. 2018, doi: 10.1109/JSSC.2017.2784761.

- [14] M. Gu, Y. Zhong, L. Jie, and N. Sun, 'A 12b 1GS/s Pipelined ADC with Digital Background Calibration of Inter-stage Gain, Capacitor Mismatch, and Kick-back Errors', in *ESSCIRC 2023- IEEE 49th European Solid State Circuits Conference (ESSCIRC)*, Lisbon, Portugal: IEEE, Sep. 2023, pp. 329–332. doi: 10.1109/ESSCIRC59616.2023.10268748.
- [15] Leon Melkonian, 'AN-804 Improving A/D Converter Performance Using Dither', *Improving A/D Converter Performance Using Dither*. Accessed: Jul. 19, 2023. [Online]. Available: <https://www.ti.com/lit/an/snoa232/snoa232.pdf>
- [16] X. He, M. Gu, H. Jiang, Y. Zhong, N. Sun, and L. Jie, '9.3 A 71dB SNDR 200MHz BW Interleaved Pipe-SAR ADC with a Shared Residue Integrating Amplifier Achieving 173dB FoMs', in *2024 IEEE International Solid-State Circuits Conference (ISSCC)*, San Francisco, CA, USA: IEEE, Feb. 2024, pp. 172–174. doi: 10.1109/ISSCC49657.2024.10454431.
- [17] J. Lagos *et al.*, 'A 10.0 ENOB, 6.2 fJ/conv.-step, 500 MS/s Ringamp-Based Pipelined-SAR ADC with Background Calibration and Dynamic Reference Regulation in 16nm CMOS', in *2021 Symposium on VLSI Circuits*, Kyoto, Japan: IEEE, Jun. 2021, pp. 1–2. doi: 10.23919/VLSICircuits52068.2021.9492354.
- [18] P. Wang, F. Li, and Z. Wang, 'A Single-Channel 8-bit 1.6-GS/s Alternate-Comparator SAR ADC With Dither-Based Background Offset Calibration in 28-nm CMOS', *IEEE Trans. Circuits Syst. I*, pp. 1–12, 2025, doi: 10.1109/TCSI.2025.3526593.
- [19] B. Murmann, 'On the use of redundancy in successive approximation A/D converters', in *Proceedings of the 10th International Conference on Sampling Theory and Applications*, Bremen, Germany: IEEE Signal Processing Society, Sep. 2013. doi: <https://doi.org/10.5281/zenodo.54476>.
- [20] T. Ogawa *et al.*, 'SAR ADC Algorithm with Redundancy and Digital Error Correction', *IEICE Trans. Fundamentals*, vol. E93-A, no. 2, pp. 415–423, 2010, doi: 10.1587/transfun.E93.A.415.
- [21] R. Bodnar *et al.*, 'A 9.3nV/rHz 20b 40MS/s 94.2dB DR Signal-Chain Friendly Precision SAR Converter', in *2024 IEEE International Solid-State Circuits Conference (ISSCC)*, San Francisco, CA, USA: IEEE, Feb. 2024, pp. 182–184. doi: 10.1109/ISSCC49657.2024.10454329.
- [22] Z. Lan, L. Dong, X. Jing, and L. Geng, 'A 12-Bit 100MS/s SAR ADC with Digital Error Correction and High-Speed LMS-Based Background Calibration', in *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*, Daegu, Korea: IEEE, May 2021, pp. 1–5. doi: 10.1109/ISCAS51556.2021.9401172.
- [23] N. Le Dortz *et al.*, 'A 1.62GS/s time-interleaved SAR ADC with digital background mismatch calibration achieving interleaving spurs below 70dBFS', in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, San Francisco, CA, USA: IEEE, Feb. 2014, pp. 386–388. doi: 10.1109/ISSCC.2014.6757481.
- [24] Y. Shen, H. Li, H. Xin, E. Cantatore, and P. Harpe, 'A 103-dB SFDR Calibration-Free Oversampled SAR ADC With Mismatch Error Shaping and Pre-Comparison Techniques', *IEEE J. Solid-State Circuits*, vol. 57, no. 3, pp. 734–744, Mar. 2022, doi: 10.1109/JSSC.2021.3135559.
- [25] Y. Lim and M. P. Flynn, 'A 1 mW 71.5 dB SNDR 50 MS/s 13 bit Fully Differential Ring Amplifier Based SAR-Assisted Pipeline ADC', *IEEE J. Solid-State Circuits*, vol. 50, no. 12, pp. 2901–2911, Dec. 2015, doi: 10.1109/JSSC.2015.2463094.
- [26] H.-Y. Lee, B. Lee, and U.-K. Moon, 'A 31.3fJ/conversion-step 70.4dB SNDR 30MS/s 1.2V two-step pipelined ADC in 0.13μm CMOS', in *2012 IEEE International Solid-State*

- Circuits Conference*, San Francisco, CA, USA: IEEE, Feb. 2012, pp. 474–476. doi: 10.1109/ISSCC.2012.6177097.
- [27] J. Gao *et al.*, ‘18.3: A 93.3dB SNDR, 180.4dB FoMs Calibration-Free Noise-Shaping Pipelined-SAR ADC with Cross-Stage Gain-Mismatch-Error-Shaping Technique and Negative-R-Assisted Residue Integrator’, in *2025 IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2025, pp. 310–312. doi: 10.1109/ISSCC49661.2025.10904557.
- [28] H. Li, Y. Shen, H. Xin, E. Cantatore, and P. Harpe, ‘A 7.3- μ W 13-ENOB 98-dB SFDR Noise-Shaping SAR ADC With Duty-Cycled Amplifier and Mismatch Error Shaping’, *IEEE J. Solid-State Circuits*, vol. 57, no. 7, pp. 2078–2089, Jul. 2022, doi: 10.1109/JSSC.2022.3168588.
- [29] D. Shamiryan, T. Abell, F. Iacopi, and K. Maex, ‘Low-k dielectric materials’, 2004.
- [30] J. Gambino, ‘Handbook of Thin Film Deposition’, in *Handbook of Thin Film Deposition: Fourth Edition*, 4th edn, Elsevier, 2018, pp. 147–194. doi: 10.1016/B978-0-12-812311-9.00006-2.
- [31] M. J. M. Pelgrom, *Analog-to-Digital Conversion*. Dordrecht: Springer Netherlands, 2010. doi: 10.1007/978-90-481-8888-8.
- [32] W. Chang and C. Yao-I, ‘Method for fabricating a hybrid low-dielectric-constant intermetal dielectric (IMD) layer with improved reliability for multilevel interconnections’, Jan. 11, 1999
- [33] K. Maex, M. R. Baklanov, D. Shamiryan, F. Iacopi, S. H. Brongersma, and Z. S. Yanovitskaya, ‘Low dielectric constant materials for microelectronics’, *Journal of Applied Physics*, vol. 93, no. 11, pp. 8793–8841, Jun. 2003, doi: 10.1063/1.1567460.
- [34] ‘Low-K Films’, SVMl. Accessed: Aug. 16, 2025. [Online]. Available: <https://svmi.com/service/low-k-films/>
- [35] W. Baomin, R. Guoping, J. Yulong, Q. Xinping, L. Bingzong, and L. Ran, ‘Capacitance–voltage characterization of fully silicided gated MOS capacitor’, *J. Semicond.*, vol. 30, no. 3, p. 034002, Mar. 2009, doi: 10.1088/1674-4926/30/3/034002.
- [36] B. Razavi, *Analysis and design of data converters*, 1st edn. Cambridge: Cambridge University Press, 2025.
- [37] R. Aparicio and A. Hajimiri, ‘Capacity limits and matching properties of integrated capacitors’, *IEEE Journal of Solid-State Circuits*, vol. 37, no. 3, pp. 384–393, Mar. 2002, doi: 10.1109/4.987091.
- [38] P. Harpe, ‘A Compact 10-b SAR ADC With Unit-Length Capacitors and a Passive FIR Filter’, *IEEE J. Solid-State Circuits*, vol. 54, no. 3, pp. 636–645, Mar. 2019, doi: 10.1109/JSSC.2018.2878830.
- [39] Sewon Lee, Hyein Kang, and Minjae Lee, ‘A 2.72fJ/conv 13b 2MS/s SAR ADC Using Dynamic Capacitive Comparator with Wide Input Common mode’, in *IEEE International Solid-State Circuits Conference*, San Francisco, CA, USA: IEEE, 2024, pp. 184–186. Accessed: Mar. 18, 2024. [Online]. Available: <https://ieeexplore.ieee.org/document/10454566>
- [40] P. J. A. Harpe *et al.*, ‘A 26 μ W 8 bit 10 MS/s Asynchronous SAR ADC for Low Energy Radios’, *IEEE J. Solid-State Circuits*, vol. 46, no. 7, pp. 1585–1595, Jul. 2011, doi: 10.1109/JSSC.2011.2143870.
- [41] V. Tripathi and B. Murmann, ‘Mismatch Characterization of Small Metal Fringe Capacitors’, *IEEE Trans. Circuits Syst. I*, vol. 61, no. 8, pp. 2236–2242, Aug. 2014, doi: 10.1109/TCSI.2014.2332264.

- [42] H. Tuinhout, I. Brunets, and A. Z. Duijnhoven, 'Test structure to assess the useful extent of regular dummy devices around high-precision metal fringe capacitor arrays', in *2019 IEEE 32nd International Conference on Microelectronic Test Structures (ICMTS)*, Kita-Kyushu City, Fukuoka, Japan: IEEE, Mar. 2019, pp. 64–68. doi: 10.1109/ICMTS.2019.8730988.
- [43] B. An, S. Huang, Z. Chen, Z. Lu, W. Lu, and Y. Zhang, 'A 16bit 1MS/s High-Bit Sampling SAR ADC with Improved Binary-Weighted Capacitive Array', in *2020 IEEE 5th International Conference on Integrated Circuits and Microsystems (ICICM)*, Nanjing, China: IEEE, Oct. 2020, pp. 267–271. doi: 10.1109/ICICM50929.2020.9292270.
- [44] S. Lei, D. Qinyuan, L. Chuangchuan, and Q. Gaoshuai, 'Analysis on Capacitor Mismatch and Parasitic Capacitors Effect of Improved Segmented-Capacitor Array in SAR ADC', in *2009 Third International Symposium on Intelligent Information Technology Application*, NanChang, China: IEEE, 2009, pp. 280–283. doi: 10.1109/IITA.2009.193.
- [45] A. Ahuja, K. Badami, C. Barbelenet, and S. Emery, 'Comparison of Capacitive DAC Architectures for Power and Area Efficient SAR ADC Designs', in *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*, Daegu, Korea: IEEE, May 2021, pp. 1–5. doi: 10.1109/ISCAS51556.2021.9401768.
- [46] A. Lopez-Angulo, A. Gines, E. Peralias, and A. Rueda, 'Redundant SAR ADCs with Split-capacitor DAC', in *2018 25th IEEE International Conference on Electronics, Circuits and Systems (ICECS)*, Bordeaux: IEEE, Dec. 2018, pp. 801–804. doi: 10.1109/ICECS.2018.8618051.
- [47] J. Liu, Y. Zhu, C.-H. Chan, S.-W. Sin, S.-P. U, and R. P. Da Silva Martins, 'Uniform Quantization Theory-Based Linearity Calibration for Split Capacitive DAC in an SAR ADC', *IEEE Trans. VLSI Syst.*, vol. 24, no. 7, pp. 2603–2607, Jul. 2016, doi: 10.1109/TVLSI.2015.2509164.
- [48] H. Tuinhout and R. Van Dalen, 'Design and use of an array-based test structure to characterize mechanical stress effects caused by WLCSP solder bumps', in *2016 International Conference on Microelectronic Test Structures (ICMTS)*, Yokohama, Japan: IEEE, Mar. 2016, pp. 62–67. doi: 10.1109/ICMTS.2016.7476175.
- [49] H. Tuinhout, A. Z. Duijnhoven, and I. Brunets, 'A test structure to reveal short-range correlation effects of mismatch fluctuations in backend metal fringe capacitors', in *2018 IEEE International Conference on Microelectronic Test Structures (ICMTS)*, Austin, TX: IEEE, Mar. 2018, pp. 87–92. doi: 10.1109/ICMTS.2018.8383771.
- [50] Y.-C. Kwon and O.-K. Kwon, 'A Precision Mismatch Measurement Technique for Integrated Capacitor Array Using a Switched Capacitor Amplifier', *IEEE Trans. Semicond. Manufact.*, vol. 26, no. 2, pp. 226–232, May 2013, doi: 10.1109/TSM.2013.2254731.
- [51] K. Kundert, 'Modelling Dielectric Absorption in Capacitors', 2004. Accessed: Sep. 28, 2021. [Online]. Available: www.designers-guide.org
- [52] J. W. Fattaruso, M. D. Wit, G. Warwar, K. S. Tan, and R. K. Hester, 'The Effect of Dielectric Relaxation on Charge-Redistribution A/D Converters', *IEEE Journal of Solid-State Circuits*, vol. 25, no. 6, pp. 1550–1561, 1990, doi: 10.1109/4.62192.
- [53] T. Matsui, K. Sento, T. Ebata, and A. Ishibashi, 'A capacitor dielectric relaxation effect cancellation circuit in a 12-bit, 1-MSps, 5.0-V SAR ADC on a 28-nm embedded flash memory microcontroller', *IEEE Solid-State Circuits Letters*, vol. 2, no. 9, pp. 95–98, Sep. 2019, doi: 10.1109/LSSC.2019.2931606.
- [54] K. Mistry *et al.*, 'A 45nm Logic Technology with High-k+Metal Gate Transistors, Strained Silicon, 9 Cu Interconnect Layers, 193nm Dry Patterning, and 100% Pb-free Packaging', in

- 2007 IEEE International Electron Devices Meeting, Dec. 2007, pp. 247–250. doi: 10.1109/IEDM.2007.4418914.
- [55] X. Chen, Z. Zhang, S. Yu, and T.-G. Zsuzsán, ‘Fringing Effect Analysis of Parallel Plate Capacitors for Capacitive Power Transfer Application’, in *2019 IEEE 4th International Future Energy Electronics Conference (IFEEEC)*, Singapore, Singapore: IEEE, Nov. 2019, pp. 1–5. doi: 10.1109/IFEEEC47410.2019.9015111.
- [56] D. Yim and B. Vincent, ‘Analysis of Material, Design & LER of Advanced BEOL Metal Lines Using Process Modeling’, in *2023 IEEE International Interconnect Technology Conference (IITC) and IEEE Materials for Advanced Metallization Conference (MAM)(IITC/MAM)*, Dresden, Germany: IEEE, May 2023, pp. 1–3. doi: 10.1109/IITC/MAM57687.2023.10154880.
- [57] S. Narinesingh, J. Leffew, and W. A. Moreno, ‘Analysis of pattern dependencies on copper damascene chemical mechanical polishing’, in *Proceedings of the Fourth IEEE International Caracas Conference on Devices, Circuits and Systems (Cat. No.02TH8611)*, Apr. 2002, pp. D033–D033. doi: 10.1109/ICDCS.2002.1004047.
- [58] Y. L. Hsieh et al., ‘Effects of BEOL copper CMP process on TDDB for direct polishing ultra-low k dielectric cu interconnects at 28nm technology node and beyond’, in *2013 IEEE International Reliability Physics Symposium (IRPS)*, Anaheim, CA: IEEE, Apr. 2013, p. BD.3.1-BD.3.5. doi: 10.1109/IRPS.2013.6532056.
- [59] T. Park et al., ‘PATTERN AND PROCESS DEPENDENCIES IN COPPER DAMASCENE CHEMICAL MECHANICAL POLISHING PROCESSES’, in *VLSI Multilevel Interconnect Conference (VMIC)*, Santa Clara, CA, Jun. 1998. [Online]. Available: <https://boning.mit.edu/wp-content/uploads/2022/11/Pattern-and-Process-Dependencies-in-Copper-Damascene-Chemical-Mechanical-Polishing-Processes.pdf>
- [60] A. Toffoli et al., ‘Test structure for characterizing metal thickness in damascene CMP technology’, in *2008 IEEE International Conference on Microelectronic Test Structures*, Edinburgh, UK: IEEE, Mar. 2008, pp. 210–213. doi: 10.1109/ICMTS.2008.4509340.
- [61] G. Marxsen, T. Merbeth, M. Nopper, T. Ortleb, M. U. Lehr, and T. Roessler, ‘Major Influences of Shaping and Profiling of Cu ECP and CMP on Feature Level’, 2007.
- [62] R. Chang, Y. Cao, and C. J. Spanos, ‘Modeling the Electrical Effects of Metal Dishing Due to CMP for On-Chip Interconnect Optimization’, *IEEE Trans. Electron Devices*, vol. 51, no. 10, pp. 1577–1583, Oct. 2004, doi: 10.1109/TED.2004.834898.
- [63] N.-C. Chen, P.-Y. Chou, H. Graeb, and M. P.-H. Lin, ‘High-density MOM capacitor array with novel mortise-tenon structure for low-power SAR ADC’, in *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2017*, Lausanne, Switzerland: IEEE, Mar. 2017, pp. 1757–1762. doi: 10.23919/DATE.2017.7927277.
- [64] W. Posch, G. Promitzer, and E. Seebacher, ‘A test structure for integrated capacitor array matching characterization’, in *2010 International Conference on Microelectronic Test Structures (ICMTS)*, Hiroshima, Japan: IEEE, Mar. 2010, pp. 152–157. doi: 10.1109/ICMTS.2010.5466833.
- [65] B. Razavi, *Design of analog CMOS integrated circuits*, Second edition. New York, NY: McGraw-Hill Education, 2017.
- [66] W. Posch and E. Seebacher, ‘Integrated Capacitor Array Matching Characterization’, *IEEE Trans. Semicond. Manufact.*, vol. 25, no. 3, pp. 331–338, Aug. 2012, doi: 10.1109/TSM.2012.2202791.

- [67]Y. Hou and Q. Wu, 'Optical Proximity Correction, Methodology and Limitations', in *2021 China Semiconductor Technology International Conference (CSTIC)*, Shanghai, China: IEEE, Mar. 2021, pp. 1–5. doi: 10.1109/CSTIC52283.2021.9461507.
- [68]H.-C. Shao, C.-W. Lin, and S.-Y. Fang, 'Data-Driven Approaches for Process Simulation and Optical Proximity Correction', in *2023 28th Asia and South Pacific Design Automation Conference (ASP-DAC)*, Jan. 2023, pp. 721–726. Accessed: Oct. 08, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/10044862/>
- [69]A. Abusleme, A. Dragone, G. Haller, and B. Murmann, 'Mismatch of lateral field metal-oxide-metal capacitors in 180 nm CMOS process', *Electron. Lett.*, vol. 48, no. 5, p. 286, 2012, doi: 10.1049/el.2011.3804.
- [70]J.-B. Shyu, G. C. Temes, and K. Yao, 'Random errors in MOS capacitors', *IEEE J. Solid-State Circuits*, vol. 17, no. 6, pp. 1070–1076, Dec. 1982, doi: 10.1109/JSSC.1982.1051862.
- [71]M. Marin, S. Cremer, J.-C. Giraudin, and B. Martinet, 'Modeling the Mismatch of High-k MIM Capacitors', in *2007 IEEE International Conference on Microelectronic Test Structures*, Tokyo: IEEE, Mar. 2007, pp. 115–119. doi: 10.1109/ICMTS.2007.374466.
- [72]H. Tuinhout and N. Wils, 'A cross-coupled common centroid test structures layout method for high precision MIM capacitor mismatch measurements', in *2014 International Conference on Microelectronic Test Structures (ICMTS)*, Udine, Italy: IEEE, Mar. 2014, pp. 243–248. doi: 10.1109/ICMTS.2014.6841500.
- [73]B. Murmann and B. E. Boser, *Digitally Assisted Pipeline ADCs: Theory and Implementation*, 1st ed. 2004. New York, NY: Imprint: Springer, 2004. doi: 10.1007/b117112.
- [74]V. Tripathi and B. Murmann, 'Mismatch characterization of small metal fringe capacitors', in *Proceedings of the IEEE 2013 Custom Integrated Circuits Conference*, Sep. 2013, pp. 1–4. doi: 10.1109/CICC.2013.6658431.
- [75]H. Omran, H. Alahmadi, and K. N. Salama, 'Matching properties of femtofarad and sub-femtofarad MOM capacitors', *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 63, no. 6, pp. 763–772, Jun. 2016, doi: 10.1109/TCSI.2016.2537824.
- [76]H. Omran, R. T. ElAfandy, M. Arsalan, and K. N. Salama, 'Direct Mismatch Characterization of Femtofarad Capacitors', *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 63, no. 2, pp. 151–155, Feb. 2016, doi: 10.1109/TCSII.2015.2468919.
- [77]C. Yang *et al.*, 'An Area-Efficient SAR ADC With Mismatch Error Shaping Technique Achieving 102-dB SFDR 90.2-dB SNDR Over 20-kHz Bandwidth', *IEEE Trans. VLSI Syst.*, vol. 29, no. 8, pp. 1575–1585, Aug. 2021, doi: 10.1109/TVLSI.2021.3087660.
- [78]H. P. Tuinhout, H. Elzinga, J. T. Brugman, and F. Postma, 'Accurate capacitor matching measurements using floating gate test structures', in *Proceedings International Conference on Microelectronic Test Structures*, Nara, Japan: IEEE, 1995, pp. 133–137. doi: 10.1109/ICMTS.1995.513960.
- [79]H. Tuinhout, F. Van Rossem, and N. Wils, 'High precision on-wafer backend capacitor mismatch measurements using a benchtop semiconductor characterization system', in *2009 IEEE International Conference on Microelectronic Test Structures*, Oxnard, CA, USA: IEEE, Mar. 2009, pp. 3–8. doi: 10.1109/ICMTS.2009.4814598.
- [80]'A JSSC classic paper: Matching properties of MOS transistors', *IEEE Solid-State Circuits Society Newsletter*, vol. 10, no. 1, pp. 6–7, Jan. 2005, doi: 10.1109/N-SSC.2005.6500068.

- [81] M. Pelgrom, H. Tuinhout, and M. Vertregt, 'A Designer's View on Mismatch', in *Nyquist AD Converters, Sensor Interfaces, and Robustness*, A. H. M. Van Roermund, A. Baschiroto, and M. Steyaert, Eds, New York, NY: Springer New York, 2013, pp. 245–267. doi: 10.1007/978-1-4614-4587-6_13.
- [82] S. Mukherjee and G. Aggarwal, 'Metal CMP process optimization for low abrasive slurry', in *2018 29th Annual SEMI Advanced Semiconductor Manufacturing Conference (ASMC)*, Saratoga Springs, NY, USA: IEEE, Apr. 2018, pp. 226–230. doi: 10.1109/ASMC.2018.8373168.
- [83] P. Harpe, 'Energy Efficient ADC Design Techniques', in *2024 IEEE Custom Integrated Circuits Conference (CICC)*, Denver, CO, USA: IEEE, Apr. 2024, pp. 1–2. doi: 10.1109/CICC60959.2024.10529035.
- [84] J. Lienig and J. Scheible, *Fundamentals of Layout Design for Electronic Circuits*. Cham: Springer International Publishing, 2020. doi: 10.1007/978-3-030-39284-0.