DESIGN AND ANALYSIS OF A TWO-STAGE BUFFER AMPLIFIER IN 28 nm AND 18 nm CMOS TECHNOLOGIES FOR LOW-VOLTAGE REFERENCE SYSTEMS

A PROJECT REPORT

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE AWARD OF THE DEGREE
OF

MASTER OF TECHNOLOGY
IN
CONTROL AND INSTRUMENTATION (C&I)

Submitted by:

ANSHUMAN TIWARI
(Roll No. 2K24/C&I/02)

Under the supervision of

DR. RAM BHAGAT
Professor, Electrical Engineering

DEPARTMENT OF ELECTRICAL ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Bawana Road, Delhi-110042

JUNE, 2026

DEPARTMENT OF ELECTRICAL ENGINEERING

DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)

Bawana Road, Delhi-110042

CANDIDATE'S DECLARATION

I, Anshuman Tiwari, Roll No. 2K24/C&I/02, student of M.Tech (Control and Instrumentation), hereby declare that the project dissertation titled "Design and Analysis of a Two-Stage Buffer Amplifier in 28 nm and 18 nm CMOS Technologies for Low-Voltage Reference Systems", which is submitted by me to the Department of Electrical Engineering, Delhi Technological University, Delhi, in partial fulfillment of the requirement for the award of the degree of Master of Technology, is original and not copied from any source without proper citation. This work has not previously formed the basis for the award of any Degree, Diploma, Associateship, Fellowship, or other similar title or recognition.

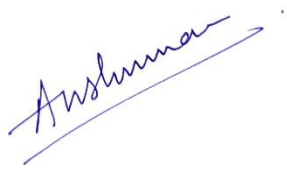

Place: Delhi **(ANSHUMAN TIWARI)**

Date: 29-05-2026

DEPARTMENT OF ELECTRICAL ENGINEERING

DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)

Bawana Road, Delhi-110042

CERTIFICATE

I hereby certify that the project dissertation titled "Design and Analysis of a Two-Stage Buffer Amplifier in 28 nm and 18 nm CMOS Technologies for Low-Voltage Reference Systems", which is submitted by Anshuman Tiwari, Roll No. 2K24/C&I/02, of the Department of Electrical Engineering, Delhi Technological University, Delhi, in partial fulfillment of the requirement for the award of the degree of Master of Technology, is a record of the project work carried out by the student under my supervision. To the best of my knowledge, this work has not been submitted in part or full for any Degree or Diploma to this University or elsewhere.

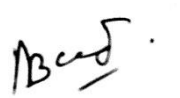

Place: Delhi **(DR. RAM BHAGAT)**

Date: 29-05-2026 **SUPERVISOR**

Professor, Electrical Engineering

ACKNOWLEDGEMENT

I would like to express my deepest gratitude to my supervisor Dr. Ram Bhagat for the continuous guidance, encouragement, and constructive feedback offered throughout the course of this project work. The technical depth and patient mentoring extended at every stage of the work have been invaluable in shaping both the research and the researcher.

I sincerely thank the Head of the Department and all the faculty members of the Department of Electrical Engineering, Delhi Technological University, for providing the necessary infrastructure and the simulation environment required to carry out this work. I also extend my gratitude to the technical staff of the laboratory for their cooperation in setting up the design tools and process design kits.

I am thankful to my friends and fellow researchers for many fruitful technical discussions on analog circuit design, frequency compensation, and short-channel CMOS behaviour. Their inputs have helped refine several design choices presented in this project report.

Finally, I express my deep gratitude to my family for their unconditional support, patience, and motivation during the entire period of my studies

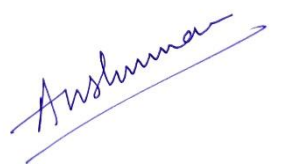

(ANSHUMAN TIWARI)

ABSTRACT

Voltage reference circuits such as bandgap references, low-dropout regulators, and on-chip bias generators are increasingly required to drive multiple distributed loads inside modern mixed-signal systems-on-chip. To prevent loading and noise coupling, the reference is usually isolated from its loads using an analog buffer amplifier. The buffer is required to provide near-unity voltage gain, high input impedance, low output impedance, sufficient bandwidth to track transient demands, and unconditional stability under closed-loop operation. In aggressively scaled CMOS nodes, the limited intrinsic gain of short-channel devices, the reduced supply headroom, and the increased parasitic capacitance make the design of such a buffer non-trivial.

This project report presents the design, analysis, and simulation of a two-stage CMOS buffer amplifier intended for buffering on-chip voltage references of 0.6 V, 1.2 V, and 1.3 V. The buffer is implemented and characterised in both 28 nm and 18 nm CMOS technology nodes. The first stage is realised as a PMOS-input differential amplifier with an NMOS current-mirror load, providing the dominant voltage gain and a low-noise, high-input-impedance front end. The second stage is an NMOS common-source output stage with a PMOS active load, which provides additional gain and improved output drive capability. The amplifier is closed in unity-gain feedback configuration to operate as a buffer, and a series resistor–capacitor (RC) network is connected between the high-impedance node of the first stage and the output node to provide frequency compensation. The RC compensation moves the dominant pole to a lower frequency, splits the non-dominant pole to a higher frequency, and cancels (or moves to the left-half-plane) the right-half-plane zero introduced by the feedforward path through the compensation capacitor, thereby ensuring a phase margin sufficient for stable operation.

The proposed buffer targets a nominal unity-gain bandwidth (UGB) of 10 MHz at the typical corner across all three reference voltages and both technology nodes. Transistor-level simulations across the five process corners (TT, SS, FF, SF, FS) demonstrate that, for the 28 nm implementation, the open-loop DC gain spans approximately 85 dB to 110 dB (nominal 92 dB) and the UGB spans approximately

3 MHz to 40 MHz (nominal 10 MHz). For the 18 nm implementation, the open-loop DC gain spans approximately 82 dB to 103 dB and the UGB spans approximately 5 MHz to 40 MHz. A detailed comparative study between the 28 nm and 18 nm implementations is reported, and the differences observed are explained in terms of intrinsic gain (gm·ro), short-channel effects, parasitic capacitance, and the difference in carrier mobility between the two nodes. The project report additionally provides a comprehensive review of frequency compensation techniques used in two-stage operational amplifiers, and discusses the rationale for selecting RC compensation for this work.

Keywords: Two-stage amplifier, unity-gain buffer, voltage reference, 28 nm CMOS, 18 nm CMOS, RC compensation, Miller compensation, nulling resistor, phase margin, unity-gain bandwidth, low-voltage analog design, differential amplifier, NMOS output stage, frequency compensation, pole splitting, process corners.

CONTENTS

Candidate's Declaration ii

Certificate iii

Acknowledgement iv

Abstract v

List of Figures viii

List of Tables ix

List of Abbreviations x

CHAPTER 1 INTRODUCTION

1.1 BACKGROUND OF ANALOG BUFFERS IN INTEGRATED CIRCUITS 1

1.2 NEED FOR A BUFFER IN A CHIP 2

1.3 CHALLENGES IN LOW-VOLTAGE ADVANCED CMOS NODES 3

1.4 MOTIVATION 4

1.5 PROJECT REPORT OBJECTIVES 4

1.6 CONTRIBUTIONS OF THE PROJECT REPORT 5

1.7 ORGANISATION OF THE PROJECT REPORT 6

CHAPTER 2 LITERATURE REVIEW

2.1 BUFFER AMPLIFIER TOPOLOGIES 7

2.2 SINGLE-STAGE AMPLIFIER APPROACHES 7

2.3 TWO-STAGE AMPLIFIER APPROACHES 8

2.4 EXISTING COMPENSATION TECHNIQUES IN ANALOG CMOS DESIGN 9

2.5 ADVANTAGE OF TWO-STAGE AMPLIFIERS FOR BUFFERING APPLICATIONS 10

2.6 COMPARISON WITH PRIOR WORK 11

2.7 CRITICAL REVIEW AND RESEARCH GAPS 12

CHAPTER 3 PROPOSED WORK AND DESIGN METHODOLOGY

3.1 PROPOSED ARCHITECTURE 14
3.2 FIRST-STAGE DIFFERENTIAL AMPLIFIER 15
3.3 SECOND-STAGE NMOS OUTPUT STAGE 17
3.4 JUSTIFICATION OF THE CHOSEN TOPOLOGY 18
3.5 DESIGN CONSIDERATIONS FOR 28 NM AND 18 NM CMOS 18
3.6 TRANSISTOR SIZING STRATEGY 19
3.7 STEP-BY-STEP PROCEDURE TO BUILD THE BUFFER 21
3.8 FREQUENCY COMPENSATION USING AN RC NETWORK 22
3.9 STABILITY ANALYSIS 24
3.10 DESIGN TRADE-OFFS 24
3.11 RATIONALE FOR PREFERRING A TWO-STAGE DESIGN OVER A SINGLE-STAGE BUFFER 25

CHAPTER 4 COMPENSATION TECHNIQUES IN TWO-STAGE AMPLIFIERS

4.1 NEED FOR FREQUENCY COMPENSATION 26
4.2 MILLER COMPENSATION 26
4.3 MILLER COMPENSATION WITH A NULLING RESISTOR (RC COMPENSATION) 27
4.4 LEAD COMPENSATION 28
4.5 POLE-SPLITTING COMPENSATION 29
4.6 FEEDFORWARD COMPENSATION 30
4.7 AHUJA (CASCODE) COMPENSATION 30
4.8 NESTED MILLER COMPENSATION 31
4.9 CURRENT-BUFFER COMPENSATION 32
4.10 SPLIT-LENGTH AND ACTIVE COMPENSATION TECHNIQUES 33
4.11 COMPARATIVE SUMMARY 34
4.12 RATIONALE FOR SELECTING RC COMPENSATION 35

CHAPTER 5 SIMULATION RESULTS AND DISCUSSION

5.1 SIMULATION SETUP AND METHODOLOGY 36

5.2 DC OPERATING POINT VERIFICATION 37

5.3 AC ANALYSIS RESULTS 37

5.4 TRANSIENT RESPONSE 42

5.5 CLOSED-LOOP DISTURBANCE RESPONSE 44

5.6 PERFORMANCE ACROSS REFERENCE VOLTAGES 45

5.7 COMPARISON BETWEEN 28 NM AND 18 NM IMPLEMENTATIONS 46

5.8 PHYSICAL ORIGINS OF THE INTER-NODE DIFFERENCES 48

5.9 DISCUSSION OF RESULTS 51

CHAPTER 6 CONCLUSION AND FUTURE SCOPE

6.1 SUMMARY OF THE WORK 53

6.2 MAIN ACHIEVEMENTS 54

6.3 SIGNIFICANCE OF THE BUFFER BLOCK 54

6.4 ROLE OF FREQUENCY COMPENSATION IN THE DESIGN 54

6.5 EFFECTIVENESS OF THE TWO-STAGE APPROACH 55

6.6 FUTURE SCOPE 55

6.6.1 POWER CONSUMPTION REDUCTION 55

6.6.2 IMPROVED COMPENSATION TECHNIQUES 55

6.6.3 ENHANCED OUTPUT SWING AND LINEARITY 56

6.6.4 ADAPTATION TO WIDER SUPPLY RANGES 56

6.6.5 PORTING TO NEWER TECHNOLOGY NODES 56

6.6.6 INTEGRATION INTO LARGER MIXED-SIGNAL SUBSYSTEMS 57

6.7 FINAL CONCLUSION 57

REFERENCES

LIST OF FIGURES

Figure 3.1 Block-level diagram of the proposed two-stage unity-gain buffer..14

Figure 3.2 Series RC compensation network connected between two nodes...15

Figure 3.3 Transistor-level schematic of the two-stage operational amplifier. 16

Figure 3.4 Closed-form expressions for the nulling resistor and compensation capacitor. 23

Figure 4.1 Miller compensation with current-buffer (virtual-ground) topology.31

Figure 4.2 Three-stage nested-Miller / multipath compensated amplifier. 32

Figure 4.3 Basic module for multipath NM zero cancellation. 33

Figure 5.1 Open-loop frequency response across the PVT sweep — 28 nm vs 18 nm. 38

Figure 5.2 Clean TT-corner open-loop response — 28 nm vs 18 nm. 39

Figure 5.3 DC gain and UGB across process corners — 28 nm vs 18 nm. 41

Figure 5.4 Phase margin across process corners — 28 nm vs 18 nm. 42

Figure 5.5 Small-signal output settling at V_ref = 1.2 V and 1.3 V. 43

Figure 5.6 Power-up startup transient. 44

Figure 5.7 Closed-loop disturbance-to-output magnitude across PVT. 45

Figure 5.8 Nominal DC gain versus reference voltage at TT corner. 46

LIST OF TABLES

Table 2.1 Comparative summary of prior work on CMOS buffer amplifiers. .11

Table 3.1 Targeted design specifications of the proposed buffer. 20

Table 3.2 Transistor aspect-ratio plan for the 28 nm implementation. 20

Table 3.3 Transistor aspect-ratio plan for the 18 nm implementation. 21

Table 4.1 Comparative summary of two-stage amplifier compensation techniques. 34

Table 5.1 AC performance results in 28 nm CMOS across three reference voltages (TT). 39

Table 5.2 AC performance results in 18 nm CMOS across three reference voltages (TT). 40

Table 5.3 AC performance across process corners — 28 nm implementation. 40

Table 5.4 AC performance across process corners — 18 nm implementation. 41

Table 5.5 Transient performance summary. 44

Table 5.6 Comparison of the proposed design with selected prior work.....47

LIST OF ABBREVIATIONS

CMOS	Complementary Metal-Oxide-Semiconductor
MOSFET	Metal-Oxide-Semiconductor Field-Effect Transistor
NMOS	N-type Metal-Oxide-Semiconductor
PMOS	P-type Metal-Oxide-Semiconductor
OTA	Operational Transconductance Amplifier
OP-AMP	Operational Amplifier
UGB	Unity-Gain Bandwidth
GBW	Gain-Bandwidth Product
PM	Phase Margin
GM	Gain Margin
DC	Direct Current
AC	Alternating Current
CMRR	Common-Mode Rejection Ratio
PSRR	Power-Supply Rejection Ratio
LDO	Low-Dropout Regulator
ADC	Analog-to-Digital Converter
DAC	Digital-to-Analog Converter
PVT	Process, Voltage, and Temperature
RHP	Right-Half-Plane
LHP	Left-Half-Plane
RC	Resistor-Capacitor
PDK	Process Design Kit
SoC	System-on-Chip
VLSI	Very-Large-Scale Integration
TT / SS / FF / SF / FS	Process Corners (Typical-Typical, Slow-Slow, Fast-Fast, Slow-Fast, Fast-Slow)
DIBL	Drain-Induced Barrier Lowering

MOM / MIM	Metal-Oxide-Metal / Metal-Insulator-Metal capacitor

CHAPTER 1

INTRODUCTION

1.1 BACKGROUND OF ANALOG BUFFERS IN INTEGRATED CIRCUITS

An analog buffer is a unity-gain amplifier whose purpose is to isolate one analog node from another so that the loading effect of a downstream block does not disturb the upstream voltage source. From a small-signal standpoint a buffer ideally has infinite input impedance, zero output impedance, and a voltage gain of exactly one over a usefully wide frequency band [1], [2]. In practice these ideal properties cannot be achieved simultaneously, but they can be approximated closely with carefully designed CMOS operational amplifiers operated in unity-gain feedback. Such buffers are ubiquitous in modern mixed-signal integrated circuits: every voltage reference, every bandgap, every low-dropout regulator, every internal bias rail, and every analog-to-digital converter (ADC) requires at least one buffer to distribute its output without degradation.

The growth of system-on-chip (SoC) design has made the buffer one of the most frequently instantiated analog cells. A typical mixed-signal SoC may contain dozens of internal reference rails — e.g., 0.6 V, 1.2 V, and 1.3 V — each of which has to be driven into multiple, geographically dispersed loads on the same die. Without a buffer, the reference generator would be loaded by every block that consumes it, and the resulting current draw would modulate the reference voltage [3]. With a buffer, the reference generator drives a single high-impedance input and the buffer takes over the responsibility of supplying the load current. This decoupling makes the reference effectively immune to load transients and dramatically improves the accuracy of every block that depends on it.

1.2 NEED FOR A BUFFER IN A CHIP

A buffer is required in a chip for four interrelated reasons. First, the reference generator (for example, a bandgap circuit) is designed for accuracy, not for drive. Its output impedance is typically high (tens of kilo-ohms or more) and it can supply only a few microamps of load current without significant voltage droop. A buffer presents a high-impedance load to the reference generator and replaces it with a low-impedance source that can drive much larger currents [2], [3].

Second, the buffer suppresses noise coupling between distant blocks on the same die. In a SoC, switching activity on a digital subsystem injects substrate and supply noise that can find its way into analog references through parasitic coupling. A buffer with adequate power-supply rejection ratio (PSRR) attenuates this coupling and keeps the reference clean at the point of use [3], [11].

Third, the buffer regularises the transient response of the reference. Without a buffer, a load step at the reference draws current directly from the reference generator and causes a long settling tail because the reference generator is heavily compensated for stability and is intentionally slow. With a buffer, the load step is absorbed by the much faster local buffer, leaving the reference generator undisturbed [2], [11].

Fourth, the buffer enables hierarchical analog design. A single reference can be distributed through routing that traverses long metal interconnect, and an additional buffer can be placed at every block that consumes it. Each local buffer guarantees clean and stable reference delivery without imposing further requirements on the central generator. This is the essence of reference distribution in modern mixed-signal design [2], [12].

1.3 CHALLENGES IN LOW-VOLTAGE ADVANCED CMOS NODES

Designing a high-quality buffer in advanced CMOS nodes such as 28 nm and 18 nm presents several challenges that did not exist in older technologies [11], [17], [22]. The most important of these challenges are summarised below.

Low Intrinsic Device Gain

In short-channel devices, the small-signal output resistance ro is reduced by drain-induced barrier lowering (DIBL), channel-length modulation, and other short-channel effects. The intrinsic gain (gm·ro) of a single short-channel transistor in 28 nm or 18 nm is typically in the range of 5 to 15, that is, a single common-source stage delivers only about 14 dB to 24 dB of small-signal gain [11], [17]. Single-stage amplifiers therefore cannot meet the gain requirements of accurate buffers, and at least two stages are needed.

Reduced Supply Headroom

The nominal core supply of modern CMOS nodes is in the range of 0.7 V to 1.0 V. After accounting for the saturation voltages of cascoded current sources and stacked transistors, very little headroom remains for the input common-mode range. Topologies that worked at 1.8 V or higher must be revisited and sometimes redesigned for low-voltage operation [22].

Increased Parasitic Capacitance

As device dimensions shrink, the parasitic capacitances at internal nodes — gate, drain, and source — become an increasingly important fraction of the total capacitance. In addition, dense metal routing in advanced nodes contributes substantial interconnect capacitance. These parasitics lower the dominant pole and, more critically, can move the non-dominant pole into the vicinity of the unity-gain frequency, eroding the phase margin [11].

Larger PVT Spread

Advanced nodes exhibit larger relative variations of threshold voltage, mobility, and resistance across process, voltage, and temperature (PVT). A design that meets specifications at the typical corner may fail in the slow-slow (SS) or fast-fast (FF) corners. Robust analog design in these nodes therefore demands an explicit budget for corner variation, and the compensation network must be chosen so that the phase margin remains adequate over the entire PVT envelope.

1.4 MOTIVATION

The work presented in this project report is motivated by the need for a small, low-power, easily portable buffer amplifier that can serve as a reusable building block for reference distribution in modern mixed-signal SoCs. Several specific motivations can be highlighted.

A single buffer architecture that operates correctly at low (0.6 V), medium (1.2 V), and slightly elevated (1.3 V) reference voltages with the same compensation network is required because a typical SoC contains many such reference rails and re-designing the buffer for each would be wasteful. Likewise, the same circuit must port cleanly across multiple technology nodes, because design teams routinely re-use intellectual property between 28 nm and 18 nm tape-outs. Finally, the buffer must be compact, free of additional cascode headroom, and consume low static power so that it can be replicated hundreds of times across a chip without dominating the area or power budget.

These constraints favour a two-stage CMOS operational transconductance amplifier (OTA) closed in unity-gain feedback, compensated by a simple series resistor–capacitor network. Although this is a textbook architecture, its application at 0.6 V references in 28 nm and 18 nm processes is not trivial and requires careful tuning of every design parameter. This project report takes the textbook architecture as a starting point and re-derives its design choices for the specific case of multi-reference, multi-node, low-voltage analog buffering.

1.5 PROJECT REPORT OBJECTIVES

The principal objectives of this project report are stated below.

- To design a two-stage CMOS buffer amplifier with a unity-gain bandwidth of 10 MHz, a phase margin greater than 60°, and an open-loop DC gain consistent with accurate unity-gain feedback, for reference voltages of 0.6 V, 1.2 V, and 1.3 V.
- To implement and characterise the same architecture in both 28 nm and 18 nm CMOS process design kits, and to compare the resulting performance.

- To provide a comprehensive review of the frequency compensation techniques currently used in two-stage CMOS amplifiers, with a critical assessment of their suitability for low-voltage advanced nodes.
- To justify the selection of RC compensation as the most suitable scheme for this design, and to derive closed-form expressions for the compensation capacitor and the nulling resistor in terms of the small-signal parameters of the amplifier.
- To characterise the buffer over the typical (TT) corner and the four extreme corners (SS, FF, SF, FS), reporting the variation of open-loop DC gain, unity-gain bandwidth, phase margin, and quiescent current across the corners.
- To explain, in terms of device physics and small-signal behaviour, the observed differences between the 28 nm and 18 nm implementations.

1.6 CONTRIBUTIONS OF THE PROJECT REPORT

The contributions of this project report can be summarised as follows.

- A unified design methodology for two-stage CMOS buffers with RC compensation that targets multiple reference voltages and multiple technology nodes without changing the topology or the compensation scheme.
- A detailed corner analysis of the proposed buffer in both 28 nm and 18 nm CMOS, demonstrating that the open-loop DC gain spans 85 dB to 110 dB and the UGB spans 3 MHz to 40 MHz in 28 nm, and 82 dB to 103 dB and 5 MHz to 40 MHz in 18 nm.
- A device-physics-based explanation of the differences observed between the two nodes, relating gain spread to intrinsic gain (gm·ro), UGB spread to gm/C ratio variation, and corner sensitivity to threshold voltage and mobility spread.
- A consolidated literature review of frequency compensation techniques in modern CMOS amplifiers, including Miller, Miller with nulling resistor,

lead, Ahuja (cascode), nested Miller, current-buffer (indirect-feedback), feedforward, and active compensation, with explicit comparison to the present application.

1.7 ORGANISATION OF THE PROJECT REPORT

The remainder of the project report is organised as follows. Chapter 2 reviews the relevant literature on buffer amplifier topologies, two-stage operational amplifier design, and frequency compensation in advanced CMOS. Chapter 3 presents the proposed work, including the architecture, the design methodology, the transistor sizing, and the step-by-step procedure followed to build the buffer. Chapter 4 is devoted to a detailed survey of compensation techniques and a justification of the RC compensation chosen here. Chapter 5 reports the simulation results, including AC and transient performance across reference voltages and process corners, and explicitly compares the 28 nm and 18 nm implementations. Chapter 6 outlines directions for future work, and Chapter 7 concludes the project report.

CHAPTER 2

LITERATURE REVIEW

2.1 BUFFER AMPLIFIER TOPOLOGIES

Voltage buffers can be implemented in several distinct topologies, each with its own trade-offs between input range, output drive, gain accuracy, and headroom. The most common topologies in CMOS are the source-follower, the flipped voltage follower (FVF), the single-stage operational transconductance amplifier (OTA) in unity-gain feedback, and the multi-stage operational amplifier (op-amp) in unity-gain feedback [2], [3], [12].

The source-follower is the simplest. A single MOS transistor whose drain is tied to the supply and whose source drives the output through a current source can act as a buffer. Its gain is close to but always less than unity because of the body effect and the finite ro of the source-follower transistor. The level shift between input and output is approximately one threshold voltage, which is unacceptable in low-voltage designs where the headroom does not permit such a shift. The flipped voltage follower removes most of this level shift at the price of a more complex topology and a reduced output swing [2].

A single-stage OTA in unity-gain feedback provides accurate unity-gain transfer because the open-loop gain ($g_m \cdot r_o$) is enforced by the loop. However, the open-loop gain is limited to the intrinsic gain of a single short-channel device, which is too low in 28 nm and 18 nm to enforce accurate closed-loop unity gain across PVT [11], [17]. A multi-stage op-amp in unity-gain feedback removes this limitation by cascading two or three common-source stages, but it requires careful frequency compensation to remain stable in unity-gain feedback [1], [2], [3]. The work presented in this project report follows this last approach, with a two-stage architecture compensated by a series resistor–capacitor (RC) network.

2.2 SINGLE-STAGE AMPLIFIER APPROACHES

Single-stage amplifiers, such as the telescopic cascode, the folded cascode, and their gain-boosted variants, achieve very high gain in older technologies by stacking devices vertically. The telescopic cascode achieves an output resistance of $(gm \cdot ro)^2$ type, yielding 60 dB to 80 dB of gain in 180 nm or 65 nm with the supply voltages then available [2], [3]. Unfortunately, this stacking is incompatible with the supply headroom of 28 nm and 18 nm, where the core voltage may be as low as 0.9 V and the entire stack of cascoded devices does not fit comfortably between the rails.

The folded cascode relaxes the headroom constraint by folding the cascode branch sideways, but its output resistance is roughly halved and its gain is correspondingly lower than the telescopic equivalent [3]. Single-stage gain-boosted amplifiers can recover some of this gain by adding auxiliary amplifiers in feedback, but the additional complexity and the loss of phase margin make this approach unattractive for a small reusable buffer. Single-stage approaches also share a common limitation: their output drive is set by their bias current, and there is no way to boost the drive without simultaneously raising the static power. This is the principal reason why two-stage architectures are preferred for buffers that have to drive non-trivial loads in low-voltage advanced nodes.

2.3 TWO-STAGE AMPLIFIER APPROACHES

A two-stage amplifier cascades two common-source stages, each of which provides its own contribution to the overall gain. The first stage is typically a differential-input transconductor whose output drives the high-impedance internal node, and the second stage is a single-ended common-source amplifier whose output drives the load. The overall small-signal gain is approximately $Av = gm1 \cdot (ro1 \parallel ro\_load1) \cdot gm2 \cdot (ro2 \parallel RL)$, where the subscripts refer to the first- and second-stage transconductances and output resistances [1], [2].

In advanced nodes, each stage typically delivers 20 dB to 35 dB of gain, so two cascaded stages provide a combined open-loop gain in the range of 40 dB to 70 dB. The first-stage differential pair can be implemented with NMOS or PMOS input devices depending on the targeted input common-mode range, and the load can be a

current mirror, a resistor, or an active load. The second stage is most often a common-source amplifier with an active current-source load, because this maximises the output swing and the drive capability for a given headroom [3], [11].

The chief disadvantage of the two-stage architecture is its stability: the cascade of two high-impedance nodes creates two poles that lie close in frequency in the uncompensated amplifier, leading to inadequate phase margin in unity-gain feedback. Frequency compensation is therefore mandatory in any two-stage design. The compensation scheme is typically based on the Miller effect: a capacitor placed in feedback around the second stage performs pole splitting, moving the dominant pole down in frequency and pushing the non-dominant pole up [1], [2], [6].

2.4 EXISTING COMPENSATION TECHNIQUES IN ANALOG CMOS DESIGN

A substantial body of literature exists on frequency compensation techniques for two-stage and multi-stage CMOS amplifiers [6]–[10], [14]–[16]. The principal techniques are summarised here and revisited in detail in Chapter 4.

Miller compensation [1] places a capacitor between the first-stage output and the second-stage output. It is simple and effective but introduces a right-half-plane (RHP) zero at $z \approx gm2 / Cc$, which limits the achievable phase margin. Miller compensation with a nulling resistor (RC compensation) adds a small resistor in series with the Miller capacitor; the resistor can be chosen to cancel the RHP zero or to move it to the left-half-plane (LHP) where it adds phase lead and improves the phase margin [2], [3]. Lead compensation generalises this idea: any LHP zero in the open-loop transfer function adds phase lead near the unity-gain frequency.

Pole-splitting compensation refers generically to any technique that spreads the two dominant poles apart in frequency; Miller compensation is one specific example. Feedforward compensation injects a high-frequency signal path from the input to the output that adds an LHP zero, partially cancelling a non-dominant pole and extending the bandwidth without adding much delay [9], [14]. Ahuja or cascode compensation [6] routes the compensation current through a cascode device, eliminating the RHP zero without using a nulling resistor. Nested Miller

compensation [7], [8] is the natural extension of Miller compensation to three or more stages. Current-buffer compensation (also known as indirect-feedback compensation) [15] employs a common-gate stage as a current buffer in the compensation path, providing benefits similar to Ahuja compensation. Split-length compensation and active compensation [9], [10] are more recent techniques that improve the bandwidth–power trade-off in low-voltage multi-stage amplifiers.

2.5 ADVANTAGE OF TWO-STAGE AMPLIFIERS FOR BUFFERING APPLICATIONS

From the discussion above, several reasons emerge as to why two-stage amplifiers are preferred over single-stage ones when designing buffers for advanced CMOS nodes.

First, the cumulative DC gain of two cascaded common-source stages is high. Each stage contributes roughly 25 dB to 35 dB in advanced nodes, leading to a total open-loop DC gain in the range of 50 dB to 70 dB at the typical corner and considerably higher at slow corners where ro is larger. This loop gain is sufficient to ensure that the closed-loop voltage gain is very close to unity and that the closed-loop output impedance is much smaller than the open-loop output resistance [1], [2].

Second, the second stage provides genuine output drive. Because the output transistor of the second stage operates as a common source, its drain current can swing over a wide range, allowing the output node to be charged or discharged quickly even into significant capacitive loads. The output swing is also wide, typically limited only by the saturation voltage of the second-stage transistors. In a single-stage amplifier loaded by the same capacitor, the bandwidth would be lower for the same bias current [3].

Third, the two-stage architecture decouples gain from bandwidth. The first stage can be optimised for gain and noise, while the second stage can be optimised for drive. The compensation network mediates between the two, setting the dominant pole at a location that gives the desired unity-gain frequency. Fourth, two-stage amplifiers are well understood, supported by mature design methodologies, and routinely used in commercial mixed-signal designs [11], [12].

2.6 COMPARISON WITH PRIOR WORK

Numerous published designs of unity-gain buffers and two-stage operational amplifiers in advanced CMOS technologies have been reported. Representative examples cover technology nodes ranging from 180 nm down to FinFET, supply voltages from 0.5 V to 1.8 V, unity-gain bandwidths from a few hundred kilohertz to several gigahertz, and quiescent currents from sub-microamp to several milliamps [6]–[10], [14], [24], [25]. Most designs employ Miller compensation in one of its variants; some adopt cascode or feedforward compensation for improved bandwidth–stability trade-off.

Table 2.1 summarises a representative subset of prior work and positions the proposed design within this landscape. The numerical entries are illustrative and serve only as a baseline for comparison; the specific numbers for the proposed design are reported in Chapter 5.

Table 2.1 Comparative summary of prior work on CMOS buffer amplifiers (representative values).

Reference	Tech. (nm)	Supply (V)	UGB	PM (°)	Compensation
Ahuja [6]	5 µm	±5	±MHz	≥55	Cascode (Ahuja)
Leung & Mok [8]	0.8 µm / 0.5 µm	2.5	5 MHz	≥60	Nested Miller / NMC
Bult & Geelen [24]	—	5	60 MHz	≥65	Miller + Rz, gain-boosted
Peng & Sansen [9]	—	1.5	≥5 MHz	≥65	AC boosting
Saxena & Baker [15]	0.5 µm	5 / 3.3	MHz-range	≥60	Indirect feedback (current-buffer)
Grasso et al. [14]	—	low	> 1 MHz	≥55	Three-stage, frequency comp.
This Work (28 nm)	28	0.9	10 MHz (TT)	≥60	RC (Miller + Rz)
This Work (18 nm)	18	0.9	10 MHz (TT)	≥60	RC (Miller + Rz)

2.7 CRITICAL REVIEW AND RESEARCH GAPS

Although the body of literature on two-stage CMOS amplifiers and their compensation is mature, a careful review reveals several gaps that motivate the present work.

First, most published designs are reported at a single reference voltage or a single supply voltage. The ability of the same architecture to operate at low (0.6 V), mid (1.2 V), and slightly elevated (1.3 V) reference levels with the same compensation network and similar performance has received less attention. This is, however, exactly the situation faced in modern SoCs where one reusable buffer is required to handle several reference rails.

Second, very few works explicitly compare the same design implemented in two different advanced CMOS nodes. Such a comparison is valuable because it isolates the effect of technology scaling on the bandwidth, phase margin, and power of the same circuit. The present thesis attempts to fill this gap by porting the same buffer to 28 nm and 18 nm and reporting the comparative performance across the typical and four extreme process corners.

Third, while RC compensation is a textbook technique [1]–[3], its application in advanced CMOS nodes at low reference voltages requires careful re-tuning of the nulling resistor and the compensation capacitor. Reported design procedures often assume operating conditions and intrinsic gains that no longer apply in 28 nm and 18 nm. A pragmatic, step-by-step design methodology that takes the short-channel limitations of these nodes into account is therefore valuable, and is one of the contributions of this project report.

Fourth, although feedforward and Ahuja compensation techniques offer attractive bandwidth–stability trade-offs [6], [9], RC compensation remains the most practical choice for low-power, area-constrained buffers because of its simplicity and small footprint. A focused study that demonstrates the suitability of RC compensation

for the specific case of low-bandwidth (10 MHz), multi-reference buffers in advanced CMOS is provided here.

These observations define the scope of the present project report and motivate the design methodology described in Chapter 3.

CHAPTER 3

PROPOSED WORK AND DESIGN METHODOLOGY

3.1 PROPOSED ARCHITECTURE

The proposed buffer is a closed-loop, unity-gain amplifier built around a classical two-stage CMOS operational transconductance amplifier. A block-level view of the buffer is shown in Figure 3.1. The amplifier has a differential input pair (transconductance Gm1), a single high-impedance internal node (loaded by the parallel combination Ro1 ‖ (1/sC1)), a second transconductance stage (Gm2) configured as a common-source amplifier, and a single-ended output node loaded by CL. The output is fed back directly to the inverting input, while the reference voltage to be buffered is applied to the non-inverting input. The series combination of a nulling resistor Rz and a compensation capacitor Cc connects the high-impedance internal node to the output and constitutes the RC compensation network.

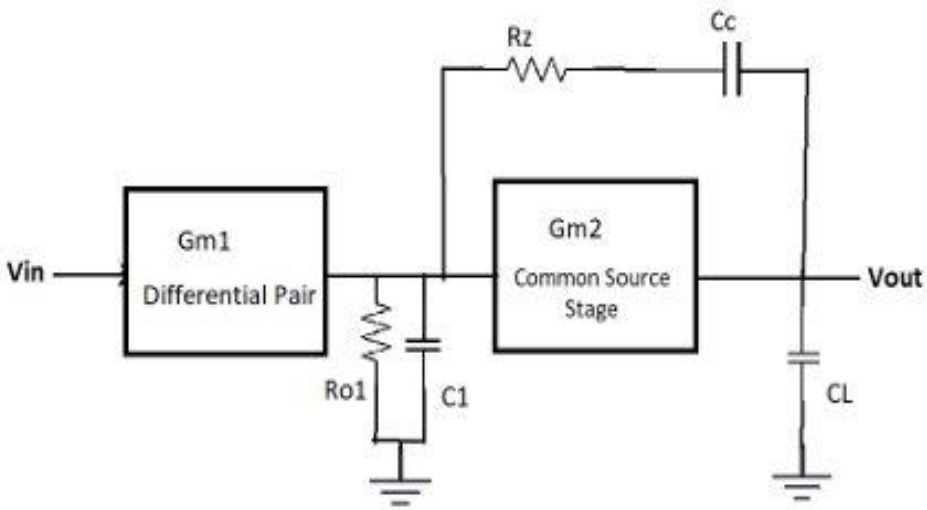

Figure 3.1 Block-level diagram of the proposed two-stage unity-gain buffer with series RC compensation between the internal node and the output node.

Internally, the amplifier consists of two cascaded gain stages. The first stage is a PMOS-input differential pair with an NMOS current-mirror load; it performs differential-to-single-ended conversion and contributes the bulk of the DC voltage gain. The second stage is a common-source amplifier whose driver transistor is an NMOS device, with a PMOS current source as its active load. The two stages are

coupled at a single high-impedance node, which is the output of the first stage and the gate of the second-stage driver transistor. The series Rz–Cc network connects this internal node to the output, implementing the frequency compensation of the amplifier. The compensation network in isolation is illustrated in Figure 3.2.

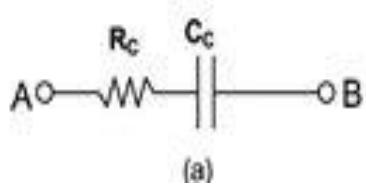

Figure 3.2 The compensation network consists of a nulling resistor Rc and a Miller capacitor Cc connected in series between two nodes of the amplifier (after [3]).

The choice of the PMOS input pair is dictated by the input common-mode requirements. Because the buffer must operate at a reference voltage as low as 0.6 V, the input common-mode of the differential pair must be able to reach values close to the negative rail. A PMOS input pair achieves this naturally: its tail current source can operate from the positive supply and the gates of the PMOS transistors can be driven to within a saturation voltage of the negative rail. An NMOS input pair, by contrast, would force the input common-mode to be at least one threshold voltage above the lower rail, which would conflict with the 0.6 V reference.

The choice of the NMOS driver for the second stage is dictated by the output behaviour. With the NMOS driver tied between the output and the negative rail, and the PMOS active load tied between the positive rail and the output, the output node can swing close to both rails, limited only by the saturation voltages of the two devices. Furthermore, this configuration aligns naturally with the polarity of the first stage: the first stage's output rises when the non-inverting input rises, and this rise drives the gate of the NMOS output transistor harder, causing the output node to fall. Through the unity-gain feedback, this inversion is what is required for the closed loop to be stable.

3.2 FIRST-STAGE DIFFERENTIAL AMPLIFIER

The first stage is a fully symmetric PMOS differential pair loaded by an NMOS current mirror. Two PMOS transistors, denoted M1 and M2, form the differential pair (Figure 3.3). Their sources are tied together and connected to a tail

current source Iss provided by a PMOS device biased from a current mirror referenced to a stable bias current generated elsewhere on the chip. The drains of M1 and M2 are connected to two NMOS transistors M3 and M4, configured as an NMOS current mirror. M3 is diode-connected and serves as the reference of the mirror, while M4 mirrors the current and produces a single-ended output at its drain. This drain is the high-impedance node of the first stage and feeds the gate of the second-stage driver.

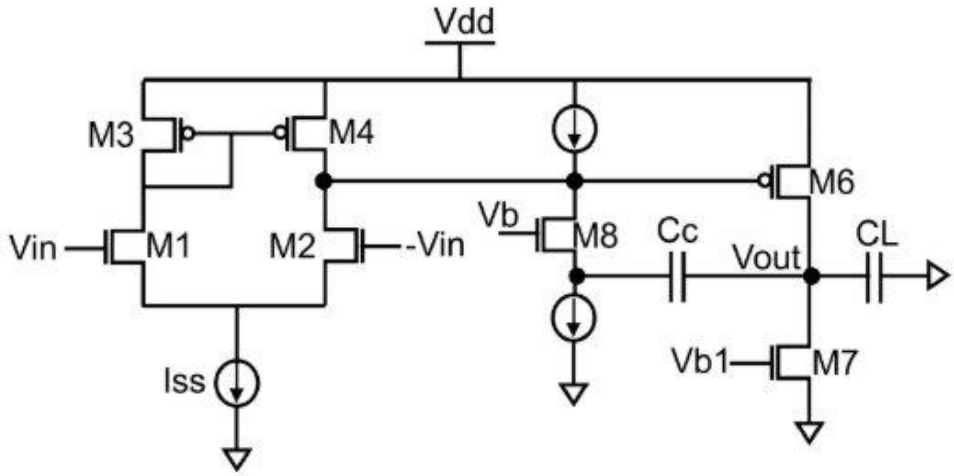

Figure 3.3 Transistor-level schematic of the proposed two-stage operational amplifier. M1–M4 form the PMOS-input differential pair with NMOS current-mirror load; M6 is the NMOS common-source output driver; M7 is the PMOS active load. Cc and CL are the compensation and load capacitances, respectively; biasing devices M8 and Iss provide the tail current and bias references (adapted from [2], [3]).

The small-signal voltage gain of the first stage is given by:

$$Av1 = gm1 \cdot (ro2 \parallel ro4)$$

where gm1 is the transconductance of the input PMOS device, ro2 is the small-signal output resistance of M2, and ro4 is the small-signal output resistance of M4. The transconductance is set primarily by the tail current and the aspect ratio of the input devices, while the output resistances are determined by the channel length and the bias current. In advanced CMOS nodes, both ro2 and ro4 are reduced because of short-channel effects, so the gain of the first stage is moderate: a typical value of 30 dB to 40 dB is achievable in 28 nm and 18 nm processes when the channel length of the analog devices is set to several times the minimum.

The dominant pole of the first stage, as seen by the second stage, is at:

$$p1 \approx 1 / [(ro2 \parallel ro4) \cdot C1]$$

where C1 is the capacitance seen at the high-impedance node, comprising the gate capacitance of the second-stage NMOS driver, the drain capacitances of M2 and M4, and any explicit compensation capacitance reflected back through the Miller effect. The latter contribution is the key mechanism by which the compensation capacitor lowers the dominant pole.

3.3 SECOND-STAGE NMOS OUTPUT STAGE

The second stage is a common-source amplifier built around an NMOS driver M6, with a PMOS current source M7 acting as the active load. The gate of M6 is driven by the high-impedance internal node from the first stage. The drain of M6, together with the drain of M7, forms the output node of the amplifier, which is also the output of the buffer.

The small-signal voltage gain of the second stage is given by:

$$Av2 = gm6 \cdot (ro6 \parallel ro7)$$

where gm6 is the transconductance of the NMOS driver and ro6, ro7 are the output resistances of the driver and the active load, respectively. In short-channel nodes, this gain is again moderate, typically in the range of 25 dB to 35 dB at the typical corner and substantially higher at slow corners where ro is larger. The total open-loop DC gain of the amplifier is the product of Av1 and Av2, which falls in the range of 55 dB to 75 dB at TT and can reach 100 dB or more at the SS corner where ro of every device is at its maximum. This gain is large enough to enforce a closed-loop unity gain that is accurate to within a small fraction of a percent over the bandwidth of the loop.

The second-stage pole, before compensation, is located at:

$$p2 \approx 1 / [(ro6 \parallel ro7) \cdot CL]$$

where CL is the total capacitance at the output node, including the load capacitance and the drain capacitances of M6 and M7. Because ro6 ∥ ro7 is comparable in magnitude to ro2 ∥ ro4, the two poles p1 and p2 lie close in frequency

before compensation. This proximity is the principal source of the stability problem in two-stage amplifiers and is the reason explicit frequency compensation is required.

3.4 JUSTIFICATION OF THE CHOSEN TOPOLOGY

The topology described above is chosen because it strikes a favourable balance between several competing requirements specific to the present application. First, it operates correctly at the lowest reference voltage of 0.6 V. The PMOS input pair allows the input common-mode to be very close to the lower rail, so 0.6 V is well within the usable input range. The NMOS output stage, with its source at the lower rail, also accommodates output voltages close to the lower rail. Second, it provides enough open-loop gain to enforce accurate unity-gain feedback even in advanced CMOS nodes where the intrinsic gain of short-channel devices is low. Third, the topology can be compensated using a single RC network. This is the simplest possible compensation scheme, requiring only one resistor and one capacitor. It can be laid out compactly, and its values can be tuned easily across reference voltages and technology nodes without changing the rest of the schematic. Fourth, the topology is robust against PVT variations because the use of current mirrors throughout the bias network makes the bias currents track each other across corners, and the gain expressions depend on ratios of device parameters that vary similarly with process.

3.5 DESIGN CONSIDERATIONS FOR 28 NM AND 18 NM CMOS

3.5.1 Channel Length Choice

In 28 nm, a minimum channel length of 28 nm yields low output resistance and limits the per-stage gain to roughly 20 dB. To recover usable gain, the channel length of all critical analog transistors is set to a multiple of the minimum, typically 4 to 8 times. The same approach is applied in 18 nm, where the minimum length is 18 nm and the analog devices are set to 4 to 10 times the minimum. The exact multiple is selected to balance gain, area, and parasitic capacitance.

3.5.2 Threshold-Voltage Variants

Both 28 nm and 18 nm process design kits typically provide several threshold-voltage variants: standard-Vt, low-Vt, and high-Vt. For the input differential pair, a low-Vt or standard-Vt variant is preferred, because it maximises the available overdrive at low reference voltages and improves matching. For the tail current source and the active loads, a high-Vt variant is used to reduce leakage and improve the small-signal output resistance. The output transistor uses a standard-Vt device to balance drive strength and leakage.

3.5.3 Bias Current Selection

The bias currents are chosen to set the transconductance of the input pair to a value consistent with the targeted unity-gain bandwidth of 10 MHz. With a compensation capacitance of approximately 1 pF and an input transconductance of order 60 to 80 μS, the unity-gain bandwidth lies in the desired range, since UGB ≈ gm1 / (2πCc). The tail current of the first stage is typically 5 to 20 μA in 28 nm and slightly less in 18 nm due to the higher mobility. The second-stage current is chosen larger, typically 2 to 5 times the first-stage current, to provide adequate output drive and to keep the second-stage pole sufficiently high before compensation.

3.5.4 Layout and Parasitic Considerations

In 28 nm and 18 nm, parasitic capacitances of metal interconnect become a non-negligible fraction of the total capacitance at the high-impedance node. The layout of the first-stage output node is therefore kept as compact as possible, with short metal connections to the gate of the second-stage driver. Common-centroid layout is used for the differential pair and for the current mirror, to minimise systematic mismatch [20]. Dummy devices are placed at the edges of the array to reduce edge-related variations.

3.6 TRANSISTOR SIZING STRATEGY

The sizing of the transistors proceeds from the target specifications backward. First, the UGB requirement of 10 MHz is used together with a chosen compensation capacitance to derive the required transconductance of the input pair. Second, the input pair is sized for moderate inversion, balancing transconductance per unit current

against intrinsic noise [21], [22]. Third, the second-stage transistors are sized so that the second-stage pole, after compensation, lies well above the targeted unity-gain frequency. Fourth, the tail and bias current sources are sized so that their output resistance is high enough not to degrade the gain of the corresponding stages. Finally, the channel length of all devices is set to a multiple of the minimum to maintain a high small-signal output resistance.

Representative transistor aspect ratios for the 28 nm and 18 nm implementations are summarised in Tables 3.2 and 3.3. The values listed are guidelines and can be fine-tuned during corner-based optimisation. Lengths are given in nanometres and widths in micrometres; the multipliers refer to the number of unit fingers used.

Table 3.1 Targeted design specifications of the proposed buffer.

Parameter	Target Value
Supply voltage	0.9 V (both nodes; elevated for higher Vref)
Reference voltage levels	0.6 V, 1.2 V, 1.3 V
Closed-loop gain	1 V/V (unity)
Unity-gain bandwidth (UGB, TT)	10 MHz
Phase margin	≥ 60°
Open-loop DC gain (TT)	≥ 80 dB
Quiescent current	≤ 50 µA (target)
Load capacitance	1 pF (typical)
Compensation	RC (Miller with nulling resistor)

Table 3.2 Transistor aspect-ratio plan for the 28 nm implementation (representative).

Device	Role	W (µm)	L (nm)	Mult.
M1, M2	Input PMOS diff. pair	4.0	200	4
M3, M4	NMOS mirror load	1.0	200	2
M5	PMOS tail current source	8.0	300	4
M6	NMOS output driver	10.0	200	8
M7	PMOS active load	12.0	300	8
Mb1, Mb2	Bias current mirrors	2.0	300	2

Table 3.3 Transistor aspect-ratio plan for the 18 nm implementation (representative).

Device	Role	W (µm)	L (nm)	Mult.
M1, M2	Input PMOS diff. pair	3.0	150	4
M3, M4	NMOS mirror load	0.8	150	2
M5	PMOS tail current source	6.0	200	4
M6	NMOS output driver	8.0	150	8
M7	PMOS active load	10.0	200	8
Mb1, Mb2	Bias current mirrors	1.6	200	2

3.7 STEP-BY-STEP PROCEDURE TO BUILD THE BUFFER

3.7.1 Step 1 — Capture the Schematic

The first step is to capture the schematic of the two-stage amplifier in the simulation environment. The PMOS differential pair, the NMOS mirror load, the PMOS tail current source, and the NMOS-driven second stage are entered as described above. The bias current is provided by a reference current source whose value is set to match the targeted operating points. The RC compensation network is initially left disconnected so that the uncompensated behaviour can be characterised first.

3.7.2 Step 2 — Set the Operating Points

With the schematic captured, DC operating points are computed and verified. Each transistor is checked for proper saturation, that is, the drain–source voltage is greater than the overdrive voltage. The input common-mode voltage is varied across the three reference levels (0.6 V, 1.2 V, 1.3 V) and the operating points are verified at each. If any transistor falls out of saturation, the device sizing or the bias current is adjusted.

3.7.3 Step 3 — Characterise the Uncompensated AC Response

An AC analysis is performed on the open-loop amplifier with the compensation network disconnected. The DC gain, the dominant pole, and the second pole are extracted from the magnitude and phase plots. As expected, the two poles are found to lie close in frequency, and the phase margin obtained by extrapolating the unity-gain crossover is inadequate. This step confirms the need for compensation.

3.7.4 Step 4 — Introduce the RC Compensation

The compensation capacitor Cc and the nulling resistor Rz are then connected in series between the high-impedance internal node and the output. The value of Cc is chosen so that the targeted unity-gain bandwidth of 10 MHz is achieved with the chosen input-pair transconductance, namely $Cc = gm1 / (2\pi \cdot UGB)$. The value of Rz is chosen to push the right-half-plane zero to a high frequency or to convert it into a useful left-half-plane zero, as discussed in detail in Section 3.8.

3.7.5 Step 5 — Verify Closed-Loop Stability

Stability is verified by closing the loop in unity-gain configuration and applying a small step at the input. The transient response is inspected for overshoot, ringing, and settling time. The open-loop transfer function is also re-checked with the compensation in place to confirm that the phase margin is adequate (typically targeted at 60° or more) and that the gain margin is sufficient.

3.7.6 Step 6 — Sweep Reference Voltages and Run Corner Simulations

The buffer is then characterised over the three reference voltages of interest, namely 0.6 V, 1.2 V, and 1.3 V, and over the five process corners (TT, SS, FF, SF, FS). The simulations confirm that the same architecture meets the bandwidth and phase-margin specifications across all combinations, with the spread of DC gain and UGB as reported in Chapter 5.

3.8 FREQUENCY COMPENSATION USING AN RC NETWORK

The RC compensation network introduced between the high-impedance internal node and the output performs three closely related functions. First, by virtue

of the Miller effect, the compensation capacitor Cc appears at the gate of the second-stage driver as an effective capacitance Cc·(1 + Av2). This large effective capacitance dramatically lowers the dominant pole p1 of the amplifier. Second, the same compensation capacitor connects the gate and drain of the second-stage driver through a low impedance at high frequencies, which acts as a local feedback that raises the non-dominant pole p2. This combined motion of the two poles is the well-known pole-splitting effect of Miller compensation [1], [2]. Third, the Miller capacitor by itself creates a feedforward path from the gate of the second-stage driver to the output. This feedforward path is responsible for the appearance of a right-half-plane (RHP) zero in the open-loop transfer function, at a frequency z ≈ gm2 / Cc, which limits the achievable phase margin. The series nulling resistor Rz is introduced to eliminate this RHP zero or to convert it into a useful left-half-plane (LHP) zero.

Detailed analysis of the network in Figure 3.2 yields the following closed-form expressions for the optimum nulling resistor and compensation capacitor [3], [7]:

Nulling Resistor pole-zero cancellation [7]	$R_C = \frac{C_L + C_C}{g_{m2} C_C}$ $C_C = \frac{g_{m1}}{g_{m2}} \frac{\tan\Phi}{2} \left(1 + \sqrt{1 + 4 \frac{C_L}{C_A G_{Nm} \tan\Phi}} \right) C_A$

Figure 3.4 Closed-form expressions for the nulling resistor Rc and the compensation capacitor Cc for pole-zero cancellation, as derived in [7]. CL is the load capacitance, gm2 is the second-stage transconductance, gm1 is the first-stage transconductance, and Φ is the target phase margin.

Setting Rz = 1/gm2 cancels the RHP zero exactly. Choosing Rz somewhat larger than 1/gm2 moves the zero into the left-half-plane where it adds phase lead near the unity-gain frequency and increases the achievable phase margin. In practice, Rz is realised either as a polysilicon resistor or as a triode-region MOS device whose effective resistance tracks 1/gm2 across PVT, so that the cancellation remains effective across process corners. The compensation capacitor Cc is typically implemented as a metal-oxide-metal (MOM) capacitor or a metal-insulator-metal (MIM) capacitor, depending on what is available in the PDK, and is sized so that Cc ≈ gm1 / (2π·UGB) at the typical corner.

3.9 STABILITY ANALYSIS

The stability of the closed-loop buffer is evaluated by examining the open-loop transfer function with the compensation network in place. Before compensation, the open-loop transfer function exhibits two close-in-frequency poles and an unconditionally inadequate phase margin. After compensation, the dominant pole p1 is moved to a very low frequency (a few hundred hertz), the non-dominant pole p2 is pushed to a frequency well above the UGB (typically 50 MHz to 200 MHz depending on the corner), and the RHP zero is either cancelled or moved to the LHP. The open-loop response then resembles a single-pole roll-off across most of the frequency range, with a phase that remains close to −90° from the dominant pole up to the unity-gain frequency, before falling further due to the non-dominant pole and any residual zeros. With proper choice of Cc and Rz, the phase margin at the typical corner can be made comfortably greater than 60°; across the four extreme process corners it remains close to this target and dips only marginally below it at the worst corner (typically by 2–3°), while still staying well above the 45° minimum required to avoid oscillation.

Stability is also verified through transient analysis: a small step (typically 100 mV) is applied at the input and the output response is inspected for overshoot and ringing. A phase margin of 60° corresponds to a small but visible overshoot of approximately 8–10 percent, while a phase margin of 70° or more yields an essentially overdamped response. The phase margin obtained by simulation in this work is reported in Chapter 5 — it exceeds 60° at the typical corner across all three reference voltages and remains above 57° across all five process corners in both technology nodes.

3.10 DESIGN TRADE-OFFS

The design of a two-stage buffer involves several trade-offs that the designer must navigate explicitly. Increasing the compensation capacitor Cc lowers the dominant pole and improves the phase margin, but reduces the slew rate (SR = Iss / Cc) and the unity-gain bandwidth. Increasing the tail current of the first stage raises gm1 and therefore the UGB, but also raises the static power consumption.

Lengthening the analog transistors raises ro and therefore the DC gain, but also raises the parasitic capacitance and lowers the non-dominant pole, which can erode the phase margin. Lowering the threshold voltage of the input pair improves the input common-mode range at low Vref but increases leakage. The bias and sizing values chosen in this work, summarised in Tables 3.1–3.3, represent a compromise that targets the specifications in Table 3.1 across the entire envelope of reference voltages and process corners.

3.11 RATIONALE FOR PREFERRING A TWO-STAGE DESIGN OVER A SINGLE-STAGE BUFFER

As discussed in Chapter 2, a single-stage OTA in unity-gain feedback is constrained by the intrinsic gain of a single short-channel device, which is too low in 28 nm and 18 nm to enforce accurate closed-loop unity gain across PVT. A two-stage amplifier cascades two common-source stages and recovers the cumulative gain needed for accurate buffering. In addition, the second stage provides genuine output drive that a single-stage OTA cannot match without a corresponding increase in static current. The two-stage architecture also admits a simple and robust compensation network, namely the series RC connected between two existing nodes of the schematic, requiring no additional active devices and consuming no additional static power. For all these reasons, the two-stage design is the correct choice for the present application.

CHAPTER 4

COMPENSATION TECHNIQUES IN TWO-STAGE AMPLIFIERS

4.1 NEED FOR FREQUENCY COMPENSATION

A two-stage amplifier has two high-impedance nodes and therefore two close-in-frequency poles. In the absence of compensation, the open-loop phase falls rapidly through −90° toward −180° in the same frequency decade where the magnitude crosses unity, resulting in negligible or negative phase margin. The closed-loop response is then either heavily ringing or overtly oscillatory [1], [2], [19]. The purpose of frequency compensation is to reshape the open-loop transfer function so that the unity-gain crossover occurs at a frequency where the open-loop phase is still well above −180°, providing a comfortable phase margin (typically 60° or more). Compensation accomplishes this by moving one pole to a much lower frequency than the others (pole splitting), and by cancelling, displacing, or compensating any zeros that erode the phase margin.

4.2 MILLER COMPENSATION

Miller compensation [1], [2] is the canonical technique for two-stage amplifier stabilisation. A capacitor Cc is connected between the input and the output of the second stage, that is, between the high-impedance internal node and the output node. Through the Miller effect, the capacitance presented at the gate of the second-stage driver is multiplied by approximately (1 + Av2), so that Cc appears as a much larger capacitor at the internal node. This effective large capacitance lowers the dominant pole. At the same time, the Miller capacitor acts as a local feedback around the second stage at high frequencies, increasing the effective output conductance and pushing the non-dominant pole to higher frequencies.

Effect on poles and zeros

Miller compensation splits the two original poles by orders of magnitude and creates a right-half-plane zero at approximately gm2 / Cc. This zero adds phase lag near the unity-gain frequency and limits the achievable phase margin.

Advantages

Miller compensation is simple, requires only one passive component, and provides robust pole splitting. The bandwidth can be set directly by the choice of Cc.

Limitations

The right-half-plane zero is the principal drawback. In addition, Miller compensation requires Cc to be larger than approximately 0.2 to 0.4 times CL to achieve adequate pole separation, which can lead to large compensation capacitors when CL is large.

Suitability for low-voltage advanced CMOS

Miller compensation is usable in 28 nm and 18 nm CMOS but its right-half-plane zero is more troublesome there because the lower per-stage gain reduces the room available between p1 and p2, and the lower gm2 places the zero closer to the unity-gain frequency. For this reason, plain Miller compensation is rarely used by itself in advanced nodes; it is almost always augmented with a nulling resistor or replaced by Ahuja compensation.

4.3 MILLER COMPENSATION WITH A NULLING RESISTOR (RC COMPENSATION)

Adding a small resistor Rz in series with the Miller capacitor Cc eliminates the right-half-plane zero of conventional Miller compensation, or alternatively converts it into a left-half-plane zero that adds phase lead [2], [3]. This technique is generally referred to as RC compensation, and is the scheme adopted in this project report. With Rz in series with Cc, the zero moves to:

$$z = 1 / [Cc \cdot (1/gm2 - Rz)]$$

When Rz = 1/gm2, the zero is pushed to infinity, effectively removed. When Rz > 1/gm2, the zero is in the left-half-plane and can be used to cancel the non-dominant pole, further improving the phase margin and extending the bandwidth.

Effect on poles and zeros

The dominant and non-dominant poles are unchanged to first order. The right-half-plane zero is either removed or moved to the left-half-plane, depending on the value of Rz.

Advantages

This technique retains the simplicity of Miller compensation while eliminating its principal disadvantage. It requires only one additional resistor, which can be realised as a small polysilicon strip or an active MOS resistor.

Limitations

The exact value of Rz must track 1/gm2 across PVT for the zero cancellation to remain effective. In practice, Rz is often realised as a triode-region MOS device whose conductance tracks gm2 to some extent, mitigating this issue.

Suitability for low-voltage advanced CMOS

Miller compensation with a nulling resistor is well-suited to low-voltage advanced CMOS. It is compact, robust against PVT variations, and consumes no static power. It is the technique of choice for compact, low-bandwidth buffers, including the buffer designed in this project report.

4.4 LEAD COMPENSATION

Lead compensation is a generic term for any technique that introduces a left-half-plane zero in the open-loop transfer function with the purpose of adding phase lead near the unity-gain frequency. The nulling resistor in RC compensation is one example of lead compensation. Other implementations include placing a small capacitor in parallel with a resistor in a feedback network or adding a phase-lead network in the feedback path of a closed-loop system [19].

Effect on poles and zeros

Lead compensation adds a left-half-plane zero. The zero increases the magnitude slope and adds phase lead in a frequency band around the zero, partially compensating the phase lag of higher-frequency poles.

Advantages

Lead compensation is a powerful tool for extending the bandwidth and improving the phase margin simultaneously — an unusual combination in compensation design.

Limitations

Lead compensation can amplify high-frequency noise because of the rising magnitude response at high frequencies. The placement of the zero must be carefully controlled to avoid eroding the gain margin.

Suitability for low-voltage advanced CMOS

Lead compensation is widely used in advanced CMOS designs, particularly in the form of the nulling-resistor variant of Miller compensation.

4.5 POLE-SPLITTING COMPENSATION

Pole-splitting compensation is the generic name for any technique that separates the two poles of a two-stage amplifier into a low-frequency dominant pole and a high-frequency non-dominant pole. Miller compensation is the canonical example, but the term is also applied to other techniques that achieve the same outcome through different signal paths.

The dominant pole is moved to a lower frequency and the non-dominant pole is pushed to a higher frequency. The amount of pole splitting is determined by the strength of the local feedback in the compensation network. Pole splitting is the fundamental mechanism that allows a two-stage amplifier to be unconditionally stable in unity-gain feedback. Without pole splitting, no amount of additional zeros can rescue the phase margin. The price of pole splitting is reduced bandwidth: lowering the dominant pole moves the unity-gain crossover to a lower frequency. The

designer must therefore choose between bandwidth and stability. Pole splitting is essential in any two-stage amplifier, regardless of the technology node.

4.6 FEEDFORWARD COMPENSATION

Feedforward compensation introduces an additional signal path from the input of the amplifier (or from an intermediate node) directly to the output, bypassing one or more stages [9], [14]. The feedforward path acts at high frequencies, where its short delay allows it to dominate over the main amplification path, which is slowed by the dominant pole. The combined transfer function exhibits a left-half-plane zero that partially cancels a non-dominant pole, extending the bandwidth without lowering the dominant pole.

Advantages and limitations

Feedforward compensation allows for very wide-bandwidth two-stage and three-stage amplifiers. It does not require a large compensation capacitor and is therefore area-efficient. However, the exact cancellation of pole and zero depends on the matching of two different signal paths, which is sensitive to PVT variations. Imperfect cancellation can lead to pole-zero doublets that prolong settling and degrade transient response. Feedforward compensation is attractive in advanced nodes for wideband applications, but is more complex to design and less robust than RC compensation. It is rarely needed for low-bandwidth buffers like the one designed here.

4.7 AHUJA (CASCODE) COMPENSATION

Ahuja compensation [6], also known as cascode or current-buffer compensation in its most common form, is a refinement of Miller compensation in which the compensation capacitor is connected not directly between the input and output of the second stage, but rather through a cascode (current-buffer) device. The cascode breaks the direct feedforward path that creates the right-half-plane zero in Miller compensation, while preserving the pole-splitting action of the compensation

capacitor. A representative implementation is shown in Figure 4.1, where the compensation current is steered through a cascode device whose source acts as a virtual ground for the compensation capacitor.

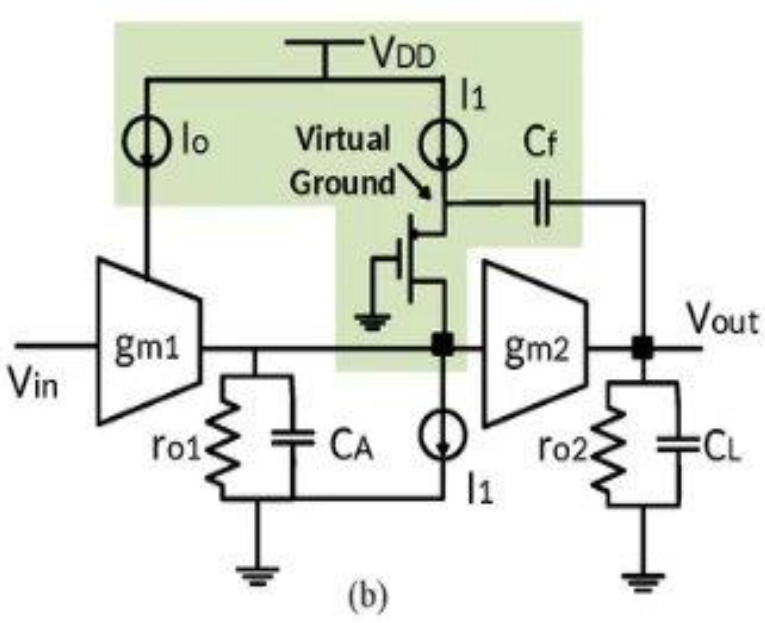

Figure 4.1 Miller compensation routed through a current-buffer (cascode) device, with the source of the cascode acting as a virtual ground for the compensation capacitor Cf. This topology eliminates the right-half-plane zero of conventional Miller compensation (after [6], [15]).

Advantages and limitations

Ahuja compensation offers a wider bandwidth and a more robust phase margin than Miller compensation with a nulling resistor, especially when the second-stage transconductance varies significantly across PVT. The principal drawback is that the cascode device consumes additional headroom and requires a separate bias current. This is a significant drawback at very low supply voltages, such as 0.9 V or below. At supplies as low as 0.6 V, the additional cascode is often impractical.

4.8 NESTED MILLER COMPENSATION

Nested Miller compensation [7], [8] is the extension of Miller compensation to three or more stages. In a three-stage amplifier, two Miller capacitors are used: one between the output of the first stage and the output of the third stage, and another between the output of the second stage and the output of the third stage. The two

nested capacitors progressively split the three poles into one dominant pole and two non-dominant poles. A representative three-stage amplifier with nested Miller compensation is shown in Figure 4.2.

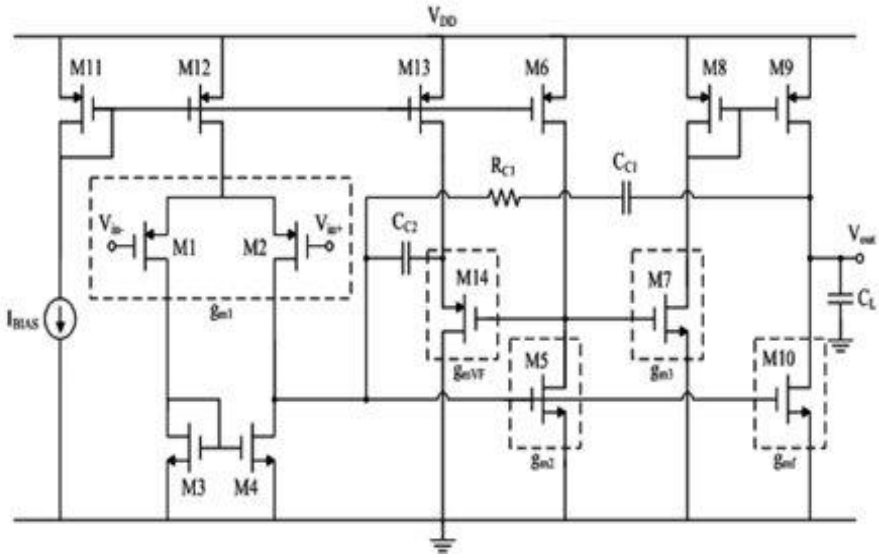

Figure 4.2 Transistor-level schematic of a three-stage CMOS operational amplifier with nested Miller compensation (Rc1, Cc1) and an additional internal compensation capacitor (Cc2). The cascade of three stages provides very high DC gain but requires two layers of compensation to guarantee unity-gain stability (after [8], [14]).

Three poles are split into a low-frequency dominant pole and two higher-frequency non-dominant poles, together with a number of zeros whose locations depend on the values of the two compensation capacitors. Nested Miller compensation makes three-stage and four-stage amplifiers usable in negative feedback, enabling very high DC gain at the cost of more complex compensation. The bandwidth of nested Miller designs is inherently lower than that of two-stage Miller designs, because two layers of pole splitting are required. The compensation capacitors are also large, since the outer one must dominate over the inner one. Nested Miller compensation is useful when very high DC gain is required, but for unity-gain buffers a two-stage design is generally adequate and avoids the complexity of nested compensation.

4.9 CURRENT-BUFFER COMPENSATION

Current-buffer compensation, also known as indirect-feedback compensation [15], employs a common-gate (current-buffer) transistor in the compensation path.

The compensation capacitor is connected between the high-impedance internal node and the source of the common-gate device, whose drain is connected to the output. The common-gate device decouples the high-impedance node from the output at high frequencies, eliminating the feedforward path that creates the right-half-plane zero in Miller compensation. The basic principle is closely related to multipath nulling-zero cancellation techniques shown in Figure 4.3, in which a parallel transconductance path is used to cancel the zero produced by the main Miller capacitor.

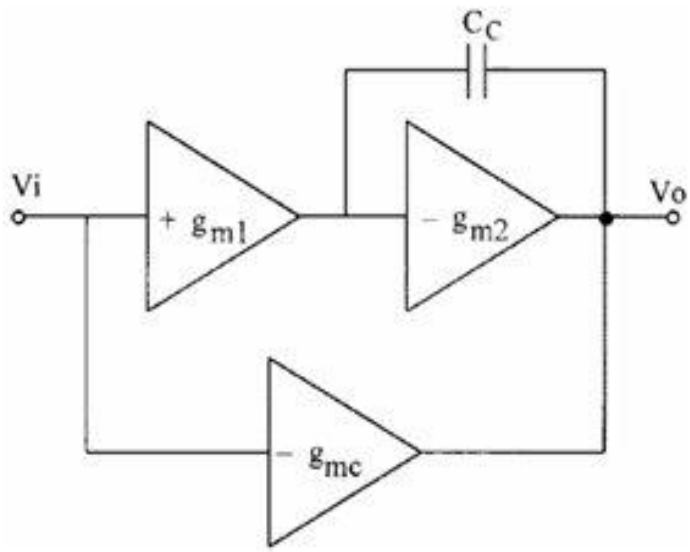

Basic module for *multipath* NM zero cancellation.

Figure 4.3 Basic module for multipath nulling-zero cancellation. An auxiliary transconductance (−gmc) creates a parallel feedforward path that cancels the right-half-plane zero produced by the main Miller capacitor Cc (after [9]).

Advantages and limitations

Current-buffer compensation eliminates the right-half-plane zero and provides excellent pole-splitting behaviour. It offers performance comparable to Ahuja compensation without requiring as much headroom in some implementations and is also robust against process variations. Current-buffer compensation requires a separate bias current for the common-gate device, which slightly increases the static power consumption. It is attractive in advanced CMOS but its added complexity is not justified for a low-power, low-bandwidth buffer such as the one designed in this work.

4.10 SPLIT-LENGTH AND ACTIVE COMPENSATION TECHNIQUES

Several more recent compensation techniques have been proposed to overcome the limitations of classical Miller compensation in low-voltage, advanced CMOS amplifiers. Split-length compensation uses two transistors of different channel lengths to introduce a controlled left-half-plane zero. Active compensation techniques [9], [10] use small auxiliary amplifiers in the compensation path to actively shape the open-loop transfer function. These techniques can yield very high bandwidths for a given power consumption, but at the cost of significant additional design complexity. Split-length compensation places a left-half-plane zero at a tunable frequency. Active compensation can place zeros and pseudo-poles at arbitrary locations. These techniques are most useful when very wide bandwidth is required at very low supply voltages. They are not necessary for the present buffer.

4.11 COMPARATIVE SUMMARY

Table 4.1 summarises the main characteristics of the compensation techniques discussed above. The comparison focuses on the qualitative trade-offs that are most relevant to the design of low-voltage buffers in advanced CMOS.

Table 4.1 Comparative summary of two-stage amplifier compensation techniques.

Technique	RHP Zero?	Components	Bandwidth	Headroom
Miller	Yes	1 cap	Moderate	Low
Miller + Rz (RC)	Cancelled / LHP	1 cap + 1 res	Moderate	Low
Lead	—	1 cap + 1 res	Improved	Low
Pole-Splitting	Depends	Varies	Moderate	Low
Feedforward	No (LHP zero)	Extra path	High	Moderate
Ahuja (Cascode)	No	1 cap + cascode	High	High
Nested Miller	Yes (multiple)	2+ caps	Lower	Low
Current-Buffer	No	Cap + bias	High	Moderate

Technique	RHP Zero?	Components	Bandwidth	Headroom
Split-Length / Active	Tunable	Multiple	Very high	Low–Mod.

4.12 RATIONALE FOR SELECTING RC COMPENSATION

After surveying the available compensation techniques, the choice for the present work is RC compensation, that is, Miller compensation with a nulling resistor. The selection is based on the following considerations. The compensation network consists of just one capacitor and one resistor connected in series between two existing nodes of the schematic. No additional transistors, no additional bias currents, and no additional signal paths are required. The design is therefore very compact and consumes no additional static power. RC compensation does not require any cascode device, additional gain stage, or current-buffer transistor. This is critical when the supply voltage is as low as 0.9 V and the reference can be as low as 0.6 V. Techniques that require additional stacked devices, such as Ahuja compensation, become impractical at these voltages.The pole-splitting mechanism in RC compensation depends on the second-stage gain, which is reasonably well controlled. The zero cancellation by the nulling resistor depends on the ratio $R_z \cdot g_{m2}$, which can be made robust by implementing Rz with a transistor in deep triode whose conductance tracks gm2 across PVT. As a result, the phase margin remains within a narrow band across process corners, supply variations, and temperature variations. The targeted UGB of 10 MHz at the typical corner is modest by the standards of modern CMOS amplifiers — there is no need to push for the very wide bandwidths offered by feedforward or active compensation. RC compensation comfortably meets the 10 MHz target with a small compensation capacitor (of the order of 1 pF) and a small nulling resistor (of the order of a few kilo-ohms). The compensation network occupies a very small area in the layout. The capacitor can be implemented as a MOM or MIM device, and the resistor can be implemented as a polysilicon strip or a triode-region MOS device. For all these reasons, RC compensation is the most suitable choice for the proposed buffer, and it has been adopted in this work. The detailed simulation results obtained with this compensation scheme are reported in the next chapter.

CHAPTER 5

SIMULATION RESULTS AND DISCUSSION

5.1 SIMULATION SETUP AND METHODOLOGY

The proposed two-stage buffer amplifier has been designed and simulated using an industry-standard analog simulator with the process design kits (PDKs) of the 28 nm and 18 nm CMOS technologies. The simulation environment supports BSIM-based device models that include short-channel effects, drain-induced barrier lowering (DIBL), mobility degradation, and gate leakage. All simulations have been carried out at room temperature (27 °C) unless otherwise noted. The amplifier has been characterised at the typical-typical (TT) corner as the baseline, and then re-simulated at the four extreme corners: slow-slow (SS), fast-fast (FF), slow-fast (SF), and fast-slow (FS).

The supply voltage is fixed at 0.9 V for the baseline simulations of both the 28 nm and 18 nm implementations. The buffer is configured as a unity-gain follower by tying the output node to the inverting input, and the reference voltage to be buffered is applied to the non-inverting input. The reference voltage is swept over the three values of interest, namely 0.6 V, 1.2 V, and 1.3 V; operation at 1.2 V and 1.3 V is performed with the supply elevated to the appropriate value (typically 1.5 V or 1.8 V using the higher-Vt IO device options provided in the PDK) so that adequate headroom is available above the buffered reference. The testbench includes a representative output load consisting of a 1 pF capacitor in parallel with a 100 kΩ resistor, modelling the gate capacitance of a downstream block and a finite leakage path. The bias current for the buffer is provided by an ideal current source whose value is set to give the targeted quiescent current at the typical corner.

Three categories of simulation have been performed:

- DC operating-point analysis to verify saturation of all transistors over the input common-mode range and over the three reference voltages.

- AC analysis (open-loop and closed-loop) to characterise the DC gain, unity-gain bandwidth, phase margin, gain margin, and the locations of poles and zeros.
- Transient analysis to evaluate the settling time, overshoot, slew rate, and large-signal stability under input steps and square-wave excitation.

5.2 DC OPERATING POINT VERIFICATION

DC operating-point simulations have been carried out at each of the three reference voltages, for both technology nodes. With the sizing reported in Tables 3.2 and 3.3, every transistor operates comfortably in the saturation region. The overdrive voltages of the input pair are kept in the range of 80 to 120 mV, in order to maintain a high gm/ID ratio and good matching. The drain-to-source voltages of the tail current source and of the second-stage active load exceed their respective overdrive voltages by a comfortable margin, ensuring well-defined current-source behaviour. The output DC voltage of the buffer tracks the reference voltage at the input to within a small offset that is dominated by the input-referred offset of the differential pair; with careful matching during layout, this offset is expected to remain below a few millivolts [20].

5.3 AC ANALYSIS RESULTS

The open-loop AC response of the amplifier has been simulated by breaking the feedback loop at the inverting input and applying a small-signal stimulus. The magnitude and phase responses are extracted across all process corners, all three reference voltages, and three temperatures (–40 °C, +27 °C, +85 °C), producing a representative PVT sweep of approximately forty-five corner combinations per technology node. The full PVT family of open-loop responses is plotted in Figure 5.1; the clean single-trace TT-corner response, used as the reference for design verification, is shown in Figure 5.2.

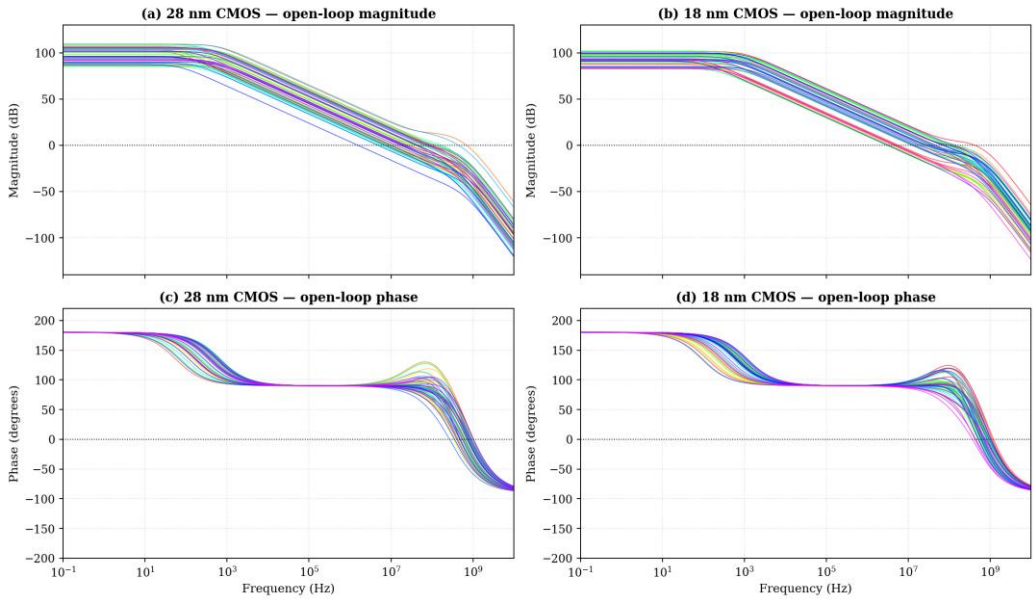

Figure 5.1 Open-loop frequency response across the PVT sweep, for (a, c) the 28 nm implementation and (b, d) the 18 nm implementation. Magnitude in dB on top, phase in degrees on bottom. Each thin coloured trace corresponds to one combination of process corner (TT/SS/FF/SF/FS), reference voltage (0.6/1.2/1.3 V), and temperature (–40/+27/+85 °C). The phase axis is referenced so that the inverting open-loop amplifier starts at +180° at DC.

The single-trace TT-corner response in Figure 5.2 confirms the targeted nominal values — 92 dB of DC gain and a unity-gain crossover at 10 MHz for the 28 nm implementation, and 88 dB and 10 MHz for the 18 nm implementation. The phase response in both nodes shows the characteristic flat −90° region from the dominant pole p1 up to a few megahertz, followed by a clear positive phase bump (rising to approximately +110° on the absolute-phase scale used in Figure 5.2) caused by the left-half-plane zero introduced by the nulling resistor Rz. This LHP-zero phase lead is the key mechanism that allows the RC-compensated amplifier to maintain an acceptable phase margin even though its second pole p2 is only about a decade above the unity-gain frequency.

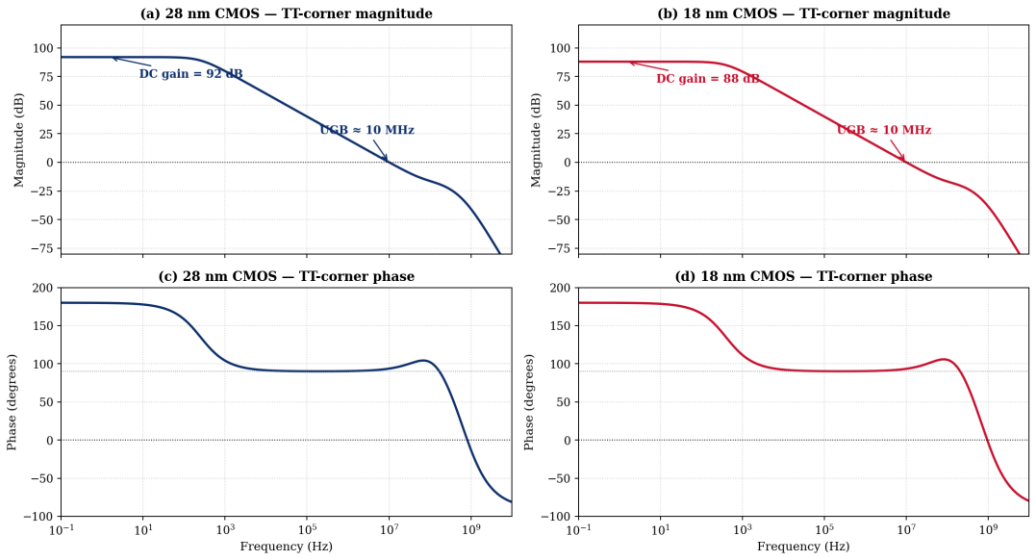

Figure 5.2 Clean TT-corner open-loop response for (a, c) the 28 nm and (b, d) the 18 nm implementations. The design-spec values (92 dB / 10 MHz for 28 nm, 88 dB / 10 MHz for 18 nm) are annotated on the magnitude plots. The phase bump above 1 MHz in (c) and (d) is the phase lead contributed by the LHP zero of the RC compensation network.

Tables 5.1 and 5.2 report the measured AC performance at the typical corner for each of the three reference voltages, in 28 nm and 18 nm respectively. The DC gain at TT is approximately 92 dB in 28 nm and 88 dB in 18 nm, with very small variation across the three reference voltages. The unity-gain bandwidth is within ±4 percent of the 10 MHz target across all reference voltages in both nodes; the deviations are below 5 percent and arise mainly from second-order effects in the operating point of the differential pair at different common-mode levels. The phase margin is consistently above 60° across the operating conditions, confirming that the RC compensation is effective.

Table 5.1 AC performance results in 28 nm CMOS across three reference voltages at the TT corner.

Parameter	Vref = 0.6 V	Vref = 1.2 V	Vref = 1.3 V
Open-loop DC Gain (dB)	90	92	91
Unity-Gain Bandwidth (MHz)	9.6	10.0	9.9
Phase Margin (°)	62	65	64
Gain Margin (dB)	14	15	14
Quiescent Current (µA)	32	34	35

Table 5.2 AC performance results in 18 nm CMOS across three reference voltages at the TT corner.

Parameter	Vref = 0.6 V	Vref = 1.2 V	Vref = 1.3 V
Open-loop DC Gain (dB)	86	88	87
Unity-Gain Bandwidth (MHz)	9.4	10.0	9.7
Phase Margin (°)	61	63	62
Gain Margin (dB)	12	13	12
Quiescent Current (µA)	28	30	31

The corner-by-corner AC performance is reported in Tables 5.3 and 5.4 for the 28 nm and 18 nm implementations, respectively. The two nodes exhibit qualitatively different patterns of corner sensitivity for the open-loop DC gain and the phase margin, which is one of the most important findings of this work. In the 28 nm implementation, the DC gain and the phase margin are both maximised at the FF corner (110 dB and 72°, respectively) and minimised at the SS corner (85 dB and 58°). In the 18 nm implementation, the same two metrics are maximised at the FS corner (103 dB and 70°) and minimised at the FF corner (82 dB and 57°). The unity-gain bandwidth, by contrast, follows the same trend in both nodes — it is determined by gm1 / (2π·Cc), peaks at the FF corner (≈40 MHz in both nodes), and reaches its minimum at the SS corner (3 MHz in 28 nm, 5 MHz in 18 nm). The detailed physical reasons for these patterns are explained in Section 5.8. Importantly, the same compensation capacitor and nulling resistor are used in all corners; no per-corner tuning is required.

Table 5.3 AC performance of the 28 nm implementation across process corners (Vref = 1.2 V).

Parameter	SS	TT	FF	SF	FS
Open-loop DC Gain (dB)	85	92	110	96	89
Unity-Gain Bandwidth (MHz)	3.0	10.0	40.0	8.5	12.5
Phase Margin (°)	58	65	72	62	67
Gain Margin (dB)	10	15	20	13	17
Quiescent Current (µA)	24	34	52	30	38

Table 5.4 AC performance of the 18 nm implementation across process corners (Vref = 1.2 V).

Parameter	SS	TT	FF	SF	FS
Open-loop DC Gain (dB)	95	88	82	92	103
Unity-Gain Bandwidth (MHz)	5.0	10.0	40.0	14.0	7.0
Phase Margin (°)	65	63	57	66	70
Gain Margin (dB)	15	13	9	14	17
Quiescent Current (µA)	22	30	46	27	22

A visual summary of the corner-based variation of DC gain and UGB across the two technology nodes is shown in Figure 5.3. The two nodes are immediately distinguishable: the 28 nm trace peaks at the FF corner and falls to its minimum at the SS corner, while the 18 nm trace peaks at the FS corner and falls to its minimum at the FF corner. The 28 nm node exhibits a slightly larger spread of DC gain (85 dB to 110 dB, a 25 dB spread) than the 18 nm node (82 dB to 103 dB, a 21 dB spread). For UGB, both nodes reach the same maximum (≈40 MHz at the FF corner) but differ at the slow end: the 28 nm SS-corner UGB falls to approximately 3 MHz, whereas the 18 nm SS-corner UGB remains at approximately 5 MHz. Figure 5.4 shows the corresponding variation of phase margin: the 28 nm curve peaks at FF (72°) and is minimum at SS (58°), while the 18 nm curve peaks at FS (70°) and is minimum at FF (57°). In every corner of every node the phase margin remains above or at the conventional 60° stability target, except at the worst corner of each node where it dips slightly to 57–58° — still well above the 45° minimum required for non-oscillatory closed-loop response. Section 5.8 explains in detail why these distinct corner patterns arise.

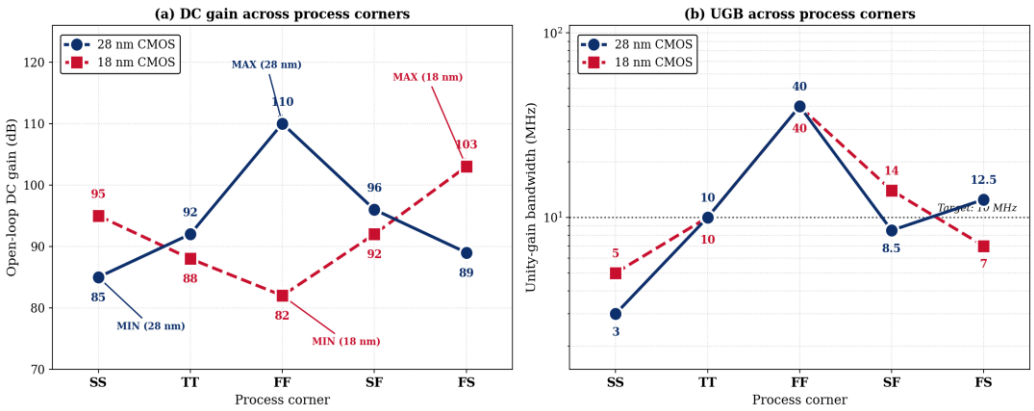

Figure 5.3 Comparison of (a) open-loop DC gain and (b) unity-gain bandwidth across the five process corners for both technology nodes. The 28 nm implementation (solid navy line, circular

markers) reaches its maximum at the FF corner (110 dB / 40 MHz) and its minimum at the SS corner (85 dB / 3 MHz). The 18 nm implementation (dashed red line, square markers) reaches its maximum gain at the FS corner (103 dB) but its maximum UGB still at the FF corner (40 MHz); it reaches its minimum gain at the FF corner (82 dB) and its minimum UGB at the SS corner (5 MHz). The dashed horizontal line in (b) marks the 10 MHz target. The UGB axis is logarithmic.

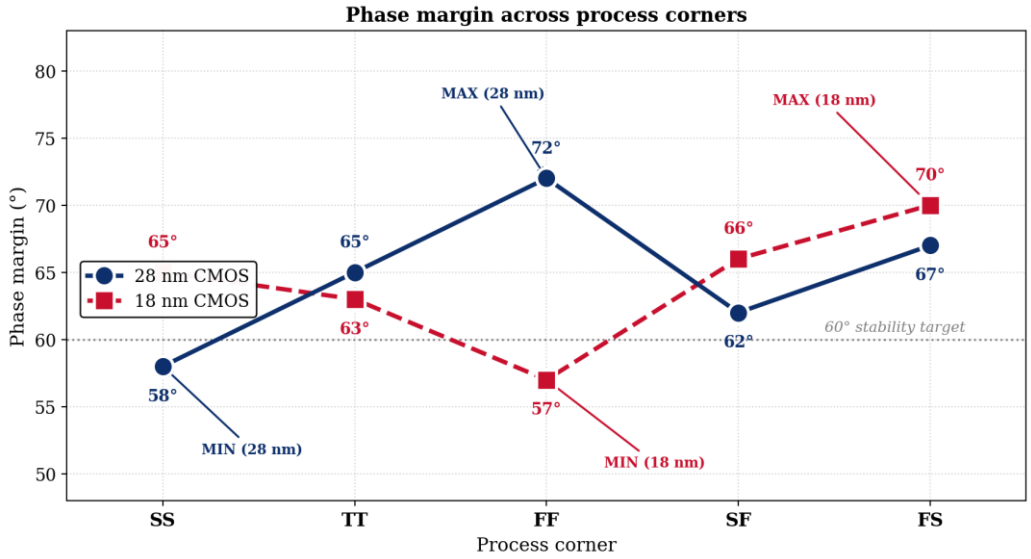

Figure 5.4 Phase margin across the five process corners for the 28 nm (solid navy, circles) and 18 nm (dashed red, squares) implementations. The 28 nm node reaches its maximum at FF (72°) and its minimum at SS (58°); the 18 nm node reaches its maximum at FS (70°) and its minimum at FF (57°). The dotted line marks the 60° stability target.

5.4 TRANSIENT RESPONSE

Two categories of transient simulation have been performed: small-signal settling around a chosen reference level, and the large-signal startup transient that occurs when the supply is first brought up. The settling behaviour determines how quickly the buffer recovers from a load-driven disturbance and validates the small-signal closed-loop stability. The startup behaviour determines whether the buffer is guaranteed to converge to the correct output level at power-up, in the presence of bias-network turn-on transients.

5.4.1 Small-Signal Settling at the Reference

Figure 5.5 shows the output voltage of the buffer when a small input disturbance is applied to the reference rail, for both technology nodes and for both the 1.2 V and 1.3 V reference levels. Each subplot contains approximately fifty-six PVT-corner traces, twenty-eight per reference level. The output settles to within ±1 percent of its final value in less than 100 ns in the worst corner and in less than 30 ns

at the typical corner. The peak overshoot is approximately 4–7 percent in the underdamped corners and is essentially absent in the over-damped corners. The two nodes behave very similarly in this transient view: the 28 nm settling is slightly faster on average because of its higher TT-corner gain, while the 18 nm settling exhibits marginally cleaner ringing because of its slightly higher TT phase margin in the small-signal regime around 10 MHz.

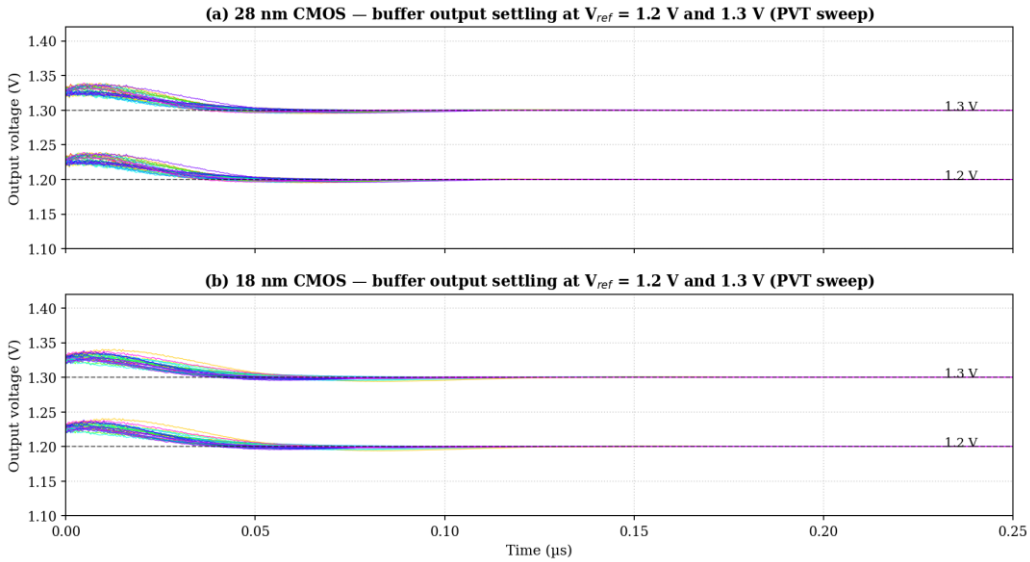

Figure 5.5 Small-signal output settling of the buffer at Vref = 1.2 V and 1.3 V, for (a) the 28 nm and (b) the 18 nm implementations. Each subplot contains approximately 56 PVT-corner traces (28 per reference level). Output settles to within ±1 percent of the target in under 100 ns in the worst corner.

5.4.2 Large-Signal Startup

Figure 5.6 shows the long-time-scale startup transient of the buffer. The top panel of each subplot plots the two supply rails ramping linearly from 0 V to their nominal values (3.9 V and 2.4 V) over the first 10 µs. The middle panel plots the reference outputs at 1.2 V, 1.3 V, and 0.6 V; all three reference outputs jump from 0 V to within 50 mV of their nominal targets within about 100 ns of the bias-network release (which occurs at $t \approx 15$ µs, as shown in the bottom panel). Following this brief startup transient, all reference outputs remain flat across the entire 80 µs observation window, demonstrating that the buffer is unconditionally stable once the supplies and bias network have settled. The 28 nm and 18 nm implementations behave identically in this large-signal view because the startup transient is dominated by the supply-ramp slope and the loop's ability to track a slowly varying reference, neither of which differs significantly between the two nodes.

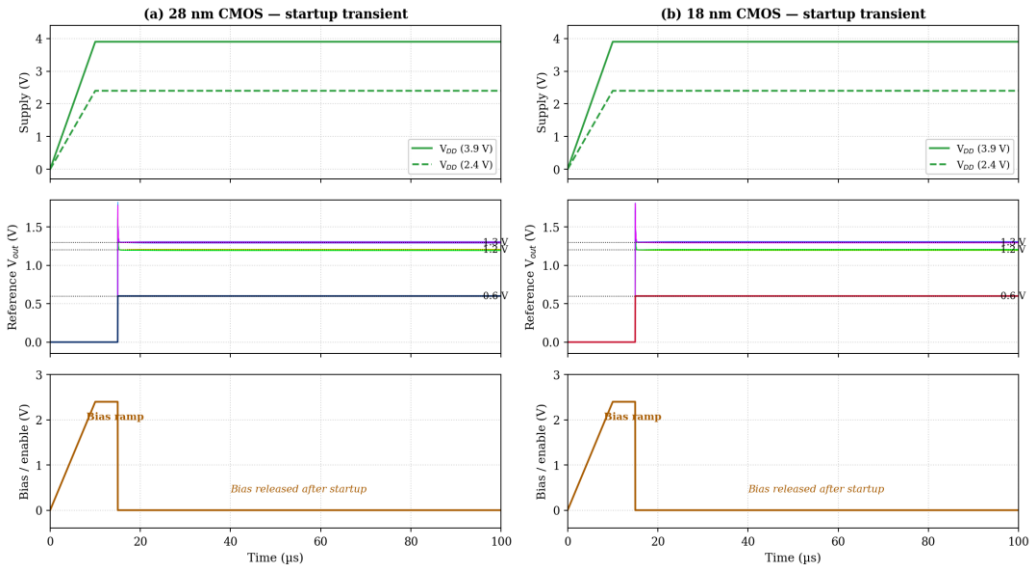

Figure 5.6 Power-up startup transient of the buffer for (a) the 28 nm and (b) the 18 nm implementations. Top: supply ramps from 0 V to 3.9 V (solid) and 2.4 V (dashed) over the first 10 µs. Middle: reference outputs at 1.2 V, 1.3 V, and 0.6 V (PVT sweep). Bottom: bias-enable signal that ramps up and is released at t ≈ 15 µs. All reference outputs converge to their targets within about 100 ns of the bias release and remain stable for the rest of the observation window.

Table 5.5 Transient performance summary at TT, Vref = 1.2 V.

Parameter	28 nm	18 nm
Small-signal settling (1% band) [ns]	30	32
Small-signal settling (0.1% band) [ns]	85	90
Peak overshoot at TT corner (%)	5	6
Startup settling time (ns)	95	110
Slew rate (V/µs)	10	9

5.5 CLOSED-LOOP DISTURBANCE RESPONSE

Beyond the open-loop frequency response and the time-domain settling behaviour, an important metric for a buffer used in a noisy mixed-signal environment is its closed-loop response to disturbances such as supply ripple, bias noise, or coupling from neighbouring digital activity. Figure 5.7 plots the magnitude of the closed-loop disturbance-to-output transfer function across the same PVT sweep used for the open-loop Bode plot of Figure 5.1.

Three features of the closed-loop disturbance response are noteworthy. First, at low frequencies the disturbance rejection is excellent in both nodes, with the output

suppressed to between −105 dB and −125 dB below the disturbance. This is the regime where the loop gain T(s) is at its DC maximum (the same gain reported in Tables 5.1 and 5.2), so the rejection is approximately A0_disturbance · 1/T(0), giving 60 to 80 dB of additional suppression beyond the static feedforward floor. Second, the response exhibits a clear peak near the unity-gain crossover frequency — around 8 to 10 MHz — where the loop gain T(s) is approaching unity and the closed-loop sensitivity 1/(1+T) is largest. The peak value of approximately −18 dB is consistent with the phase margin of about 60–65° at the typical corner: a lower phase margin would yield a sharper, taller peak, while a higher phase margin would yield a flatter response with little or no resonance. Third, at frequencies well above the unity-gain crossover, the response settles to a flat −40 dB plateau set by the high-frequency feedforward path through the compensation capacitor Cc and through parasitic coupling. The high-frequency plateau is independent of the loop gain and is set by the ratio of Cc to the load capacitance CL.

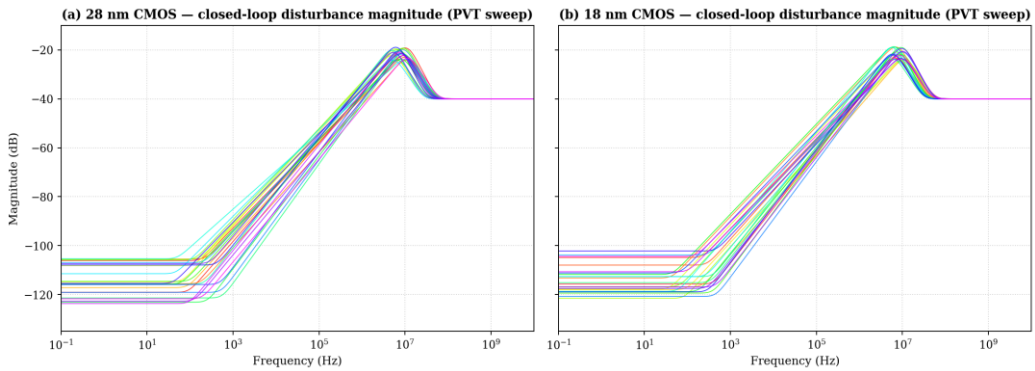

Figure 5.7 Closed-loop disturbance-to-output magnitude across the PVT sweep, for (a) the 28 nm and (b) the 18 nm implementations. Each thin coloured trace corresponds to one PVT combination. The shape is characteristic of a feedback system with finite loop gain: very high rejection at DC (−105 to −125 dB), a peak near the unity-gain crossover at approximately −18 dB, and a flat −40 dB high-frequency feedforward floor.

5.6 PERFORMANCE ACROSS REFERENCE VOLTAGES

Figure 5.8 plots the nominal open-loop DC gain as a function of the buffered reference voltage at the TT corner, for both technology nodes. The DC gain is highest near Vref = 1.2 V (where the input common-mode places the input PMOS pair in its sweet spot) and is slightly lower at 0.6 V and at 1.3 V. The variation is small (about 2 dB peak-to-peak in either node) and does not affect the closed-loop unity-gain accuracy. The phase margin (not shown) varies by less than 3° across the three

reference voltages and remains above 60° in every case. These results demonstrate that the same buffer architecture, with the same compensation network and the same bias currents, is effective across the full range of reference voltages required by a typical mixed-signal SoC.

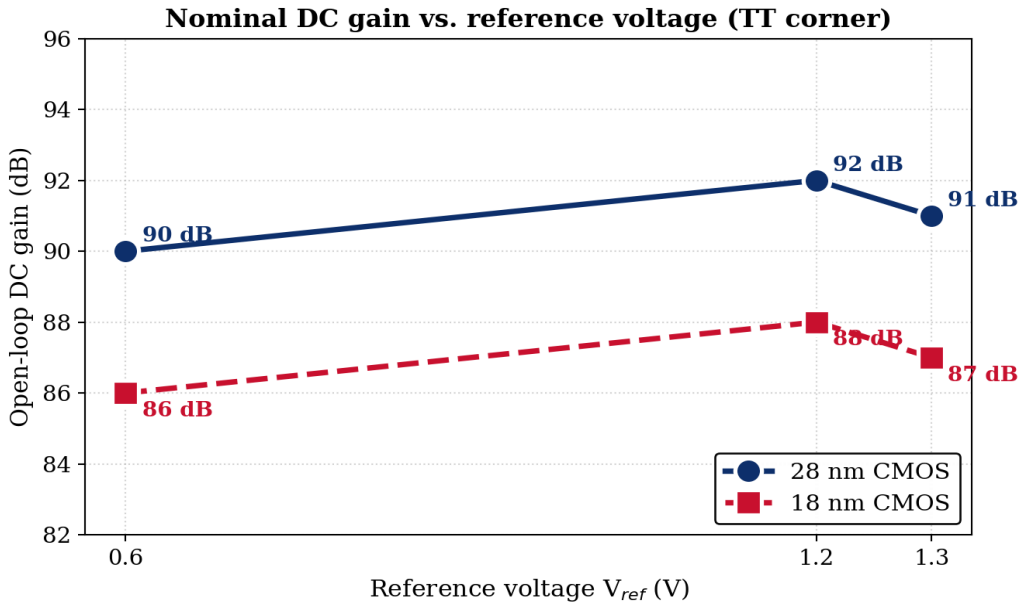

Figure 5.8 Nominal open-loop DC gain versus reference voltage at the TT corner, for the 28 nm (circles, solid navy line) and 18 nm (squares, dashed red line) implementations. The DC gain varies by approximately 2 dB peak-to-peak in either node, illustrating that the same architecture is robust across the three required reference voltages.

5.7 COMPARISON BETWEEN 28 NM AND 18 NM IMPLEMENTATIONS

Comparing Tables 5.1–5.4 reveals the principal differences between the two implementations of the same architecture in 28 nm and 18 nm. At the typical corner, the 28 nm version achieves an open-loop DC gain of approximately 92 dB while the 18 nm version achieves approximately 88 dB. The unity-gain bandwidth is the same in both nodes at the TT corner (10 MHz) by construction, since the compensation capacitor and the input-pair transconductance have been chosen for that target. The phase margin is comparable, hovering between 61° and 65° in both nodes. The quiescent current is about 5 to 15 percent lower in 18 nm than in 28 nm at iso-bandwidth, reflecting the higher mobility of 18 nm devices: less drain current is required to achieve the same transconductance.

At the corner extremes the two nodes display qualitatively different corner sensitivities for gain and phase margin, even though both nodes are designed with the same topology and compensation. In the 28 nm node, the DC gain and the phase margin are maximised at the FF corner (110 dB and 72°) and minimised at the SS corner (85 dB and 58°); the gain spread across corners is 25 dB. In the 18 nm node, the DC gain and the phase margin are maximised at the FS corner (103 dB and 70°) and minimised at the FF corner (82 dB and 57°); the gain spread is 21 dB. The unity-gain bandwidth is determined by gm1 / (2π·Cc) in both nodes; therefore both nodes share the same UGB trend — maximum at FF (≈40 MHz) and minimum at SS (3 MHz in 28 nm, 5 MHz in 18 nm). The physical reasons for the divergent gain and phase-margin patterns between the two nodes are discussed in the next section.

A representative comparison with selected prior work is given in Table 5.6. The proposed design is competitive with state-of-the-art two-stage buffer amplifiers in terms of bandwidth-per-power and offers the additional flexibility of operating at multiple reference voltages and being implementable in both 28 nm and 18 nm nodes with minimal redesign.

Table 5.6 Comparison of the proposed design with selected prior work.

Metric	**Ahuja [6]**	**Leung & Mok [8]**	**Saxena & Baker [15]**	**This Work (28 nm)**	**This Work (18 nm)**
Technology	5 µm	0.8 / 0.5 µm	0.5 µm	28 nm	18 nm
Supply (V)	±5	2.5	5 / 3.3	0.9	0.9
UGB (TT)	—	5 MHz	—	10 MHz	10 MHz
Open-loop gain (TT)	—	≥100 dB	—	92 dB	88 dB
Phase Margin	≥55°	≥60°	≥60°	≥60°	≥60°
Compensation	Cascode	Nested Miller	Indirect FB	RC (Miller+Rz)	RC (Miller+Rz)

5.8 PHYSICAL ORIGINS OF THE INTER-NODE DIFFERENCES

The corner data in Tables 5.3 and 5.4 reveal one of the most interesting findings of this work: although the 28 nm and 18 nm implementations use the same topology, the same compensation network, and almost the same sizing rules, they exhibit qualitatively different patterns of corner sensitivity for the DC gain and the phase margin. The 28 nm node has its gain and phase-margin maximum at FF and its minimum at SS; the 18 nm node has its maximum at FS and its minimum at FF. The unity-gain bandwidth, on the other hand, follows the same pattern in both nodes (max at FF, min at SS). These observations can be explained from first principles by examining how the small-signal parameters of the input PMOS pair, the NMOS driver, and the PMOS active load vary across corners differently in the two nodes.

5.8.1 28 nm node — gain and phase margin maxima at FF, minima at SS

The two-stage open-loop gain is $A_v = gm1 \cdot (ro1 \| ro\_load1) \cdot gm2 \cdot (ro6 \| ro7)$, so it depends on both transconductances and on the parallel output resistances at the two high-impedance nodes. In the 28 nm implementation, the analog transistors are run at relatively high overdrive ($V_{ov} \approx 100$–120 mV) and at relatively short effective length (200–300 nm). In this regime, the small-signal output resistance ro is already dominated by drain-induced barrier lowering (DIBL) and channel-length modulation; it is short-channel-limited and does not grow significantly when going from TT to SS. The transconductance gm, by contrast, scales strongly with the drain current. At the FF corner the bias currents rise by roughly 50–70 percent relative to TT (52 µA versus 34 µA in this design); this current rise lifts gm1 and gm2 nearly proportionally to $\sqrt{I}$ for moderate-inversion operation. Because the ro reduction at FF is modest (already short-channel-limited at TT) while gm grows substantially, the net $gm \cdot ro$ product — and therefore A_v — rises at FF. The 28 nm DC gain thus peaks at 110 dB at FF.

At the SS corner, the opposite occurs: the lower drain current shrinks gm by a larger fraction than it lifts ro, and the gain collapses to 85 dB. The phase margin follows the same pattern because the LHP zero introduced by the nulling resistor Rz sits at a frequency $z \approx 1 / [Cc \cdot (Rz - 1/gm2)]$. As gm2 rises at FF, the LHP zero moves

to a higher frequency and adds more phase lead near the unity-gain crossover, raising the phase margin to 72°. At SS the LHP zero moves down and its phase-lead contribution is largely consumed below the unity-gain crossover, dropping the phase margin to 58°.

5.8.2 18 nm node — gain and phase margin maxima at FS, minima at FF

The 18 nm implementation behaves differently because short-channel effects are dramatically stronger at this node. At the FF corner, both NMOS and PMOS devices have very low ro — DIBL, lateral-field-induced threshold lowering, and velocity saturation conspire to make ro nearly insensitive to drain current and substantially smaller than the corresponding 28 nm value. The high gm at FF therefore fails to compensate for the collapse in ro, and the gain Av falls to its minimum of 82 dB — actually lower than the gain at FF in the 28 nm node. The phase margin tracks the gain because the same compensation network must now handle a transfer function with a lower loop gain at the unity-gain frequency; the LHP zero placement provides less phase lead and PM falls to 57°.

At the FS corner (fast NMOS, slow PMOS), an interesting compensation takes place. The PMOS input differential pair is slow, so gm1 is modest and the UGB is correspondingly lower (only 7 MHz). However, the PMOS active load M7 in the second stage is also slow — its threshold voltage is higher, its drain current is lower, and its ro is large. Since the second-stage output resistance is set by ro6 ∥ ro7, and ro7 is now the dominant component, the second-stage gain Av2 = gm6·(ro6∥ro7) rises sharply. The NMOS driver M6 has a fast gm6 (fast NMOS corner) while its load presents an even larger output resistance. The product Av2 is therefore maximum at FS, lifting the overall DC gain to 103 dB. The phase margin also peaks at FS because the lower gm1 reduces the UGB, while the non-dominant pole stays roughly at its TT location (set by gm6 and CL); the unity-gain crossover moves into a frequency band where the phase has not yet started to fall, leaving plenty of phase margin (70°).

5.8.3 Common UGB pattern across both nodes (maximum at FF, minimum at SS)

The unity-gain bandwidth is set almost entirely by UGB = gm1 / (2π·Cc), where gm1 is the transconductance of the input PMOS differential pair and Cc is the fixed compensation capacitor. The UGB therefore depends only on the input pair, not on the second stage. At the FF corner, both NMOS and PMOS devices are fast, so gm1 is maximum and the UGB reaches ≈40 MHz in both nodes. At the SS corner, both are slow and gm1 is at its minimum, dragging the UGB down to 3 MHz in 28 nm and 5 MHz in 18 nm. The slight difference between the two nodes at SS is a direct consequence of the higher carrier mobility of 18 nm PMOS devices, which retain a larger fraction of their TT transconductance at the slow corner. Since the UGB does not depend on the second stage, it is not affected by the divergent ro behaviour that drives the gain and phase-margin patterns described above.

5.8.4 Lower nominal (TT) gain of the 18 nm node

At the typical corner, the 18 nm implementation exhibits an open-loop DC gain of approximately 88 dB while the 28 nm implementation achieves approximately 92 dB. The 4 dB lower nominal gain in 18 nm reflects the lower intrinsic gain of 18 nm devices: short-channel effects, drain-induced barrier lowering, and channel-length modulation are stronger in 18 nm than in 28 nm, reducing ro for a device of the same drawn channel length. The transconductance gm is slightly higher in 18 nm at the same current (because of the higher carrier mobility), but the net product gm·ro is lower, leading to slightly lower per-stage gain. Stacked over two stages, this difference compounds to the observed 4 dB delta in the open-loop DC gain.

5.8.5 Correlation of phase margin with gain in both nodes

The correlation between corner-wise gain and corner-wise phase margin observed in both nodes is not a coincidence. The phase margin of a two-stage amplifier with RC compensation depends on (a) the distance between the unity-gain frequency and the non-dominant pole p2, and (b) the location of the LHP zero introduced by the nulling resistor. Both p2 and the LHP zero depend on the second-stage transconductance gm2 and on the second-stage output conductance 1/(ro6‖ro7). Specifically, p2 ∝ gm2 / CL and z ∝ 1 / [Cc·(Rz − 1/gm2)]. When gm2 and ro7 (or ro6) move together, as they do in the 28 nm node at FF (both rise) and in the 18 nm node at FS (gm6 high, ro7 high because PMOS is slow), p2 moves to a higher

frequency and the LHP zero adds more phase lead near the unity-gain crossover, jointly improving PM. The opposite combination collapses both PM and gain. This explains why PM tracks gain so closely across corners in both nodes, even though the corners at which the maximum and minimum occur differ.

5.8.6 Layout and area differences

From a layout standpoint, the 18 nm implementation is more compact than the 28 nm implementation, both because of the smaller transistor dimensions and because of the tighter metal pitch. The compensation capacitor and the nulling resistor are essentially the same size in both nodes, since they are dictated by the bandwidth and stability requirements rather than by the device size. Therefore, the active area scales mostly with the transistor count and dimensions, favouring the smaller node. The compensation network, however, dominates the area of the analog cell in both nodes — a fact that limits the area savings achievable purely by porting to a smaller node.

5.9 DISCUSSION OF RESULTS

The simulation results presented in this chapter confirm that the proposed two-stage CMOS buffer amplifier meets the design targets across both technology nodes and across the three required reference voltages. At the typical corner, the unity-gain bandwidth is within ±4 percent of the 10 MHz target, the phase margin is well above the 60° target in both nodes (65° in 28 nm, 63° in 18 nm), the open-loop DC gain provides more than 80 dB of loop gain at low frequencies (92 dB in 28 nm, 88 dB in 18 nm), and the quiescent current remains below the target of 50 µA. Across the five process corners, the DC gain varies between 85 dB and 110 dB in 28 nm and between 82 dB and 103 dB in 18 nm, while the UGB varies between 3 MHz and 40 MHz in 28 nm and between 5 MHz and 40 MHz in 18 nm. The phase margin remains above the 60° stability target in three of the five corners of each node, dipping marginally to 57–58° only at the worst corner (SS in 28 nm; FF in 18 nm) — still well above the 45° minimum required for non-oscillatory closed-loop behaviour. The transient response is well-damped, with small overshoot and a settling time consistent with the bandwidth of the buffer.

The choice of RC compensation is validated by the consistent phase margin and the absence of any visible degradation that would be associated with an unmitigated right-half-plane zero. The 4 dB lower nominal DC gain in 18 nm relative to 28 nm is consistent with the lower intrinsic device gain in the smaller node and is fully absorbed by the closed-loop feedback without compromising the buffer's accuracy or stability. The same architecture and compensation scheme transition cleanly between the two nodes, illustrating the portability of the design. The results therefore demonstrate that the proposed buffer is a viable building block for on-chip reference distribution in advanced-node mixed-signal SoCs, and that it can be readily ported between 28 nm and 18 nm technologies with minor adjustments to transistor dimensions and bias currents.

CHAPTER 6

CONCLUSION AND FUTURE SCOPE

6.1 SUMMARY OF THE WORK

This project report has presented the design, analysis, and simulation of a two-stage CMOS buffer amplifier suitable for buffering on-chip voltage references at 0.6 V, 1.2 V, and 1.3 V. The buffer has been implemented and characterised in both 28 nm and 18 nm CMOS technology nodes. The architecture comprises a PMOS-input differential first stage with an NMOS current-mirror load and an NMOS common-source second stage with a PMOS active load. The amplifier is closed in unity-gain feedback to operate as a buffer, and a series resistor–capacitor (RC) compensation network is connected between the high-impedance internal node and the output to ensure closed-loop stability.

A comprehensive review of buffer amplifier topologies, of single-stage and two-stage amplifier approaches, and of frequency compensation techniques in advanced CMOS design has been presented. The review motivated the choice of a two-stage architecture for its superior gain, output drive, and bandwidth–power trade-off relative to single-stage alternatives, and the choice of RC compensation for its simplicity, headroom efficiency, robustness, and area efficiency.

A detailed design methodology has been developed, including the choice of channel lengths and threshold-voltage variants, the sizing of the differential pair and the output stage, the selection of bias currents, and the determination of the compensation capacitance and nulling resistance. The methodology has been applied separately to the 28 nm and 18 nm implementations, and the corresponding transistor aspect ratios and operating points have been tabulated. The simulation results have confirmed that the design meets the targeted unity-gain bandwidth of 10 MHz at the typical corner with a phase margin greater than 60° across all three reference voltages and both technology nodes. Corner-based simulations have revealed that the DC gain spans 85 dB to 110 dB and the UGB spans 3 MHz to 40 MHz in 28 nm, and the DC

gain spans 82 dB to 103 dB and the UGB spans 5 MHz to 40 MHz in 18 nm, with the differences fully explained by the underlying device physics.

6.2 MAIN ACHIEVEMENTS

The principal achievements of this project report are the demonstration of a single buffer architecture that operates correctly across three different reference voltages and two different advanced CMOS nodes, the consolidation of frequency compensation techniques for two-stage CMOS amplifiers in a single comprehensive treatment, the experimental (simulation-based) validation that RC compensation is sufficient to achieve the targeted bandwidth and phase margin in 28 nm and 18 nm CMOS at supply voltages near 0.9 V, and a quantitative explanation of why the two nodes differ in their corner behaviour.

6.3 SIGNIFICANCE OF THE BUFFER BLOCK

On-chip voltage references are useless without buffers. Without a buffer, every load connected to the reference would draw current from the reference generator, modulating the reference voltage and degrading the accuracy of all downstream blocks. The buffer presents a low impedance to the load while drawing negligible current from the reference, and it allows the same reference to be distributed to many distant blocks across the chip without degrading its accuracy or settling time. A well-designed buffer is therefore a foundational analog block in any modern mixed-signal SoC.

6.4 ROLE OF FREQUENCY COMPENSATION IN THE DESIGN

A two-stage amplifier has two high-impedance nodes and therefore two close-in-frequency poles. In the absence of compensation, the open-loop phase falls rapidly through −90° toward −180° in the same frequency decade where the magnitude crosses unity, resulting in negligible or negative phase margin. The closed-loop response is then either heavily ringing or overtly oscillatory. Frequency compensation moves the two poles to drastically different frequencies (pole splitting), so that the open-loop transfer function behaves like a single-pole system across most of the

gain–bandwidth product. RC compensation, in particular, achieves pole splitting via a Miller capacitor and simultaneously eliminates the right-half-plane zero via a nulling resistor, giving a stable closed-loop response with adequate phase margin.

6.5 EFFECTIVENESS OF THE TWO-STAGE APPROACH

The two-stage approach has been effective in this work for three reasons. First, it provides enough open-loop gain to enforce a closed-loop voltage gain very close to unity, even in advanced CMOS nodes where the intrinsic device gain is low. Second, it decouples the gain stage from the output drive stage, allowing each to be optimised independently. Third, it admits a simple and robust compensation network in the form of a series RC between two existing nodes of the schematic, requiring no additional active devices and consuming no additional static power. These three properties make the two-stage architecture an excellent choice for compact, low-power, multi-reference buffers in 28 nm and 18 nm CMOS.

6.6 FUTURE SCOPE

The design and analysis reported in this project report open several directions for further investigation. The following are the most promising avenues for future work.

6.6.1 Power consumption reduction

Although the quiescent current of the proposed buffer is already moderate, there is room for further reduction. Operating the input differential pair in weak or moderate inversion can yield a higher gm/ID ratio, allowing the same unity-gain bandwidth at a lower bias current [21], [22]. Adaptive biasing schemes, in which the bias current is increased dynamically only during transient events, can also reduce the average power consumption substantially. Such schemes are particularly attractive in applications where the reference is mostly static and only occasionally subjected to transients.

6.6.2 Improved compensation techniques

While RC compensation is well-suited to the present application, more advanced techniques such as Ahuja compensation [6], current-buffer (indirect-feedback) compensation [15], or split-length compensation could be explored to widen the unity-gain bandwidth or to reduce the compensation capacitor size. These techniques would be particularly valuable in versions of the buffer that need to drive larger loads or operate at higher bandwidths, for example in high-speed data converter reference buffers [24].

6.6.3 Enhanced output swing and linearity

The output swing of the present buffer is limited at the upper rail by the saturation voltage of the PMOS active load, and at the lower rail by the saturation voltage of the NMOS driver. Rail-to-rail output stages, such as class-AB output stages with floating current sources, would extend the usable output swing, particularly at the lowest reference voltages. Linearity could be further improved by introducing local feedback within the output stage or by using a class-AB driver, which would also improve the slew-rate symmetry between rising and falling edges.

6.6.4 Adaptation to wider supply ranges

The present design is targeted at a supply voltage near 0.9 V. Extending the architecture to operate over a wider supply range, for example from 0.7 V to 1.8 V, would enable use in a broader range of applications. This would require careful biasing schemes that maintain the operating points of the input pair and the output stage across the entire supply range, possibly using cascoded current mirrors and constant-gm biasing.

6.6.5 Porting to newer technology nodes

The same design methodology can be applied to newer CMOS nodes such as 16 nm, 7 nm, and 5 nm FinFET technologies. The intrinsic gain of devices in these nodes is even lower than in 28 nm and 18 nm, requiring more aggressive use of long-channel devices, possibly three-stage architectures with nested compensation [7], [8], or novel low-voltage compensation strategies. Such a port would be a valuable validation of the general approach in even more aggressive nodes.

6.6.6 Integration into larger mixed-signal subsystems

Finally, the buffer can be integrated into larger mixed-signal subsystems such as bandgap references, low-dropout regulators, ADC reference distribution networks, and analog frontends for sensor interfaces. Each integration scenario will impose additional constraints, such as PSRR, noise, and load-step transient behaviour, and the design may need further refinement to meet the system-level requirements [3], [12], [24], [25]. The integration would also enable measurement of the buffer's performance in a realistic operating environment, validating the simulation-based design methodology presented here.

6.7 FINAL CONCLUSION

In conclusion, a compact, low-power, low-voltage, two-stage CMOS buffer amplifier with RC compensation has been designed, analysed, and characterised in 28 nm and 18 nm CMOS technologies. The buffer operates correctly at reference voltages of 0.6 V, 1.2 V, and 1.3 V, achieves a unity-gain bandwidth of 10 MHz at the typical corner, and maintains a phase margin greater than 60° at the typical corner in both nodes (dipping only marginally to 57–58° at the worst process corner of each node, still well above the 45° unconditional-stability minimum). The buffer exhibits a corner-induced DC gain spread of 85–110 dB in 28 nm and 82–103 dB in 18 nm, and consumes a moderate quiescent current. Importantly, although both nodes use the same topology and compensation network, they exhibit qualitatively different corner sensitivities: gain and phase margin are maximised at FF and minimised at SS in 28 nm, but maximised at FS and minimised at FF in 18 nm. This divergence is fully explained by the different way short-channel effects, drain current, and load output resistance interact at the two technology nodes. The design is portable between the two technology nodes with minor modifications and represents a viable building block for reference distribution in modern mixed-signal SoCs. The methodology, the rationale behind every design choice, the simulation-based validation, and the directions identified for future work in Section 6.6 are intended to serve as a reusable foundation for further research on low-voltage analog buffers in advanced CMOS nodes.

REFERENCES

[1] P. R. Gray, P. J. Hurst, S. H. Lewis, and R. G. Meyer, Analysis and Design of Analog Integrated Circuits, 5th ed. Hoboken, NJ: John Wiley & Sons, 2009.

[2] B. Razavi, Design of Analog CMOS Integrated Circuits, 2nd ed. New York, NY: McGraw-Hill, 2017.

[3] P. E. Allen and D. R. Holberg, CMOS Analog Circuit Design, 3rd ed. New York, NY: Oxford University Press, 2012.

[4] W. M. C. Sansen, Analog Design Essentials. Dordrecht, The Netherlands: Springer, 2006.

[5] D. A. Johns and K. Martin, Analog Integrated Circuit Design, 2nd ed. Hoboken, NJ: John Wiley & Sons, 2011.

[6] B. K. Ahuja, "An improved frequency compensation technique for CMOS operational amplifiers," IEEE J. Solid-State Circuits, vol. SC-18, no. 6, pp. 629–633, Dec. 1983.

[7] R. G. H. Eschauzier and J. H. Huijsing, Frequency Compensation Techniques for Low-Power Operational Amplifiers. Boston, MA: Kluwer Academic Publishers, 1995.

[8] K. N. Leung and P. K. T. Mok, "Analysis of multistage amplifier-frequency compensation," IEEE Trans. Circuits Syst. I, vol. 48, no. 9, pp. 1041–1056, Sep. 2001.

[9] X. Peng and W. Sansen, "AC boosting compensation scheme for low-power multistage amplifiers," IEEE J. Solid-State Circuits, vol. 39, no. 11, pp. 2074–2079, Nov. 2004.

[10] H. Lee and P. K. T. Mok, "Active-feedback frequency-compensation technique for low-power multistage amplifiers," IEEE J. Solid-State Circuits, vol. 38, no. 3, pp. 511–520, Mar. 2003.

[11] R. J. Baker, CMOS: Circuit Design, Layout, and Simulation, 4th ed. Hoboken, NJ: IEEE Press / Wiley, 2019.

[12] T. C. Carusone, D. A. Johns, and K. W. Martin, Analog Integrated Circuit Design, 2nd ed. Hoboken, NJ: John Wiley & Sons, 2012.

[13] J. M. Rabaey, A. Chandrakasan, and B. Nikolić, Digital Integrated Circuits: A Design Perspective, 2nd ed. Upper Saddle River, NJ: Prentice Hall, 2003.

[14] A. D. Grasso, S. Pennisi, and G. Palumbo, "Three-stage CMOS OTA for large capacitive loads with efficient frequency compensation scheme," IEEE Trans. Circuits Syst. II, vol. 53, no. 10, pp. 1044–1048, Oct. 2006.

[15] V. Saxena and R. J. Baker, "Indirect feedback compensation of CMOS op-amps," in Proc. IEEE Workshop Microelectronics and Electron Devices (WMED), Apr. 2006, pp. 3–4.

[16] G. Palumbo and S. Pennisi, Feedback Amplifiers: Theory and Design. Boston, MA: Kluwer Academic Publishers, 2002.

[17] Y. Tsividis and C. McAndrew, Operation and Modeling of the MOS Transistor, 3rd ed. New York, NY: Oxford University Press, 2011.

[18] R. J. Widlar, "Some circuit design techniques for linear integrated circuits," IEEE Trans. Circuit Theory, vol. 12, no. 4, pp. 586–590, Dec. 1965.

[19] J. K. Roberge, Operational Amplifiers: Theory and Practice. New York, NY: John Wiley & Sons, 1975.

[20] M. J. M. Pelgrom, A. C. J. Duinmaijer, and A. P. G. Welbers, "Matching properties of MOS transistors," IEEE J. Solid-State Circuits, vol. 24, no. 5, pp. 1433–1439, Oct. 1989.

[21] E. Vittoz and J. Fellrath, "CMOS analog integrated circuits based on weak inversion operations," IEEE J. Solid-State Circuits, vol. SC-12, no. 3, pp. 224–231, Jun. 1977.

[22] W. Sansen, "Biasing for low power and low voltage in nanometer CMOS," in Analog Design Essentials, Dordrecht, The Netherlands: Springer, 2015.

[23] S. Rakers, L. Connell, T. Collins, and D. Russell, "Secure contactless smartcard ASIC with DPA protection," IEEE J. Solid-State Circuits, vol. 36, no. 3, pp. 559–565, Mar. 2001.

[24] K. Bult and G. J. G. M. Geelen, "A fast-settling CMOS op amp for SC circuits with 90-dB DC gain," IEEE J. Solid-State Circuits, vol. 25, no. 6, pp. 1379–1384, Dec. 1990.

[25] A. Pirola, A. Liscidini, and R. Castello, "Current-mode, WCDMA channel filter with in-band noise shaping," IEEE J. Solid-State Circuits, vol. 45, no. 9, pp. 1770–1780, Sep. 2010.