Politecnico di Torino

Politecnico di Torino

Master of Science in Electronic Engineering
A.a. 2025/2026
Graduation Session July 2026

Design Validation of an Active Neural Detector

Supervisors:

Paolo Motto Ros
Paolo Stefano Crovetti
Danilo Demarchi

Candidate:

Liang Wang

Abstract

This thesis focuses on the design validation of an Active Neural Detector integrated circuit intended for miniaturized neural interfaces. The circuit is based on an existing architecture developed within the CEREBRO project and combines a two-stage operational transconductance amplifier, a current-starved voltage-controlled oscillator and an inverter-based output buffer.

The purpose of this work is to verify the behavior of the reference design through circuit-level, post-layout and system-level simulations. The OTA is evaluated in terms of gain, bandwidth, noise and offset-rejection capability. The VCO is analyzed for its voltage-to-frequency conversion behavior, frequency sensitivity and process-corner dependence. The output buffer is considered as the final driving stage of the signal chain.

Additional sinusoidal-input simulations are included to validate the complete Active Neural Detector operation. These simulations show how low-frequency neural-voltage variations are translated into frequency modulation at the oscillator output. The results confirm that the system operates correctly within the intended neural-signal band, while the modulation depth is limited at very low and high input frequencies by the offset-compensation path and the OTA bandwidth.

Overall, the thesis validates the functionality of the reference Active Neural Detector and discusses its main performance limitations in terms of bandwidth, noise, process variation and frequency-conversion behavior.

Summary

This thesis validates an ultra-low-power implantable neural front-end in TSMC 180 nm CMOS technology. The circuit combines OTA amplification, VCO-based frequency modulation and output buffering to support weak neural-signal detection and tissue transmission.

Neural signals are at the tens-of-microvolt level, but they are usually superimposed on electrode DC offsets of several tens of millivolts. Direct high-gain amplification can therefore lead to saturation, while flicker noise and strict power and area limits make low-frequency acquisition difficult. The reference detector uses a two-stage OTA with approximately 49 dB gain and band-pass behavior in the neural band, followed by a current-starved VCO and a four-stage output buffer.

System simulations show that a 30 microvolt, 30 Hz neural component superimposed on a 30 mV DC offset can be converted into megahertz-range frequency variations. The work therefore validates the feasibility of translating weak low-frequency neural information into a high-frequency timing signal, which can be more suitable for transmission through tissue and for compact implantable sensing interfaces.

Acknowledgements

I would like to express my gratitude to Prof Paolo Motto Ros,Paolo Stefano Crovetti and to the research group for their guidance on low-power integrated-circuit design, neural-interface modelling and post-layout verification. I also thank my classmates for useful discussions on tissue propagation, simulation setup and measurement constraints. Finally, I thank my family and friends for their continuous support during the preparation of this thesis.

Table of Contents

Abstract I

Summary II

Acknowledgements III

List of Figures VII

List of Tables VIII

1 Introduction 1
- 1.1 Motivation . . . 1
- 1.2 Thesis objective and scope . . . 2
- 1.3 Why frequency translation is useful . . . 2
- 1.4 Structure of the thesis . . . 3
- 1.5 Thesis contribution . . . 3
- 1.6 Design questions addressed in the work . . . 4
- 1.7 Acronyms used in the thesis . . . 5

2 Background on Neural Recording and Implantable Front-Ends 6
- 2.1 Neural signals and electrode interface . . . 6
- 2.2 Noise constraints in neural front-ends . . . 6
- 2.3 Offset rejection and band-pass acquisition . . . 7
- 2.4 Implantable-device constraints . . . 7
- 2.5 Relevant integrated-circuit approaches . . . 8
- 2.6 Implications for this thesis . . . 8
- 2.7 Signal characteristics relevant to the detector . . . 9
- 2.8 Noise as a system-level constraint . . . 9
- 2.9 Low-power CMOS constraints . . . 10
- 2.10 Frequency encoding and robustness . . . 10

3 Reference Active Neural Detector Architecture and Specifications 12
- 3.1 End-to-end signal chain . . . 12

3.2 Design specifications . . . 14
3.3 Architecture-level trade-offs . . . 15
3.4 Frequency conversion principle . . . 16
3.5 Simulation and verification hierarchy . . . 16
3.6 Requirement allocation across the signal chain . . . 17
3.7 Electrode and load abstraction . . . 17
3.8 Architecture-level performance metrics . . . 18

4 OTA Design, Sizing and Layout Validation 19
4.1 Role of the OTA . . . 19
4.2 gm/ID sizing methodology . . . 19
4.3 Transistor-level sizing . . . 20
4.4 Open-loop first-stage behavior . . . 21
4.5 Buffer-mode OTA and band-pass first-stage response . . . 22
4.6 Noise behavior . . . 22
4.7 Layout and parasitic considerations . . . 23
4.8 Detailed implementation sequence . . . 24
4.9 Cadence verification flow . . . 25
4.10 OTA current and power consumption . . . 26
4.11 Input-pair sizing strategy . . . 27
4.12 Compensation and stability interpretation . . . 27
4.13 Noise calculation perspective . . . 27
4.14 Layout considerations for the OTA . . . 28

5 VCO and Buffer Design Validation 29
5.1 Current-starved VCO principle . . . 29
5.2 VCO sensitivity . . . 30
5.3 Phase noise and detectability . . . 30
5.4 Temperature dependence . . . 31
5.5 Buffer design . . . 31
5.6 Power consumption . . . 32
5.7 VCO design procedure and calibration interpretation . . . 33
5.8 From phase noise to neural-event resolution . . . 34
5.9 Design implications . . . 35
5.10 Current-starved oscillator design reasoning . . . 35
5.11 Phase noise and neural information . . . 35
5.12 Buffer sizing and load isolation . . . 36
5.13 Temperature and process sensitivity . . . 36

6 Post-Layout Simulation and System-Level Validation 38
6.1 Post-layout validation purpose . . . 38
6.2 OTA and VCO validation summary . . . 38
6.3 Complete system simulation . . . 38
6.4 Monte Carlo validation . . . 39
6.5 Quantitative performance summary . . . 40
6.6 Layout-driven interpretation . . . 41
6.7 Conclusion . . . 42

7 Additional Functional Validation under Sinusoidal Neural Inputs 43
7.1 Objective of the additional validation . . . 43
7.2 Simulation setup . . . 43
7.3 VCO operating point under sinusoidal excitation . . . 44
7.4 Representative transient responses . . . 45
7.4.1 Low-frequency input conditions . . . 46
7.4.2 Neural-band input conditions . . . 48
7.4.3 High-frequency input conditions . . . 49
7.5 Discussion . . . 52
7.6 Chapter summary . . . 53

A Supplementary Design Notes 54
A.1 Approximate neural-to-frequency conversion calculation . . . 54
A.2 Dynamic power of the output buffer . . . 54
A.3 Notes on figures . . . 55

Bibliography 56

List of Figures

1.1 Overall System Architecture of the Active Neural Detector. 3

3.1 System architecture of the active neural detector. 12
3.2 Transistor-level schematic of the two-stage OTA. 12
3.3 Transconductance feedback model of the OTA input stage. 13
3.4 Transistor-level schematic of the current-starved VCO. 13
3.5 Four-stage inverter output buffer. 14
3.6 Simulation flow of the active neural detector. 17

4.1 gm/ID sizing trade-off for the OTA input stage. 21
4.2 First-stage OTA frequency response. 22
4.3 OTA output and input noise simulation. 23
4.4 Layout of the active neural detector. 24

5.1 Transistor-level schematic of the current-starved VCO. 29
5.2 VCO simulation results. 31
5.3 Four-stage tapered inverter buffer. 32

6.1 System-level transient simulation results. 39
6.2 Monte Carlo simulation results. 40

7.1 VCO tuning characteristic under a 30 Hz neural input. 45
7.2 Active neural detector response for a 30 Hz sinusoidal input. 46
7.3 Active neural detector response for a 1 Hz sinusoidal input. 47
7.4 Active neural detector response for a 10 Hz sinusoidal input. 48
7.5 Active neural detector response for a 100 Hz sinusoidal input. 49
7.6 Active neural detector response for a 1 kHz sinusoidal input. 50
7.7 Active neural detector response for a 10 kHz sinusoidal input. 51
7.8 Active neural detector response for a 100 kHz sinusoidal input. 52

List of Tables

1.1 Main acronyms and symbols used throughout the thesis. 5

3.1 Main design specifications used for the active neural detector. . . . 15

4.1 Design of the active neural detector. N is the number of fingers. . . 20
4.2 OTA verification checklist used to organize the schematic and layout simulations. 26

5.1 Interpretation of the VCO simulation results for detector-level use. . 34

6.1 Summary of simulated performance values. 41

7.1 Sinusoidal-input validation conditions. 44
7.2 Qualitative summary of the sinusoidal-input validation. 53

Chapter 1

Introduction

1.1 Motivation

Miniaturized neural interfaces are increasingly important in neuroscience, neuroprosthetics and implantable medical devices because they make it possible to acquire neural activity closer to the tissue region of interest. A recording interface placed near the signal source can reduce the influence of cable parasitics and external interference, but the reduction in distance introduces strict limits on silicon area, heat generation and available energy. This thesis studies an active neural detector related to the CEREBRO implantable-system concept [1]. The detector is based on three compact analog and mixed-signal blocks: an input OTA, a VCO, and a tapered buffer that can drive a capacitive or tissue-like load.

The need for this architecture comes from a fundamental conflict in neural acquisition. Extracellular neural signals are small, often in the microvolt to millivolt range, whereas the electrode–electrolyte interface can introduce DC offsets that are orders of magnitude larger. A conventional high-gain voltage amplifier cannot simply amplify the entire electrode signal, because the DC offset would saturate the front-end long before the useful neural information is recovered. At the same time, the input-referred noise of the amplifier must remain low enough not to mask the low-frequency neural component. The adopted front-end therefore combines high mid-band gain, low-frequency offset compensation and frequency translation. The frequency-translation step is especially useful because the neural information is moved from a low-frequency voltage signal to a high-frequency oscillation, which can be more convenient for propagation or remote sensing in a miniaturized implantable platform.

Neural-recording amplifiers have been widely studied, from classical capacitively coupled front-ends to low-power DC-coupled and chopper-stabilized implementations [2, 4–6]. The circuit considered in this thesis belongs to the family of ultra-low-power neural front-ends, but it differs from a conventional recording chain because its output is not a digitized amplitude sample. Instead, the detector behaves as an active transducer: the OTA produces a voltage proportional to the neural signal after offset rejection, and the VCO converts this voltage into an oscillation-frequency variation. The thesis develops the circuit-level reasoning,

design-flow details, layout considerations and broader literature context needed to present the detector as a complete integrated-circuit design.

1.2 Thesis objective and scope

The objective is for design validation of the active neural detector as a complete integrated-circuit study. The work is not limited to reporting isolated simulation curves. It reconstructs the engineering logic that connects the system requirement to transistor-level choices and then to post-layout behavior. The specific goals are as follows. First, the detector architecture is described as an end-to-end signal chain, from a microelectrode input with DC offset to a high-frequency output signal. Second, the OTA is analyzed as the main noise- and gain-defining block. Third, the current-starved VCO is discussed as a sensitivity-dominant conversion block. Fourth, the output buffer and layout floorplan are included because they are essential for a practical implantable device. Finally, the post-layout simulation results are organized into independent, clearly labelled figures, improving readability and supporting quantitative comparison.

The design target is defined as follows: a 180 nm CMOS implementation, a supply of $\pm 0.9\,\mathrm{V}$, an OTA bias current around 50 nA, a first-stage gain close to 50 dB, an input-noise target below approximately 10 $\mu\mathrm{V}/\sqrt{\mathrm{Hz}}$ at 10 Hz, and an occupied area below $200\,\mu\mathrm{m} \times 200\,\mu\mathrm{m}$. These figures are deliberately demanding because the intended use is a miniaturized neural implant. Power consumption must remain below the microwatt level for the active sensing part, not only to satisfy wireless-power limitations but also to reduce local heating and tissue-safety concerns.

1.3 Why frequency translation is useful

In a classical amplitude-recording front-end, the signal must remain above the noise floor while travelling through the following readout path. This is difficult in very small implants because the available output swing and drive strength are limited. The active detector studied here translates the neural amplitude into a frequency deviation. Frequency is often more robust than amplitude to attenuation and gain uncertainty: if the receiving system can detect zero crossings or spectral peaks, the information can be recovered even when the waveform amplitude changes along the propagation path. The price to pay is that the oscillator must have sufficient sensitivity, low phase noise around relevant offset frequencies, and stable behavior under process and temperature variations.

This conversion can be interpreted as a form of analog preprocessing. The weak

neural voltage is first amplified in a controlled band and then used to modulate an oscillator. The VCO does not need to generate a precision communication carrier; rather, it must produce a measurable frequency shift that is larger than the uncertainty caused by phase noise, supply fluctuation and circuit mismatch. This design philosophy explains why the thesis pays attention both to the OTA mid-band gain and to the VCO sensitivity K_{VCO}.

1.4 Structure of the thesis

Chapter 2 introduces the biological and circuit background, including electrode offsets, noise mechanisms and the role of wireless power constraints. Chapter 3 develops the active detector architecture and summarizes the design specifications. Chapter 4 focuses on the two-stage OTA, its gm/ID design flow, its transistor-level sizing and its layout. Chapter 5 discusses the current-starved VCO and the tapered buffer. Chapter 6 presents the post-layout simulation results and the complete system-level validation. The appendix reports additional calculations and supporting design tables.

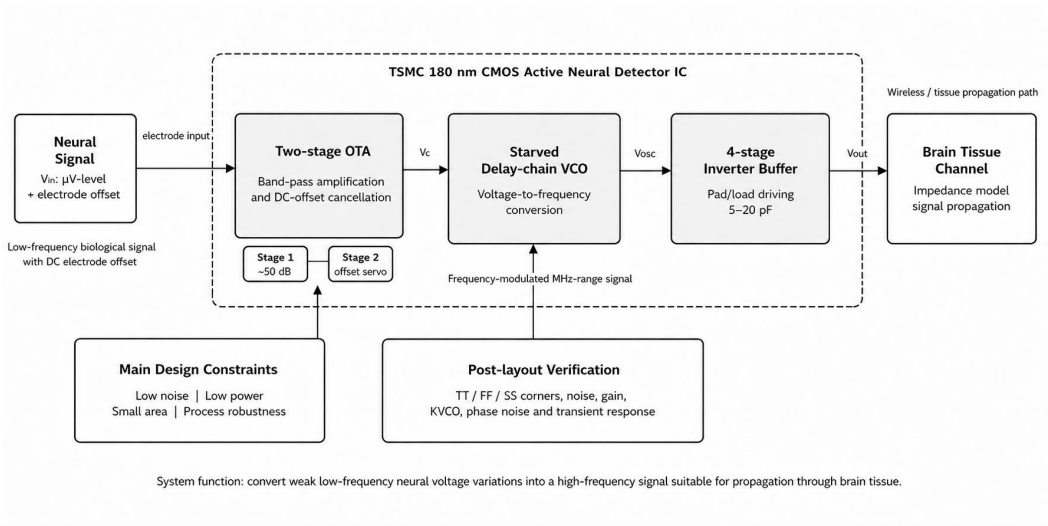

Figure 1.1: Overall System Architecture of the Active Neural Detector.

1.5 Thesis contribution

The main contribution of this thesis is to present the active detector as a coherent design study, linking the system requirement, the transistor-level implementation, the layout constraints and the final simulations. Architecture diagrams, floorplan diagrams and simulation plots are organized in a consistent engineering style, while the discussion emphasizes the design rationale that connects schematic

development with verification. The bibliography is expanded to include neural front-end circuits, analog CMOS design methodology, oscillator phase-noise theory and tissue-electromagnetic references, giving the detector design a broader academic context.

1.6 Design questions addressed in the work

The central design question is not only whether the circuit can amplify a neural signal, but whether the complete detector can preserve neural information after it has been translated into a different physical quantity. This point is important because the input neural signal is a voltage measured at an electrode, while the final observable variable is an oscillation frequency. Therefore the design must satisfy three linked conditions. The first condition is sensitivity: a change of a few tens of microvolts at the electrode must produce a frequency deviation that can be distinguished from noise and drift. The second condition is bounded offset response: a DC electrode offset of several tens of millivolts must not consume the signal swing or force the OTA into saturation. The third condition is practical integrability: the active area, bias current and output loading must remain compatible with an implantable device rather than a bench-top recording channel.

A useful way to read the thesis is to follow these three conditions through the chapters. In the OTA chapter, sensitivity appears as mid-band transconductance gain and input-referred noise. In the VCO chapter, it appears as K_{VCO} and as the linearity of the control curve. In the layout and simulation chapters, it appears as the ability of the extracted circuit to maintain these quantities after parasitic capacitances and resistances have been introduced. The same logic applies to offset handling. Offset is first a property of the electrode interface, then a constraint on the OTA input range, and finally a system-level transient condition that must be rejected before frequency encoding becomes meaningful. This thesis therefore presents the design as a chain of constraints rather than as a disconnected set of blocks.

1.7 Acronyms used in the thesis

Table 1.1: Main acronyms and symbols used throughout the thesis.

Acronym or symbol	Meaning
OTA	Operational transconductance amplifier
VCO	Voltage-controlled oscillator
CMOS	Complementary metal–oxide–semiconductor technology
LFP	Local field potential
AP	Action potential
K_{VCO}	VCO voltage-to-frequency sensitivity
g_m/I_D	Transconductance efficiency used for analog sizing
PVT	Process, voltage and temperature variation

Chapter 2

Background on Neural Recording and Implantable Front-Ends

2.1 Neural signals and electrode interface

Extracellular neural recording relies on the small voltage variations generated by ionic currents near active neurons. Depending on the electrode size, distance from the neuron, and tissue environment, the useful signal can range from local field potentials at low frequencies to spike-like activity at higher frequencies. In this design, the target neural component is represented by a low-amplitude signal of approximately 30 µV at tens of hertz, superimposed on a much larger DC offset. This simplified input is not meant to reproduce every biological waveform; it is a compact stress test for the analog front-end. If the detector can recover a microvolt-scale sinusoidal component in the presence of a tens-of-millivolts offset, the essential offset-rejection and sensitivity requirements are captured.

The electrode–tissue interface is one of the main reasons why neural front-ends differ from generic low-noise amplifiers. The electrode can be modeled by a half-cell potential, a double-layer capacitance, a charge-transfer resistance and a spreading impedance. These elements vary with material, geometry, implantation time and tissue condition. The resulting DC offset may be much larger than the neural signal, and it may drift slowly over time. A purely open-loop high-gain amplifier would therefore amplify both neural activity and offset, causing saturation. This motivates architectures that provide high gain only inside a useful band while suppressing low-frequency components.

2.2 Noise constraints in neural front-ends

Noise is a dominant design constraint for integrated neural amplifiers [2, 3, 7]. Thermal noise is related to bias current and transconductance, whereas flicker noise depends strongly on device area and technology parameters. Low-frequency operation makes flicker noise especially important in this detector. A simplified expression for the flicker-noise voltage spectral density is

$$S_V(f) = \frac{K}{C_{ox} W L f}, \tag{2.1}$$

where K is the process-dependent flicker-noise coefficient, C_{ox} is the gate-oxide capacitance per unit area, and W and L are the channel width and length. Equation (2.1) explains why the input pair of the OTA uses relatively large dimensions. Increasing the gate area reduces flicker noise, but it increases input capacitance and layout area. This is a typical analog trade-off: reducing noise by increasing area can damage compactness and bandwidth.

The detector also has to satisfy a power constraint. For a fixed transconductance, increasing bias current usually reduces thermal noise and increases bandwidth, but it increases power. For a fixed current, operating devices in a moderate or weak inversion region can improve transconductance efficiency. The gm/ID methodology is therefore suitable for this design because it provides a systematic relation between current, inversion level, transconductance, intrinsic gain and device geometry [9, 10]. The design must not only meet a nominal noise value but also maintain acceptable behavior under process corners.

2.3 Offset rejection and band-pass acquisition

A useful neural front-end must reject or compensate slow offset while preserving the target signal band. The detector achieves this by combining the first OTA stage with a low-frequency feedback path through the second stage operating as a buffer. The first stage is responsible for the high mid-band gain. The closed-loop behavior of the full OTA suppresses very low-frequency offset and avoids saturation. This produces a band-pass response at the first-stage output: the lower cut-off is determined by the DC-compensation loop, whereas the upper cut-off is determined by the open-loop gain and compensation capacitances.

This architecture is particularly relevant for miniaturized implants because it avoids adding large passive components. In many biopotential amplifiers, a very low high-pass corner would require a large resistor or a pseudo-resistor. The pseudo-resistor can realize a large effective resistance in a small area, but it is nonlinear and process dependent. The implementation uses pseudo-resistor devices M12 and M13 between the OTA stages to achieve low-frequency compensation without consuming large silicon area.

2.4 Implantable-device constraints

Implantable systems are constrained by size, power delivery, thermal safety and tissue propagation. Wireless power transfer is often considered for miniaturized implants because replacing or wiring a battery is impractical. However, when the receiver becomes smaller, the available power decreases and the transfer efficiency

becomes sensitive to depth, alignment and tissue properties. The system context includes several wireless power transfer approaches, such as acoustic, capacitive, inductive and magnetoelectric methods. The detector itself is only one part of the system, but its sub-microwatt target is a direct consequence of the power-delivery challenge.

The signal emitted by the detector also interacts with tissue impedance. Tissue is dispersive and lossy, and its dielectric properties vary with frequency [18]. Frequency translation can help because the information is carried by timing rather than by absolute amplitude. Still, the buffer must drive a capacitive and conductive load without consuming excessive power. This is why the design uses a tapered inverter buffer: small inverters reduce area and power at the VCO output, while progressively larger devices increase drive strength near the load.

2.5 Relevant integrated-circuit approaches

The literature on integrated neural interfaces includes low-noise analog front-ends, multiplexed recording arrays, wireless microimplants and frequency-based sensing approaches. Harrison and Charles demonstrated a low-power low-noise CMOS amplifier for neural recording, establishing many of the design trade-offs that remain relevant today [2]. Wattanapanitch and collaborators developed energy-efficient micropower amplifiers, showing the importance of noise efficiency in implantable channels [3]. More recent systems integrate amplifiers, digitizers and wireless interfaces on compact chips [5, 6]. The detector studied here should be read in that context: it is not a complete neural-recording system but a compact active transducer that can be embedded in a larger implantable architecture.

2.6 Implications for this thesis

The background discussion leads to four design implications. First, the input pair of the OTA is the most important noise contributor and must use large enough devices. Second, offset compensation must be integrated without large external passives. Third, the VCO must provide a frequency shift that is clearly larger than jitter and phase-noise uncertainty. Fourth, layout parasitics and process variations must be included because the target power is very small. These implications guide the architecture described in the following chapter.

For this reason, the following chapters keep the detector topology consistent while focusing on a complete engineering explanation. The emphasis is on how the circuit blocks satisfy the noise, offset, power and frequency-conversion constraints of the active neural detector.

2.7 Signal characteristics relevant to the detector

Extracellular neural signals are usually grouped into low-frequency local field potentials and higher-frequency spike activity. The exact bandwidth depends on the electrode geometry, the tissue region and the recording method, but the design problem always contains the same difficulty: the useful signal is small, while the electrode interface can generate a relatively large slowly varying offset. For an active detector, this means that the useful variable is not the absolute electrode voltage but the small variation around a large baseline. A front-end that amplifies the baseline together with the signal would quickly lose headroom, especially when the supply is only around 1.8 V end-to-end.

The detector addresses this difficulty by separating the roles of low-frequency stabilization and signal-band gain. The OTA must provide enough gain in the useful band, but it must not treat the electrode offset as a signal that deserves the same amplification. In practice, this means that the frequency response of the front-end should include an effective high-pass or offset-cancellation behavior at very low frequency, while maintaining large gain in the band where neural fluctuations are expected. The implementation uses a buffer-style feedback path for this purpose. From a system viewpoint, this is equivalent to assigning the very slow component to a compensation loop and the neural component to the VCO modulation path.

2.8 Noise as a system-level constraint

Noise cannot be treated as a final verification step only. In a weak-signal neural detector, noise affects the architecture from the beginning. The input devices of the OTA contribute thermal noise, flicker noise and bias-current-related mismatch. The VCO adds phase noise and converts voltage noise into timing uncertainty. The buffer adds edge uncertainty and load-dependent supply disturbance. If the OTA noise is too high, the oscillator may faithfully translate noise rather than neural information. If the oscillator phase noise is too high, an otherwise clear neural modulation becomes difficult to recover from the frequency-domain output.

A compact way to express this coupling is to write the small-signal frequency deviation as

$$\Delta f(t) = K_{VCO} A_{OTA}(s) v_n(t), \tag{2.2}$$

where $v_n(t)$ is the electrode neural component, $A_{OTA}(s)$ is the OTA transfer function in the signal band and K_{VCO} is the oscillator sensitivity. The same expression also applies to input-referred noise if $v_n(t)$ is replaced by the equivalent noise source.

This equation explains why the thesis discusses gain, input-referred noise and oscillator sensitivity together: the complete signal-to-noise ratio at the output is determined by the product of the analog gain and the frequency-conversion gain, not by either block alone.

2.9 Low-power CMOS constraints

Implantable and miniaturized detectors cannot rely on high bias currents to solve every analog problem. A higher current would reduce thermal noise and increase bandwidth, but it would also increase power density and simplify neither wireless powering nor thermal management. The 180 nm CMOS node used in this implementation is therefore a reasonable compromise. It offers enough analog headroom and device matching for low-frequency neural acquisition while remaining compatible with compact mixed-signal blocks and mature layout rules. The goal is not to maximize speed; the goal is to obtain adequate noise, gain and robustness at very low current.

The gm/ID methodology is especially suitable under these conditions. Instead of starting from transistor dimensions only, the designer chooses an inversion level that balances transconductance efficiency, intrinsic gain and parasitic capacitance. Weak inversion provides high transconductance efficiency, but it may suffer from bandwidth and mismatch limitations. Strong inversion provides speed and matching but costs current. Moderate inversion is often the practical choice for low-power neural amplifiers, because it allows the input pair to maintain useful transconductance without excessively increasing the device area or bias current.

2.10 Frequency encoding and robustness

The use of a VCO can be interpreted as a local analog-to-time conversion. Rather than preserving the neural signal as a voltage amplitude, the circuit preserves it as a change in oscillation frequency. This is attractive when the following medium or readout channel attenuates amplitude but still permits timing or spectral information to be observed. In such a system, the receiver does not need to know the exact gain of the propagation path; it needs to estimate the time-varying carrier frequency. This is the same general reason why frequency modulation is robust in many communication systems, although the present detector is not a conventional radio transmitter.

However, frequency encoding creates its own requirements. The oscillator must start reliably, remain monotonic over the control range and have phase noise low enough that small frequency deviations are not hidden. The buffer must drive the load without pulling the oscillator frequency. The supply network and substrate must be laid out carefully because oscillator switching can couple back into the low-noise OTA. These issues motivate the separate treatment of the VCO and buffer in Chapter 5 and the system-level verification in Chapter 6.

Chapter 3

Reference Active Neural Detector Architecture and Specifications

3.1 End-to-end signal chain

The active neural detector contains three main blocks: an OTA, a VCO and a buffer. The OTA receives the neural input and the electrode offset. Its first stage provides high gain for the useful signal band, while the second-stage feedback configuration compensates low-frequency offset. The VCO receives the OTA output as a control voltage and converts voltage variations into frequency variations. The buffer then drives the external capacitive or tissue-equivalent load. Figures 3.1–3.5 present the architecture and the main circuit blocks using separate vector graphics.

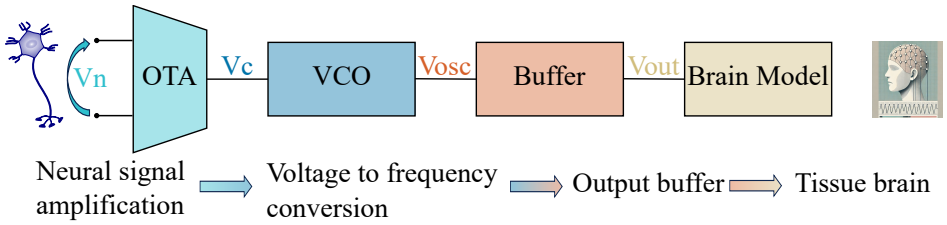

Figure 3.1: System architecture of the active neural detector.

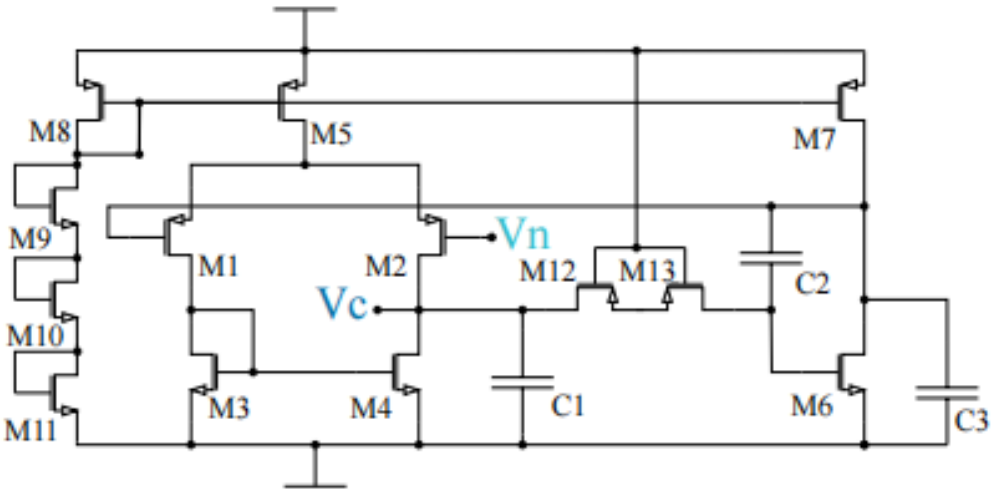

Figure 3.2: Transistor-level schematic of the two-stage OTA.

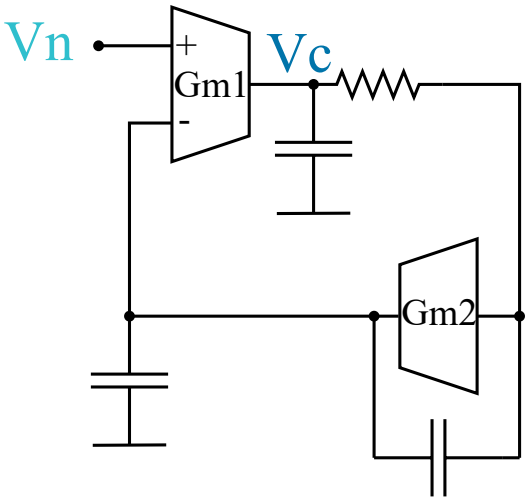

Figure 3.3: Transconductance feedback model of the OTA input stage.

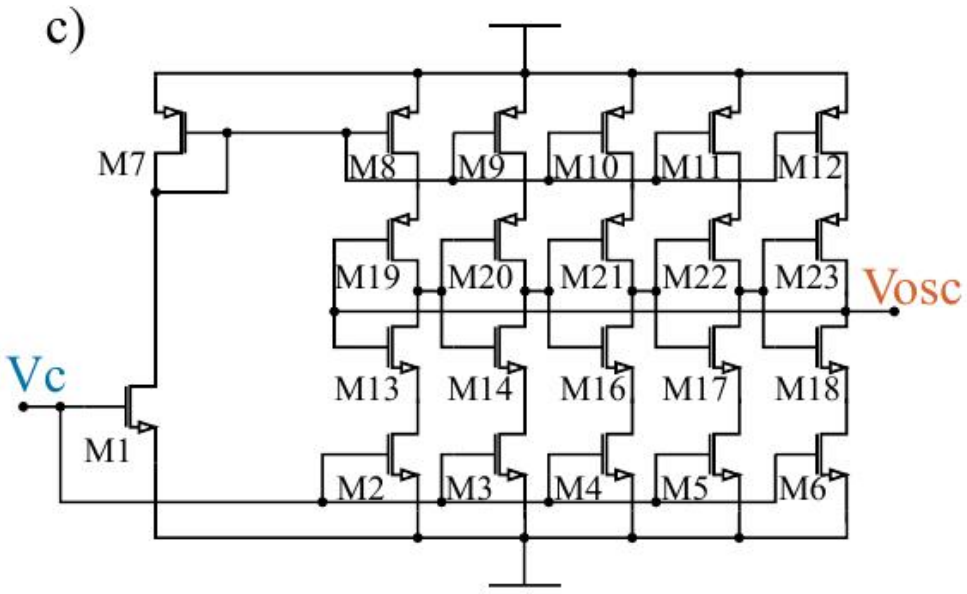

Figure 3.4: Transistor-level schematic of the current-starved VCO.

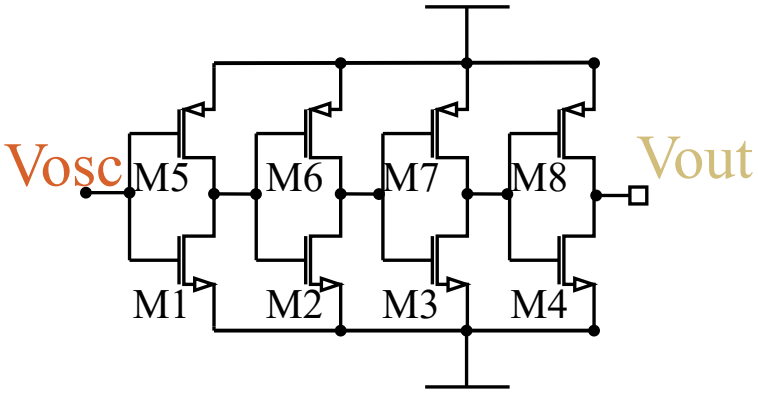

Figure 3.5: Four-stage inverter output buffer.

The signal path can be summarized as

$$V_n + V_{off} \longrightarrow V_{OTA} \longrightarrow f_{osc}(t) \longrightarrow V_{buf}(t), \tag{3.1}$$

where V_n is the neural component, V_{off} is the electrode offset, V_{OTA} is the control voltage provided by the amplifier, and $f_{osc}(t)$ is the instantaneous oscillation frequency. The central idea is that the information originally contained in a low-frequency voltage amplitude is encoded in the instantaneous frequency of a high-frequency output waveform.

3.2 Design specifications

The key specifications are defined for this design. Table 3.1 summarizes the targets and their system-level interpretation. The first-stage gain target is approximately 50 dB, corresponding to an amplitude multiplication of about 316. A neural input of 30 µV is therefore converted to a few millivolts at the first-stage output, which is large enough to modulate the VCO. The OTA noise target is below 10 $\mu V/\sqrt{Hz}$ at 10 Hz. The silicon-area target is below 200 µm×200 µm, while the final layout is about 162 µm×86 µm.

Table 3.1: Main design specifications used for the active neural detector.

Item	Meaning	Target value
Technology	CMOS process used for transistor-level implementation	180 nm
Supply	Low-voltage analog and mixed-signal operation	± 0.9 V
OTA bias current	Nominal current budget of the analog front-end	50 nA
First-stage gain	Amplification of the neural component before VCO control	50 dB target, 49 dB simulated
Input noise	Low-frequency input-referred noise requirement	$< 10\ \mu\mathrm{V}/\sqrt{\mathrm{Hz}}$ at 10 Hz
Area	Compact implantable active detector	$< 200\ \mu\mathrm{m} \times 200\ \mu\mathrm{m}$
VCO control range	Control voltage range used in transient simulations	150 mV–300 mV
Buffer load	Capacitive load associated with pads and measurement/tissue interface	5 pF–20 pF

3.3 Architecture-level trade-offs

The architecture balances several trade-offs. A larger OTA input pair reduces flicker noise but increases capacitance and area. Higher bias current improves speed and noise but violates the power target. A higher VCO sensitivity improves detectability of small neural signals but can make the oscillator more sensitive to supply and device noise. A stronger buffer improves load driving but increases dynamic power and switching disturbance. The transistor-level implementation selects transistor dimensions that keep the active detector below the area and power limits while still producing a measurable frequency shift.

A crucial architectural choice is to avoid full analog-to-digital conversion inside the implant. Instead of producing digital samples locally, the circuit produces a frequency-coded waveform. This reduces functional complexity and can be

compatible with a small implant, but it moves part of the burden to the receiver, which must estimate frequency over time. The detector therefore acts as a minimal active sensing node rather than a complete data-acquisition channel.

3.4 Frequency conversion principle

The current-starved VCO is governed by the approximate relation

$$f_{osc} = \frac{I_{ctrl}}{2CV_{pp}}, \tag{3.2}$$

where I_{ctrl} is the current controlling the inverter delay, C is the effective switching capacitance, and V_{pp} is the voltage swing. The sensitivity is

$$K_{VCO} = \frac{\partial f_{osc}}{\partial V_c}. \tag{3.3}$$

A large K_{VCO} means that a small OTA-output voltage produces a large frequency deviation. In the simulation results, the TT sensitivity around 220 mV is about 45 MHz/V, while process corners introduce a wide range from the slow to fast cases. This confirms that VCO behavior must be evaluated under PVT variations, not only under nominal conditions.

3.5 Simulation and verification hierarchy

The detector is verified in a sequence of simulations. First, the first-stage OTA is simulated in open-loop configuration to verify the mid-band gain. Second, the two-stage OTA is simulated in open-loop configuration to evaluate phase margin and unity-gain behavior. Third, the OTA in buffer configuration is simulated to check offset compensation and band-pass behavior at the first-stage output. Fourth, noise analysis is performed for input-referred and output noise. Fifth, transient simulations sweep the VCO control voltage to extract K_{VCO}, phase noise and temperature dependence. Finally, the complete system is simulated using a neural signal plus DC offset and a tissue-equivalent load.

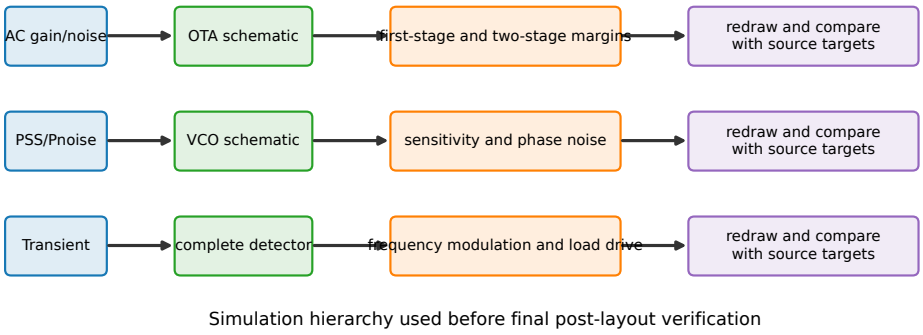

Figure 3.6: Simulation flow of the active neural detector.

3.6 Requirement allocation across the signal chain

The system requirements can be allocated to the blocks by following the signal energy from the electrode to the output. The OTA receives the weakest signal, so it is the dominant block for input-referred noise, electrode-offset tolerance and the first level of gain. The VCO receives a larger internal voltage and is therefore less sensitive to microvolt input noise directly, but it dominates the conversion from voltage information to timing information. The buffer is placed after the signal has already been encoded in the oscillator waveform; its responsibility is therefore load isolation and edge preservation rather than neural amplification. This division of responsibilities prevents over-designing one block while neglecting another.

The area budget follows the same logic. The OTA input devices and compensation elements require enough area to suppress noise and mismatch. The VCO occupies area because multiple inverter stages and bias devices must be laid out with matched parasitics. The output buffer must be wide enough to drive the load, but excessive width increases dynamic power and substrate noise. The final floorplan therefore needs to separate sensitive analog nodes from switching nodes while still keeping the entire active detector inside the compact target area.

3.7 Electrode and load abstraction

A complete electrochemical electrode model is outside the scope of this circuit-level thesis, but the main interface effects are still important for understanding what the detector must tolerate. The electrode can be abstracted as a neural voltage source, a DC offset source, a series impedance and a shunt capacitance associated with the interface. At very low frequency, the offset may drift slowly with time

and electrochemical conditions. In the signal band, the electrode impedance can attenuate or filter the neural component before it reaches the OTA. On the output side, the tissue-equivalent load can be represented as a capacitive or lossy medium that receives a high-frequency signal from the buffer.

This abstraction leads to two design principles. First, the input node must be high impedance so that the electrode is not loaded. Second, the output node must be low impedance enough to make the carrier observable through the load. The detector therefore contains a high-input-impedance analog amplifier followed by a relatively strong digital buffer. The separation is not arbitrary: it is the circuit expression of the difference between sensing a weak bioelectric source and exciting a load with a modulated carrier.

3.8 Architecture-level performance metrics

The most meaningful performance metrics for this detector are not isolated transistor parameters. The first metric is the input-referred noise density and integrated input-referred noise within the neural band. The second metric is the mid-band gain of the OTA and the associated offset-rejection behavior. The third metric is the VCO sensitivity around the chosen bias point. The fourth metric is power consumption, because each additional micro-watt affects energy autonomy and tissue heating. The fifth metric is area, because a miniaturized implant cannot rely on large passive components or extensive routing margins.

A final metric is interpretability of the output. A detector may have adequate gain and low power, yet still be difficult to use if its output frequency is too sensitive to temperature or process variation. This is why the VCO plots include sensitivity and temperature trends rather than only a single nominal frequency. The output is useful only if the receiving algorithm can distinguish neural modulation from slow drift and random phase fluctuation.

Chapter 4

OTA Design, Sizing and Layout Validation

4.1 Role of the OTA

The OTA is the most critical analog block in the detector. It must amplify the neural component, suppress low-frequency offset and keep input-referred noise low. The detector architecture uses a two-stage OTA. The first stage is a differential stage that provides the main gain and dominates the input noise. The second stage behaves as a common-source amplifier and participates in the feedback path that compensates DC offset. The two stages are separated by pseudo-resistor devices M12 and M13, allowing low-frequency feedback without large passive resistors.

The target first-stage gain is close to 50 dB. This value is not arbitrary. A 30 µV neural signal amplified by 49 dB becomes about 8.5 mV, which is sufficiently large to modulate the VCO control voltage. At the same time, the gain should not be so large that output swing is exhausted under expected offsets and process variation. The low-frequency compensation path therefore ensures that the electrode DC offset does not appear as a high-gain signal at the VCO input.

4.2 gm/ID sizing methodology

The OTA sizing follows the gm/ID methodology. In this method, the designer selects an inversion level through the transconductance efficiency g_m/I_D. Weak inversion gives high transconductance efficiency and low current, but it may reduce speed and increase mismatch sensitivity. Strong inversion gives better speed and matching in some cases but at higher current. Moderate inversion often offers a useful compromise for ultra-low-power analog circuits. The gm/ID methodology is well suited to modern CMOS analog design because it separates current efficiency from geometry selection [9, 10].

For the OTA input pair, flicker noise pushes the design toward larger transistor area. The OTA sizing uses PMOS input transistors M1 and M2 with $W = 7.9\,\mu\text{m}$, $L = 5.5\,\mu\text{m}$ and 12 fingers. PMOS devices are commonly preferred for low-frequency input pairs because they can exhibit lower flicker noise than NMOS devices in many CMOS processes. The large channel length also improves intrinsic gain and

reduces flicker noise. The cost is increased parasitic capacitance and area, which must be managed through multifinger layout.

4.3 Transistor-level sizing

Table 4.1 summarizes the transistor dimensions. The table is included in the thesis because it is the link between system specifications and the physical implementation. Although the VCO and buffer use minimum or near-minimum dimensions for compactness, the OTA input and load devices are much larger to satisfy noise and gain constraints.

Table 4.1: Design of the active neural detector. N is the number of fingers.

Device	W **(µm)**	L **(µm)**	N	**Type**
VCO (Voltage Controlled Oscillator)				
M1–M6, M13–M18	0.22	0.30	1	NMOS mvt
M7–M12, M19–M23	0.22	0.25	1	PMOS mvt
Buffer				
M1	0.22	0.18	1	NMOS
M2	0.33	0.18	2	NMOS
M3	0.60	0.18	3	NMOS
M4	1.00	0.18	6	NMOS
M5	0.44	0.18	1	PMOS
M6	0.66	0.18	2	PMOS
M7	1.32	0.18	3	PMOS
M8	1.98	0.18	6	PMOS
OTA (Operational Transconductance Amplifier)				
M1, M2	7.90	5.50	12	PMOS
M3, M4	17.16	16.50	12	NMOS
M5	0.30	1.00	1	PMOS
M6	2.68	2.00	1	NMOS mvt
M7	0.42	1.00	1	PMOS
M8	0.22	1.00	1	PMOS
M9–M11	2.00	1.55	1	NMOS
M12, M13	0.22	0.18	1	PMOS
Capacitors (OTA)				
C1		3 pF		
C2		7 pF		
C3		1 pF		

4.4 Open-loop first-stage behavior

The first-stage open-loop simulation verifies that the neural component is amplified by the intended amount. Figure 4.2 is obtained from the AC simulation data. The mid-band gain is approximately 49 dB, close to the 50 dB specification. The gain remains flat over a wide low-frequency region and then rolls off at high frequency. The reported unity-gain bandwidth is about 24 kHz. This behavior ensures that the neural band can be amplified without excessive attenuation.

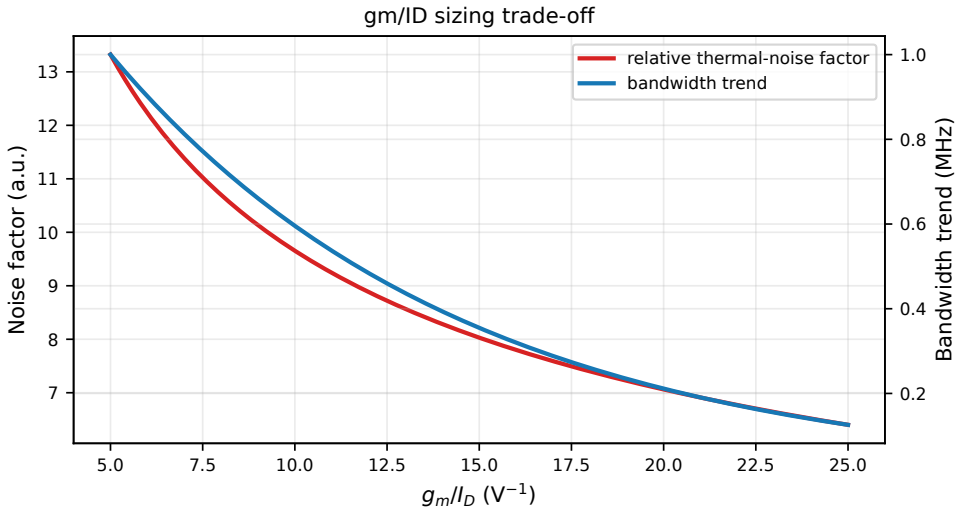

Figure 4.1: gm/ID sizing trade-off for the OTA input stage.

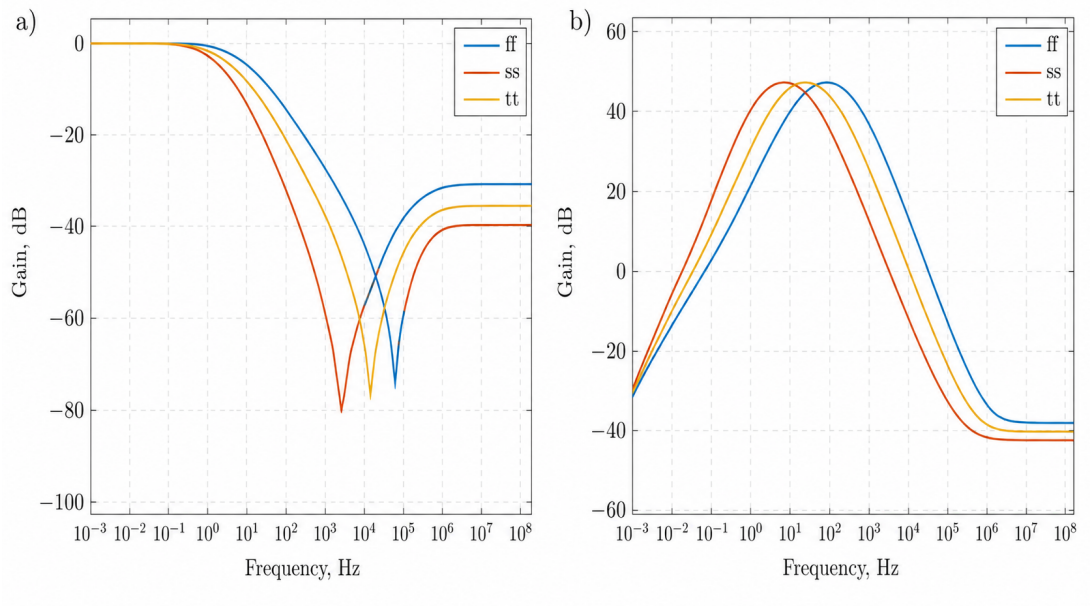

Figure 4.2: First-stage OTA frequency response.

4.5 Buffer-mode OTA and band-pass first-stage response

When the two-stage OTA is configured as a buffer, it corrects low-frequency offset and prevents the large microelectrode DC offset from saturating the first-stage output. When the compensation loop is active, the useful neural band remains amplified while very-low-frequency drift is suppressed. This separation of time constants is the key function of the OTA feedback path.

4.6 Noise behavior

Figure 4.3 presents the OTA noise behavior from the post-layout simulations. The dominant low-frequency component is flicker noise, while the wideband floor depends on thermal noise and transconductance. The design uses large input devices to reduce flicker noise and a very small bias current to remain below the power target.

The noise result should be interpreted together with the signal gain. A high first-stage gain helps the VCO because the neural signal becomes a millivolt-level control variation. However, the same gain also amplifies input noise. The relevant quantity is therefore the signal-to-noise ratio at the VCO control node and the resulting frequency uncertainty. The simulation indicates that the neural-induced frequency variation remains distinguishable from electronic-noise-induced variation in the complete system simulation.

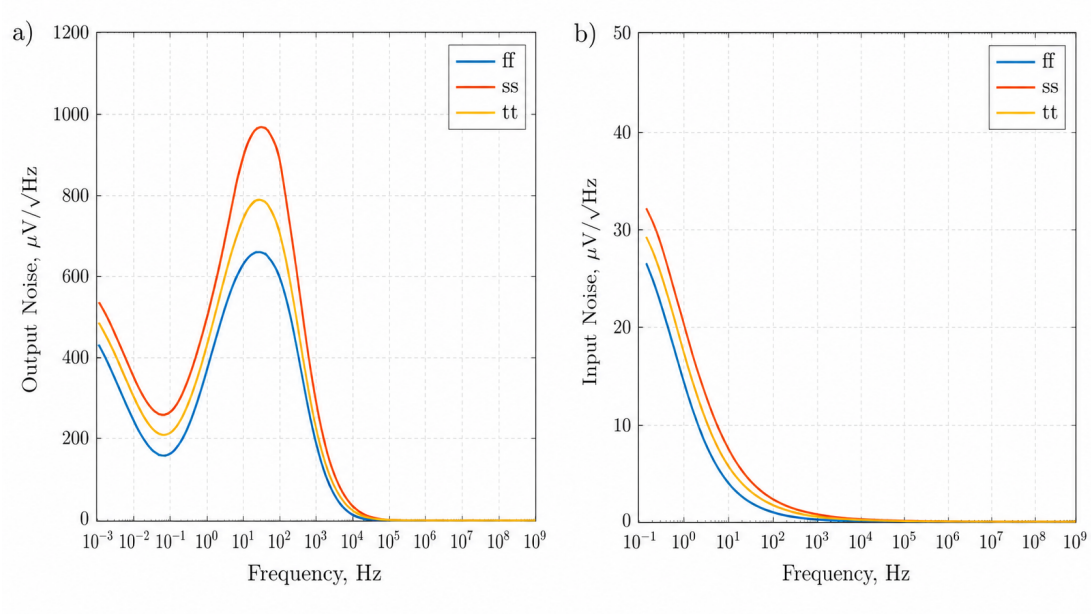

Figure 4.3: OTA output and input noise simulation.

4.7 Layout and parasitic considerations

The layout is a critical part of this design because the target area is extremely small. The layout occupies approximately 162 µm×86 µm. The OTA dominates the silicon area, while the VCO and buffer occupy much smaller regions. Figure 4.4 presents the floorplan as a vector figure. The OTA uses a multifinger layout technique to reduce area and improve matching. Multifinger devices also reduce gate resistance and make routing more compact, but they require careful symmetry and common-centroid or interdigitated placement for matched pairs where appropriate.

Post-layout extraction is necessary because parasitic capacitances can shift pole locations and reduce gain. In this detector, parasitics affect both the OTA bandwidth and the VCO oscillation frequency. The compact VCO relies on small capacitances and minimum transistors, so even small layout parasitics can change the delay per inverter stage. The final performance values discussed in Chapter 6 are therefore interpreted as post-layout-oriented results rather than ideal schematic-only behavior.

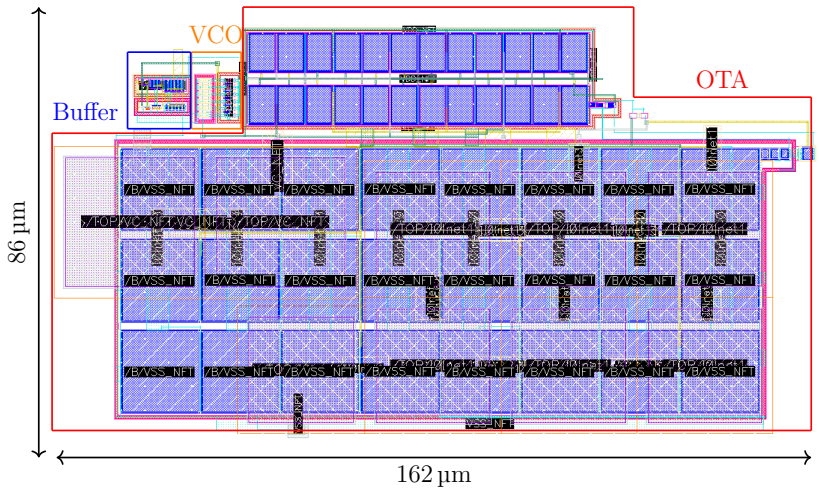

Figure 4.4: Layout of the active neural detector.

4.8 Detailed implementation sequence

The OTA implementation was treated as a sequence of schematic-level, simulation-level and layout-level refinements rather than as a single transistor-sizing exercise. First, the required low-frequency gain was translated into a minimum transconductance and output-resistance product. The input differential pair was biased in weak-to-moderate inversion to increase transconductance efficiency at nanowatt power, while the mirror and cascode devices were kept long enough to provide intrinsic gain and to reduce flicker noise contribution. Second, the compensation capacitor and second-stage current were co-adjusted so that the unity-gain response stayed stable when the OTA was used as a buffer and when it drove the subsequent voltage-controlled oscillator control node. Third, the schematic was swept over process corners, temperature and supply variations. The resulting graphs are reported with consistent line width, font size and caption style, allowing the gain, stability and noise trends to be compared directly.

A practical design detail is that the active detector does not require a high-speed OTA. The extracellular neural signal of interest is concentrated in the low-frequency band, so bandwidth beyond the kilohertz region mainly increases noise and unnecessary current. This observation changes the design priority: the OTA must be quiet, stable and power efficient, while absolute high-frequency speed is secondary. In the schematic, the input pair was therefore made sufficiently large to reduce mismatch and flicker noise, and the current budget was distributed to avoid starving the first stage. The output branch was designed only as strong

as necessary to settle the VCO control voltage during the frequency modulation process. This sizing philosophy is consistent with low-power neural interfaces, where the energy per channel and the electrode-area constraint are often more important than maximum small-signal bandwidth [2, 20, 21].

After a nominal operating point was obtained, the most important verification step was the loop-stability check in the buffer configuration. Since the OTA is not used as an isolated open-loop amplifier in the complete active detector, the buffer-mode phase margin gives a more meaningful indication of practical behavior than the open-loop gain alone. The simulated response was therefore interpreted in two layers. The open-loop result confirms that the two-stage topology has enough intrinsic gain and that the dominant pole is placed at a sufficiently low frequency. The closed-loop result confirms that the control voltage can follow slow neural envelopes without ringing or excessive overshoot. A stable buffer response also reduces the probability that process variation pushes the detector into an oscillatory analog state before the intentional VCO conversion stage.

4.9 Cadence verification flow

The Cadence verification flow followed the order normally used in analog integrated-circuit design. The first simulation was a DC operating-point analysis, used to check the bias current, transistor region, drain-source voltage headroom and output common-mode level. The second simulation was an AC analysis, first in open loop and then in buffer mode. The third simulation was a noise analysis, referred to the OTA input and evaluated around the low-frequency neural band. Finally, parametric sweeps were used to examine the effect of compensation capacitance, load capacitance and bias current. The repeated conclusion of these sweeps is that a small increase of bias current can improve phase margin and input-referred noise, but the improvement becomes inefficient beyond a moderate value because the total detector power is dominated by the oscillator and buffer blocks as well.

For layout verification, the same schematic was converted into a matched layout using multi-finger devices. The input pair was placed in a symmetric geometry so that gradient-induced mismatch is averaged as much as possible. Current mirrors were routed with short and balanced metal paths, and the sensitive OTA input nodes were separated from the VCO switching nodes. This separation is important because the VCO intentionally produces periodic transitions; if those transitions capacitively couple back into the OTA input, they can be interpreted as false neural activity. Guard rings and local decoupling were therefore treated as part of the analog design, not as an afterthought. The final layout abstraction in this chapter

Table 4.2: OTA verification checklist used to organize the schematic and layout simulations.

Verification item	Main observation	Design action
DC operating point	Devices remain biased in intended inversion region under the nominal 0.9 V supply.	Keep long-channel devices in mirror branches and reserve headroom for the second stage.
Open-loop AC gain	Low-frequency gain reaches the required order of magnitude but bandwidth is intentionally limited.	Use compensation to place the dominant pole and avoid excessive high-frequency noise.
Buffer stability	Closed-loop phase margin is around 87° in the nominal reconstruction.	Verify corners and avoid loading the output node with unnecessary parasitic capacitance.
Noise analysis	Input-referred low-frequency noise is near the neural-detection requirement.	Increase input device area and keep bias devices away from the most sensitive nodes.
Layout parasitics	Coupling between the VCO switching path and OTA input path is the main physical risk.	Separate analog input routing, add shielding and preserve symmetry in the matched pair.

highlights this block-level separation and labels the estimated core area.

4.10 OTA current and power consumption

The simulation results report OTA current consumption ranging from approximately 78 nA in the SS process to 411 nA in the FF process, with a TT value of 183 nA. Under the TT process, the OTA power consumption is approximately 329 nW. These values confirm that the OTA remains within the sub-microwatt target. The spread across process corners is significant, but such spread is expected in ultra-low-current analog circuits where threshold voltage and mobility variations strongly influence the operating point.

A practical design would include additional biasing and calibration considerations if the circuit were fabricated as part of a larger system. For example, a current reference must be robust under supply and temperature variation, and the pseudo-resistor operating point must be compatible with expected electrode offsets. In this thesis, the design is analyzed through transistor sizing, layout floorplan and post-layout simulation behavior.

4.11 Input-pair sizing strategy

The OTA input pair is the most important device group for noise and transconductance. In a low-current design, the input pair is normally made relatively wide to reduce flicker noise and mismatch. Increasing width lowers the flicker-noise coefficient at the expense of larger parasitic capacitance. Since the neural signal band is low compared with RF circuits, this trade-off is acceptable up to the point where compensation and stability become difficult. The design therefore favors a wide input pair biased at low current, using gm/ID reasoning to preserve transconductance efficiency.

The tail-current source must also be sized carefully. A poorly isolated tail node converts common-mode disturbance into differential output noise. At the same time, a very large tail device consumes area and adds capacitance. The practical solution is to select a device length larger than the minimum length for output resistance and matching, while keeping the layout compact. This is consistent with the design philosophy: the circuit is not a high-speed operational amplifier, but a low-frequency sensing amplifier that must remain stable, quiet and small.

4.12 Compensation and stability interpretation

The two-stage OTA requires frequency compensation because each gain stage introduces a pole. The compensation capacitor creates a dominant pole and pushes the non-dominant behavior to higher frequency. In this detector, the compensation objective is slightly different from a general-purpose op-amp. The OTA does not need to drive a large arbitrary external capacitance over a wide closed-loop bandwidth. It needs to provide reliable gain in the neural band and a stable internal voltage for the VCO. Therefore the phase-margin target can be selected around robustness and settling, not around maximum unity-gain bandwidth.

A high OTA gain is helpful because it improves the conversion of the weak input signal, but the compensation path must not introduce ringing that would appear as artificial neural activity at the VCO input. The correct design is therefore a separation of time constants: the offset loop acts below the neural band, while the OTA gain remains high over the signal band.

4.13 Noise calculation perspective

For a simplified input-referred noise analysis, the dominant input contribution of a MOS differential pair may be expressed as a combination of thermal and flicker

components,

$$\overline{v_{n,in}^2}(f) \approx \frac{4kT\gamma}{g_m} + \frac{K_f}{C_{ox}WLf}. \tag{4.1}$$

This expression is not sufficient to predict the complete post-layout curve, but it explains the sizing direction. Increasing g_m reduces the thermal term, while increasing WL reduces the flicker term. Both operations cost area or current. The reported bias current near 50 nA indicates that the circuit is designed at the edge of the ultra-low-power regime; therefore the input devices must be used efficiently rather than simply biased harder.

The output-noise curve and input-referred-noise curve are separated in this thesis because they answer different questions. Output noise tells the designer what voltage fluctuation is presented to the VCO. Input-referred noise tells the system designer what neural signal can be resolved at the electrode. A large output noise may still be acceptable if the gain is also large, but a large input-referred noise directly limits the minimum detectable signal. This distinction is often lost when two curves are compressed into one panel; therefore the two results are discussed separately.

4.14 Layout considerations for the OTA

The OTA layout should minimize mismatch in the differential pair and current mirrors. Common-centroid or interdigitated arrangements are commonly used when matching is critical. Guard rings and separated supply routing reduce substrate coupling from the VCO and buffer. Sensitive high-impedance nodes should be short, shielded where possible and kept away from fast switching edges. The compensation capacitor should be placed close to the corresponding internal nodes to minimize unintended parasitic poles. The bias network should be routed symmetrically so that local voltage drops do not create systematic input offset.

Post-layout verification is necessary because the extracted parasitic capacitances are comparable to the small intentional capacitances used in ultra-low-power analog design. A parasitic node that would be negligible in a milliampere op-amp can change the pole position of a nanoampere OTA. This is why the thesis places the layout floorplan before the final system results. The floorplan is not only a mechanical drawing; it is a design decision that determines coupling, parasitics and therefore the final frequency response.

Chapter 5
VCO and Buffer Design Validation

5.1 Current-starved VCO principle

The VCO converts the OTA output voltage into an oscillation frequency. A current-starved ring oscillator is used because it is compact, compatible with digital CMOS devices and easy to control with a low-voltage signal. A ring oscillator consists of an odd number of inverting stages connected in a loop. The oscillation period is approximately twice the total propagation delay around the loop. In a current-starved implementation, transistors limit the current available to charge and discharge the inverter load capacitances. The control voltage changes this current, thereby changing the delay and the oscillation frequency.

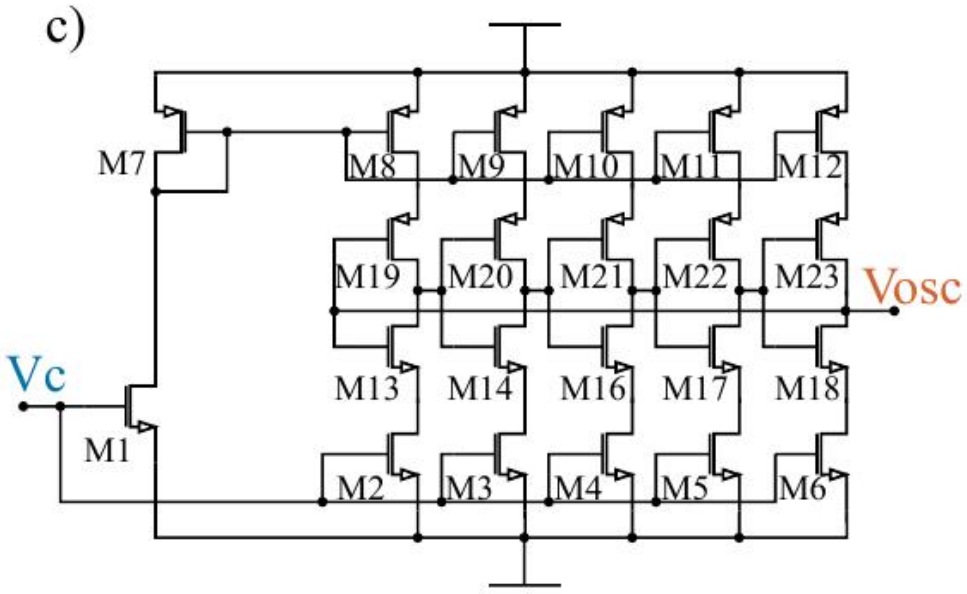

Figure 5.1: Transistor-level schematic of the current-starved VCO.

The approximate expression in Equation (3.2) shows that frequency increases when the control current increases and decreases when the switching capacitance or voltage swing increases. Although the expression is simplified, it captures the design trade-off. Minimum-size devices reduce capacitance and area, increasing the potential frequency and reducing area. However, minimum devices are more

sensitive to mismatch and process variations. Minimum transistor dimensions are used for the VCO to minimize area, accepting process-dependent sensitivity as shown in the simulation results.

5.2 VCO sensitivity

Figure 5.2 presents the VCO sensitivity K_{VCO} as a function of control voltage for process corners. The control voltage is swept from 150 mV to 300 mV in 10 mV steps in the simulations. At the nominal control voltage of 220 mV, the reported sensitivity ranges from about 6.4 MHz/V to 289 MHz/V depending on the process corner, with a TT value of approximately 45 MHz/V. This wide range is an important result. It indicates that the detector can produce a large frequency shift, but it also means that process variation can significantly change the conversion gain.

The expected neural-induced frequency deviation can be estimated. If the OTA first-stage gain is approximately 49 dB, a 30 µV input becomes approximately 8.5 mV. With a TT sensitivity of 45 MHz/V, the corresponding frequency deviation is on the order of 380 kHz. Under faster process conditions the deviation can be larger, whereas under slow conditions it can be smaller. This supports the conclusion that a neural signal of a few tens of microvolts can produce a frequency variation from hundreds of kilohertz to several megahertz around a carrier in the megahertz range.

5.3 Phase noise and detectability

VCO phase noise determines how precisely frequency can be estimated. If phase noise spreads the oscillator spectrum too much, small neural-induced frequency deviations can be hidden. Figure 5.2 presents the phase-noise curve. The simulations indicate approximately −60 dBc/Hz at a 10 kHz offset and approximately −86 dBc/Hz at a 100 kHz offset. These values suggest that the oscillator noise does not mask the expected frequency deviation, which is on the order of hundreds of kilohertz or more.

Phase-noise theory for ring oscillators is more complex than the simplified delay model. It depends on device noise, switching waveform shape, supply sensitivity and impulse sensitivity functions [15, 16]. In this detector, absolute phase-noise optimization is not the only objective. The oscillator must also be compact and low power. Therefore, the VCO design is a compromise: it provides enough frequency sensitivity and acceptable phase noise while occupying a very small area and consuming only hundreds of nanowatts.

5.4 Temperature dependence

Since the detector is intended for implantation, temperature is expected to remain close to body temperature, with smaller variation than in many external electronic systems. Nevertheless, temperature affects mobility, threshold voltage and leakage, which in turn affect the VCO delay. Figure 5.2 presents the temperature sweep for the TT process. The simulations indicate a frequency shift of approximately 50 kHz/°C from 25 °C to 45 °C. This shift is smaller than the neural-induced deviation estimated above, but it is not negligible for long-term operation. A receiver or calibration scheme could track slow carrier-frequency drift while preserving fast neural modulation.

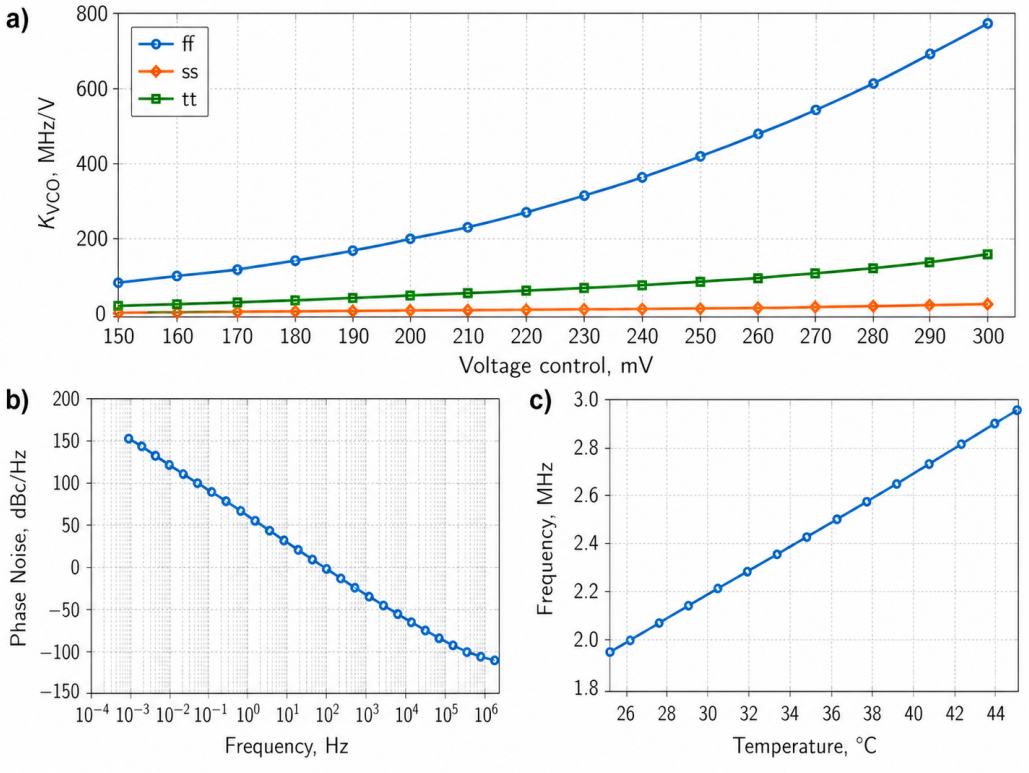

Figure 5.2: VCO simulation results.

5.5 Buffer design

The buffer is placed between the VCO output and the bonding pads or tissue-equivalent load. It uses four inverters with progressively increased transistor widths and constant channel length. This is a classical tapered-buffer idea: small stages avoid loading the oscillator, while larger stages provide the final drive. The

reported load range is approximately 5 pF–20 pF, which includes measurement instrumentation and possible tissue-coupled capacitance.

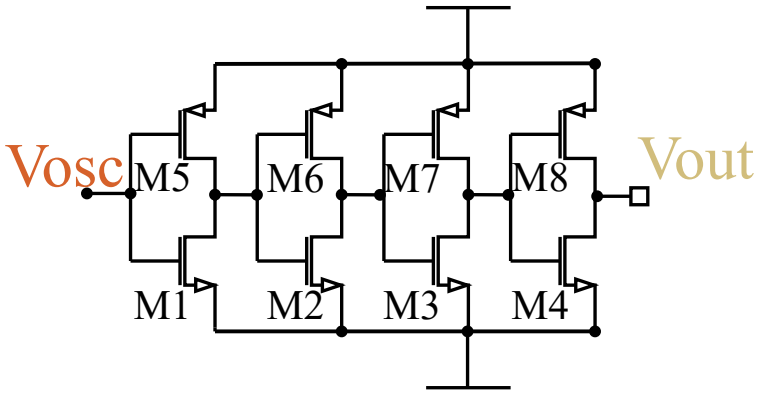

Figure 5.3: Four-stage tapered inverter buffer.

The sizing table in Chapter 4 shows that the NMOS widths increase from 0.22 µm to 1 µm, while PMOS widths increase from 0.44 µm to 1.98 µm. The PMOS devices are wider because hole mobility is lower than electron mobility, and a larger PMOS width helps balance rise and fall times. The last stage must drive the largest capacitance, but making it too large would increase input capacitance seen by the previous stage and increase power. The buffer therefore uses modest tapering rather than a large high-speed digital output driver.

5.6 Power consumption

The simulations indicate VCO power consumption of approximately 609 nW under TT process at a control voltage of 220 mV. Combined with the OTA TT power of approximately 329 nW, the active sensing and conversion chain remains close to 1 µW, excluding the buffer consumption. This is consistent with the implantable-system requirement. In practice, the buffer dynamic power depends on switching frequency, load capacitance and supply voltage according to

$$P_{dyn} \approx \alpha C_L V_{DD}^2 f_{osc}, \tag{5.1}$$

where α is the switching activity factor. Because the VCO operates in the megahertz range, the buffer can dominate power if the load capacitance is large. This is why the final buffer strength must be chosen according to the expected interface rather

than overdesigned.

5.7 VCO design procedure and calibration interpretation

The oscillator was sized after the OTA because the useful control-voltage range depends on the OTA output swing. The design target was not to maximize the oscillation frequency, but to obtain a monotonic and sufficiently sensitive mapping between neural input amplitude and output frequency. In the current-starved topology, the charging and discharging currents of the inverter delay cells determine the oscillation period. A larger control voltage increases the available current, reduces the delay of each stage and therefore increases the oscillation frequency. This simple monotonic relationship makes the oscillator attractive for an active neural detector, because the output can be measured as a frequency or counted digitally over a time window rather than digitized by a conventional voltage-mode analog-to-digital converter.

The calibration procedure can be described in four steps. First, the VCO is biased at a mid-range control voltage and the nominal frequency is recorded. Second, the input-referred OTA response is swept with a small sinusoidal signal and the resulting frequency deviation is measured. Third, the slope of the frequency–voltage curve is used to select a counting window: a longer window improves frequency resolution but reduces temporal resolution. Fourth, the counter threshold is adjusted so that noise-only activity does not trigger a detection event. This procedure connects the analog simulation to the eventual digital interpretation of the detector output. It also explains why VCO linearity is useful but not strictly required; the detector primarily needs repeatable monotonic conversion and sufficient separation between the noise-only state and the active neural state.

In the simulations, the sensitivity is highest at the low-control-voltage side and decreases at higher voltages. This trend is typical for current-starved oscillators because transistor overdrive, velocity saturation and parasitic capacitance change the effective delay nonlinearly. The most practical operating point is therefore selected away from the extremely sensitive edge, where small noise on the control node would produce large frequency jitter, and away from the high-voltage saturation region, where additional neural amplitude would not produce enough frequency separation. The resulting usable control region balances dynamic range, power and detectability.

Table 5.1: Interpretation of the VCO simulation results for detector-level use.

Simulation result	Circuit meaning	Detector-level implication
Monotonic tuning curve	Control voltage changes the current-starved delay.	Frequency counting can replace high-resolution voltage digitization.
Sensitivity variation	Low-control-voltage region gives larger gain.	Operating point must avoid excessive jitter sensitivity.
Phase noise floor	Timing uncertainty appears as count uncertainty.	Counting window and threshold must be selected together.
Temperature slope	Bias and delay vary with temperature.	Baseline frequency should be calibrated or tracked.
Buffer power	The buffer isolates the VCO output.	Digital loading should not disturb the analog oscillator core.

5.8 From phase noise to neural-event resolution

Phase noise matters because the detector output is interpreted through time-domain edge counting. Random timing fluctuation broadens the measured frequency distribution, so two input amplitudes that produce close nominal frequencies may become indistinguishable if the oscillator jitter is too high. The phase-noise curve was therefore translated qualitatively into a detection-resolution requirement. At small frequency offsets, flicker and bias noise dominate. At larger offsets, thermal noise and switching-noise mechanisms dominate. A compact implantable detector cannot eliminate these mechanisms, but it can reduce their impact by choosing a stable control-voltage operating point, limiting unnecessary high-frequency gain in the OTA and avoiding parasitic coupling in layout.

The VCO also interacts with temperature. The temperature sweep shows that frequency increases approximately linearly in the considered range, with a slope of about 50 kHz/°C. For a laboratory prototype this drift can be measured and subtracted during calibration. For an implantable or wearable detector, however, temperature drift should be considered together with biological and electrode-interface variation. A robust implementation would periodically refresh the baseline frequency during quiet intervals and then detect short-term deviations relative to that baseline. This baseline-tracking interpretation is compatible with the active detector concept because the output is already in a frequency-coded form.

5.9 Design implications

The VCO and buffer transform the OTA from a conventional analog amplifier into an active frequency-output detector. The VCO sensitivity links neural voltage to frequency deviation. The phase noise and temperature drift determine how accurately the receiver can recover the modulation. The buffer determines whether the waveform reaches the external or tissue load without excessive attenuation. These three aspects cannot be optimized independently. A stronger VCO drive may reduce some timing uncertainty but increase power. A larger buffer may improve amplitude at the load but load the oscillator and consume more energy. The resulting design reflects a compact compromise suitable for exploratory miniaturized implants.

5.10 Current-starved oscillator design reasoning

A current-starved ring oscillator is suitable for this detector because it provides a compact monotonic relationship between control voltage and oscillation frequency. The control voltage adjusts the available charging and discharging current in each inverter stage. When the current is small, the internal nodes charge slowly and the oscillation frequency decreases. When the current increases, the delay per stage is reduced and the frequency rises. The simplified delay relation can be written as

$$f_{osc} \approx \frac{I_{ctrl}}{2NC_{node}V_{sw}}, \tag{5.2}$$

where N is the number of stages, C_{node} is the effective internal-node capacitance and V_{sw} is the switching voltage swing. This equation is approximate, but it explains why extracted capacitance and bias-current linearity matter directly for the final sensitivity.

The VCO should not be interpreted as a precision frequency reference. It is a conversion element that transforms the amplified neural voltage into a time-varying frequency. Absolute frequency drift can be calibrated if the receiver observes the carrier over time, but local sensitivity must remain sufficiently stable around the operating point. For this reason, Figure 5.2 is included to show monotonicity, sensitivity, phase noise and temperature drift from the simulation results.

5.11 Phase noise and neural information

Phase noise creates uncertainty in zero-crossing times and broadens the spectrum of the oscillator output. In a frequency-encoded detector, this uncertainty competes directly with the useful neural modulation. If the neural signal produces a small

periodic frequency deviation, but the phase noise spreads the carrier power over a comparable offset range, the receiver may not recover the modulation reliably. This is why phase-noise discussion is included even though the detector is not a communication transmitter in the usual sense.

The oscillator phase noise depends on device noise, supply noise, waveform slope and the energy stored in the switching nodes. Low power generally worsens phase-noise performance because less energy is available per cycle. The design must therefore choose a carrier frequency and bias point that are high enough for propagation and frequency measurement, but not so high that dynamic power and phase noise become unacceptable. The phase-noise plot in Figure 5.2 provides a qualitative basis for this discussion.

5.12 Buffer sizing and load isolation

The buffer is necessary because the VCO core should not be directly loaded by the tissue-equivalent medium or by a large capacitive node. A direct load would change the oscillator delay, reduce the waveform slope and introduce amplitude-dependent frequency pulling. A tapered inverter chain is a compact solution. Each stage increases the drive strength gradually, reducing the delay penalty that would occur if a tiny VCO node drove a large output transistor directly. The tapering factor is selected by balancing propagation delay, power consumption and output swing.

Because the output signal is a carrier, the buffer does not need to preserve a small analog amplitude. It needs to preserve timing and provide enough voltage swing to be detected after loading. This justifies a digital-style buffer after the analog VCO. Nevertheless, the buffer is not independent of the analog front-end. Its switching current can disturb the supply and substrate, so the layout must place the buffer away from the OTA input and provide adequate supply decoupling and routing. In this sense, buffer design is both an output-drive problem and a noise-isolation problem.

5.13 Temperature and process sensitivity

Temperature changes affect mobility, threshold voltage and bias currents. Process variations affect transistor dimensions, oxide thickness and device mismatch. Both mechanisms shift the VCO free-running frequency and can also change K_{VCO}. A robust detector should therefore tolerate a moving carrier baseline. The neural information is the modulation around that baseline, not necessarily the absolute carrier frequency. In a practical receiver, slow baseline drift can be estimated and

subtracted, whereas rapid random phase fluctuation is more difficult to remove. This is why the temperature curve in Figure 5.2 is discussed together with the phase-noise curve.

For the buffer, process variation changes output resistance and delay. A faster corner may produce sharper edges but also higher dynamic current and stronger substrate coupling. A slower corner may weaken the drive and make the output more sensitive to load capacitance. The architecture remains acceptable if these variations do not destroy the frequency-encoded information or force the VCO outside the measurable range. The design should therefore be verified not only at TT but also under representative SS and FF conditions.

Chapter 6

Post-Layout Simulation and System-Level Validation

6.1 Post-layout validation purpose

For an ultra-low-power analog and mixed-signal neural detector, schematic-level simulation alone is not sufficient. Routing capacitance, device mismatch and layout parasitics can shift the OTA poles, change the VCO delay and alter the final output waveform. The validation in this chapter uses the post-layout simulation figures to verify the OTA–VCO active-detector concept.

The complete detector is verified through three main blocks: the OTA front-end, the current-starved VCO and the output path through the tissue-equivalent load. The OTA provides weak-signal gain and offset control. The VCO converts the amplified voltage into a frequency variation. The buffer then drives the load while preserving the timing information.

6.2 OTA and VCO validation summary

The OTA frequency response in Figure 4.2 verifies that the first stage provides the voltage gain required before frequency conversion. The OTA noise simulation in Figure 4.3 shows the input-referred noise behavior of the low-current front-end. These two results are the most important OTA-level checks because the detector must amplify a microvolt-level input without allowing the electrode offset or device noise to dominate the useful signal.

The VCO simulation results in Figure 5.2 show the voltage-to-frequency conversion behavior, phase noise and temperature dependence of the current-starved oscillator. The key function of this block is not to generate a perfectly fixed clock. Instead, the VCO intentionally converts the analog control voltage into a measurable frequency shift, so the neural signal is represented in the timing domain.

6.3 Complete system simulation

Figure 6.1 shows the complete system-level transient simulation. The input is a weak neural component superimposed on a much larger DC offset. After

OTA conditioning, the VCO output frequency varies with the neural signal. The propagated output waveform then verifies that the signal can still be observed after the tissue-equivalent load.

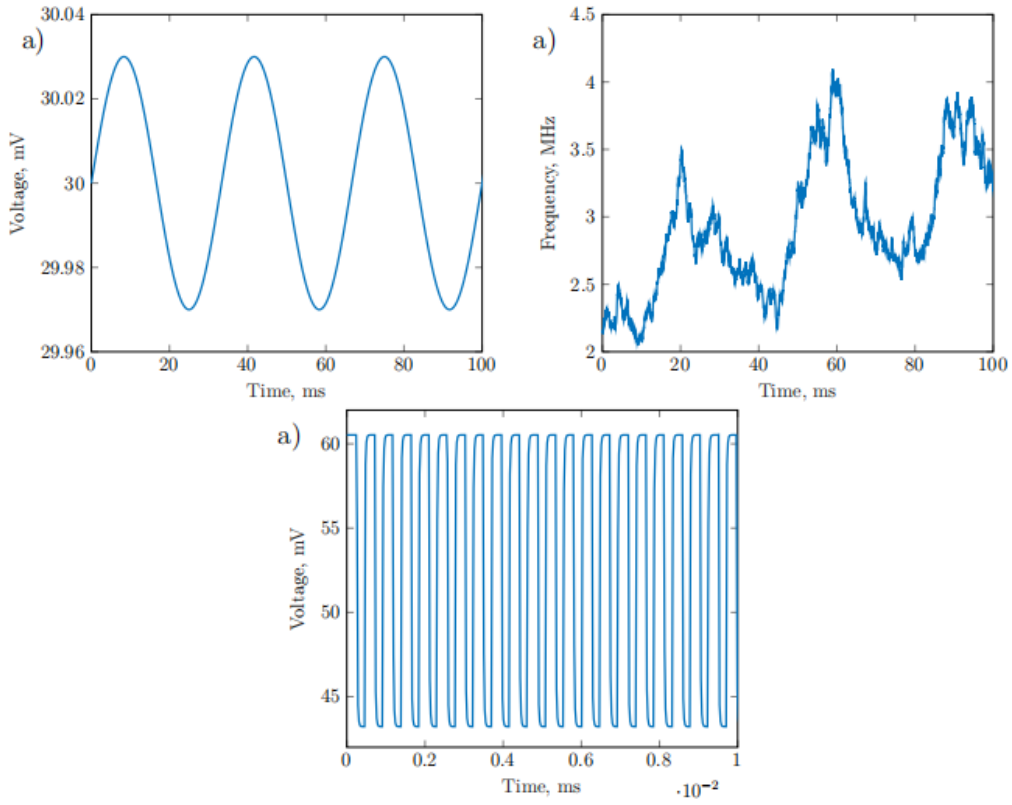

Figure 6.1: System-level transient simulation results.

This result is the main evidence for the active-detector concept. The useful information is no longer carried only by a small low-frequency voltage. It is translated into the frequency variation of a higher-frequency waveform, which is easier to observe after buffering and loading.

6.4 Monte Carlo validation

The Monte Carlo simulation result in Figure 6.2 adds mismatch-level robustness information. The spread of the simulated output parameters indicates that the final detector will require either design margin or calibration, especially if absolute carrier frequency is used. For neural detection, the more important quantity is the relative frequency deviation around the local baseline, which can be tracked by the receiver.

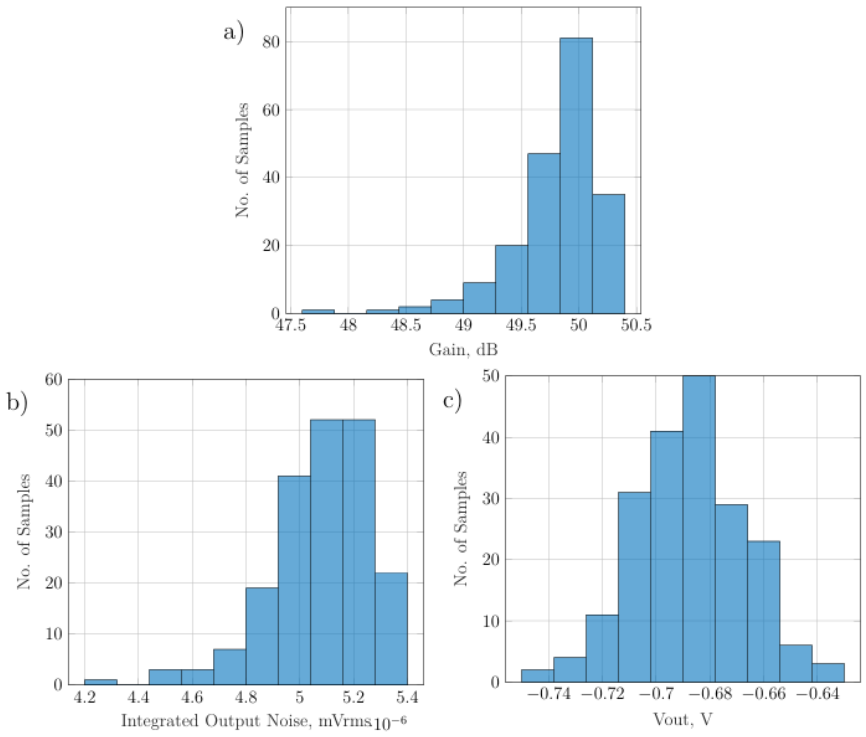

Figure 6.2: Monte Carlo simulation results.

6.5 Quantitative performance summary

Table 6.1 summarizes the main performance values used in this thesis. The OTA and VCO together remain close to the microwatt-level target, while the layout area remains compact enough for a miniaturized neural-interface front-end.

Table 6.1: Summary of simulated performance values.

Metric	Function in the detector	Reported value
First-stage OTA gain	Neural-band voltage amplification	49 dB
First-stage unity-gain bandwidth	High-frequency limit of first-stage response	24 kHz
Two-stage phase margin	Stability of offset-control configuration	87°
First-stage band-pass range	Useful neural band in TT condition	5 Hz–150 Hz
OTA current	Low-power analog front-end bias	183 nA
OTA power	Front-end power consumption	329 nW
VCO control range	Voltage sweep used for conversion	150 mV–300 mV
VCO sensitivity	Frequency shift near nominal control voltage	45 MHz/V
VCO phase noise	Carrier-noise reference values	−60 dBc/Hz at 10 kHz; −86 dBc/Hz at 100 kHz
VCO power	Oscillator power in TT condition	609 nW
Layout area	Final detector floorplan	162 µm×86 µm

6.6 Layout-driven interpretation

The layout in Figure 4.4 shows that the OTA occupies the largest portion of the active detector. This is consistent with the design objective: the input stage must use sufficiently large devices and compensation elements to reduce low-frequency noise and stabilize the front-end. The VCO and buffer are more compact because they rely mainly on inverter-based switching structures.

Parasitics are most important at the OTA compensation network and the VCO delay nodes. The OTA compensation capacitors are intentional and relatively large, so moderate parasitic variation is tolerable. In contrast, the VCO delay nodes use small capacitances to obtain megahertz oscillation at low power, so parasitic capacitance can noticeably shift the carrier frequency. This explains why post-layout validation is necessary.

6.7 Conclusion

The simulation results support the proposed active neural detector architecture. The OTA amplifies the weak input and controls low-frequency offset, the VCO translates the conditioned voltage into frequency variation, and the buffer allows the output waveform to drive the load. The Monte Carlo result shows that absolute carrier frequency can vary, but the frequency-encoded neural information remains the central usable output. This makes the OTA–VCO architecture a compact low-power approach for miniaturized neural-interface sensing.

Chapter 7

Additional Functional Validation under Sinusoidal Neural Inputs

7.1 Objective of the additional validation

The previous chapters validate the main building blocks of the active neural detector, including the OTA front-end, the current-starved VCO and the output buffer. This chapter adds a functional validation step based on the additional sinusoidal-input simulations requested during the thesis activity. The objective is to verify the complete signal path when the input is a weak sinusoidal neural voltage superimposed on an electrode DC offset.

The simulation campaign focuses on the end-to-end behavior of the detector. The input signal is applied at the neural input node, the OTA generates the VCO control voltage V_c, the VCO converts this voltage into an oscillation-frequency variation, and the output buffer provides the final square-like waveform. The relevant observation nodes are therefore the neural input voltage, the control voltage V_c, the oscillator output and the buffered output. This validation is complementary to the block-level simulations because it checks whether the frequency-encoding mechanism remains visible when the detector is driven by time-varying neural inputs.

7.2 Simulation setup

The sinusoidal simulations use a neural input with an amplitude of 30 µV and a DC offset of 30 mV. The input frequency is varied over a wide range in order to test the detector from slowly varying neural activity to frequencies well above the intended neural band. The minimum transient window is selected so that at least three periods of the input sinusoid can be observed after the initial settling interval. Table 7.1 summarizes the simulation conditions.

Table 7.1: Sinusoidal-input validation conditions.

Parameter	Value or setting
Input waveform	Sinusoidal neural voltage with DC offset
Input amplitude	30 µV
Electrode offset	30 mV
Frequency range	1 Hz to 100 kHz
Observed nodes	V_{neural}, V_c, V_{osc} and V_{out}
Validation target	Frequency-modulated oscillator output and square-like buffered output

7.3 VCO operating point under sinusoidal excitation

Since the VCO control voltage is generated by the OTA output, its modulation amplitude depends on the frequency response of the OTA. Around the nominal operating point, V_c is centered at approximately 220 mV. For input signals within the OTA passband, such as 10 Hz–100 Hz, the control voltage exhibits a modulation amplitude of several millivolts. In contrast, low-frequency components around 1 Hz are affected by the DC-offset compensation path, while high-frequency components above the neural band are suppressed by the OTA low-pass behavior. Therefore, the VCO frequency modulation is expected to be most evident within the neural signal band.

Figure 7.1 shows the VCO tuning characteristic around the nominal operating point under the 30 Hz neural-input condition. The simulated operating point is located close to $V_c = 220$ mV, where the oscillation frequency is approximately 2.09 MHz. The local slope of the tuning curve gives a control-to-frequency sensitivity of about 47 MHz/V. This value indicates that a millivolt-level variation of the OTA output is sufficient to generate a measurable frequency deviation at the oscillator output.

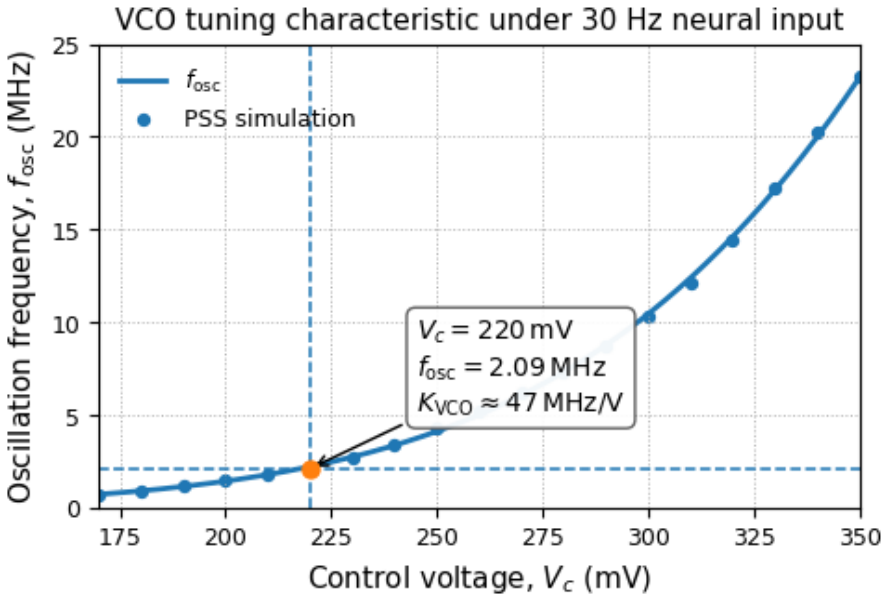

Figure 7.1: VCO tuning characteristic under a 30 Hz neural input.

7.4 Representative transient responses

Figure 7.2 reports the reference transient response for a 30 Hz sinusoidal input. The input voltage contains the microvolt-level neural component and the millivolt-level electrode offset. The frequency plot shows that the detector converts the analog input variation into a time-varying oscillation frequency. The buffered output preserves a square-like waveform, confirming that the final output remains compatible with timing-domain readout.

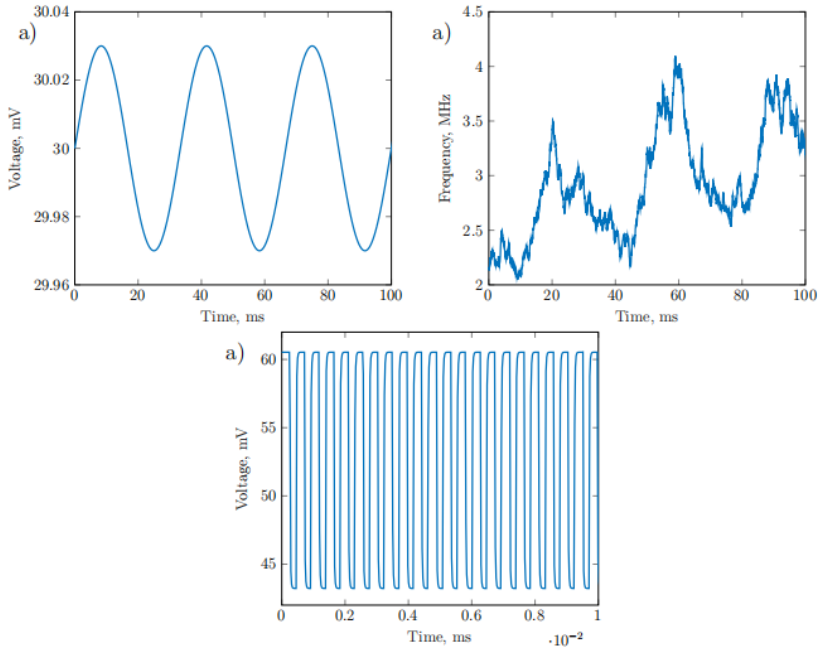

Figure 5.9: a) Input signal composed of neural amplitude and DC offset. b) Frequency response of active neural recorder output signal. c) Propagation of output signal in the brain tissue.

Figure 7.2: Active neural detector response for a 30 Hz sinusoidal input.

7.4.1 Low-frequency input conditions

The low-frequency simulations are useful for verifying the interaction between the electrode-offset cancellation path and the useful neural component. Figures 7.3 and 7.4 show representative responses at 1 Hz and 10 Hz, respectively. At these frequencies, the input envelope is slow with respect to the megahertz carrier, and the VCO output can be interpreted as a carrier whose instantaneous frequency follows the low-frequency control signal.

For the 1 Hz case, the response evolves over a long time window. The frequency variation remains observable, but the response is influenced by the low-frequency compensation behavior of the OTA. For the 10 Hz case, the control voltage falls closer to the useful neural band and the frequency modulation becomes more regular. The buffered output remains square-like in both cases.

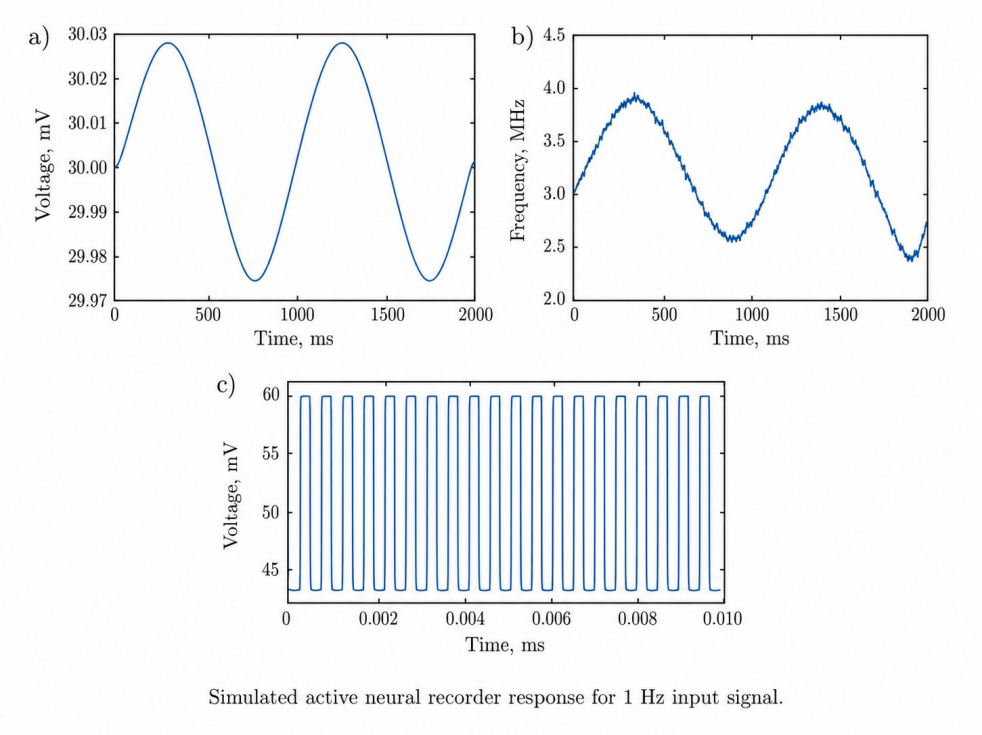

Figure 7.3: Active neural detector response for a 1 Hz sinusoidal input.

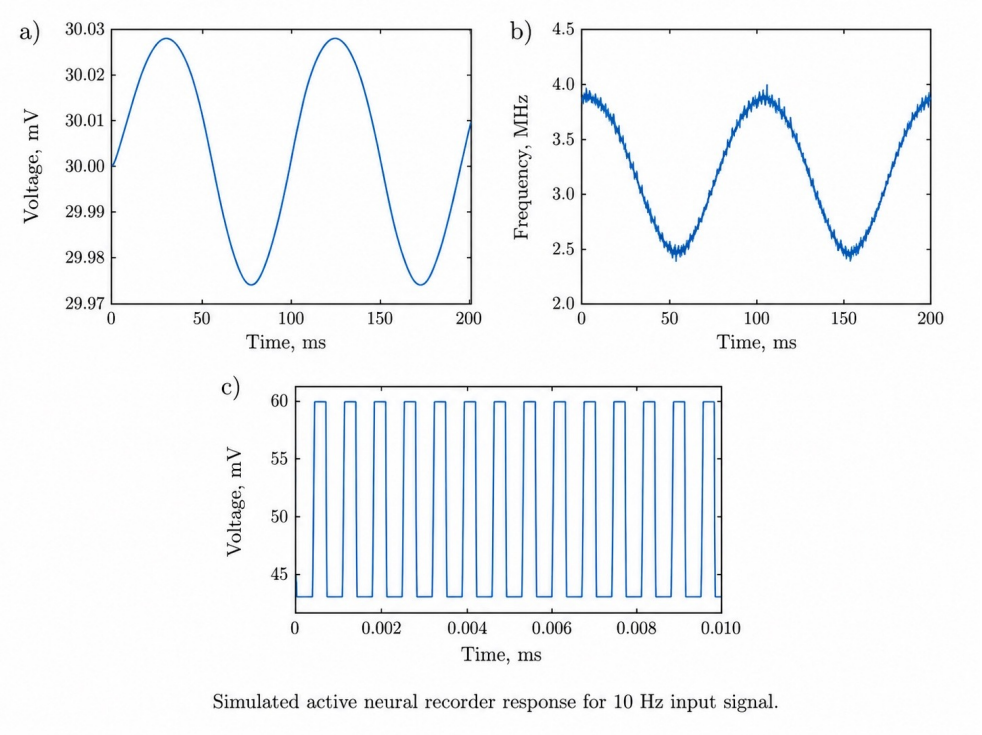

Figure 7.4: Active neural detector response for a 10 Hz sinusoidal input.

7.4.2 Neural-band input conditions

Figure 7.5 shows the transient response for a 100 Hz input. This condition is representative of the upper part of the neural-frequency band considered in the front-end validation. The frequency variation follows the input envelope more clearly than in the very-low-frequency case, indicating that the OTA output provides an effective control signal for the VCO. The buffered waveform remains stable and does not show a loss of switching behavior.

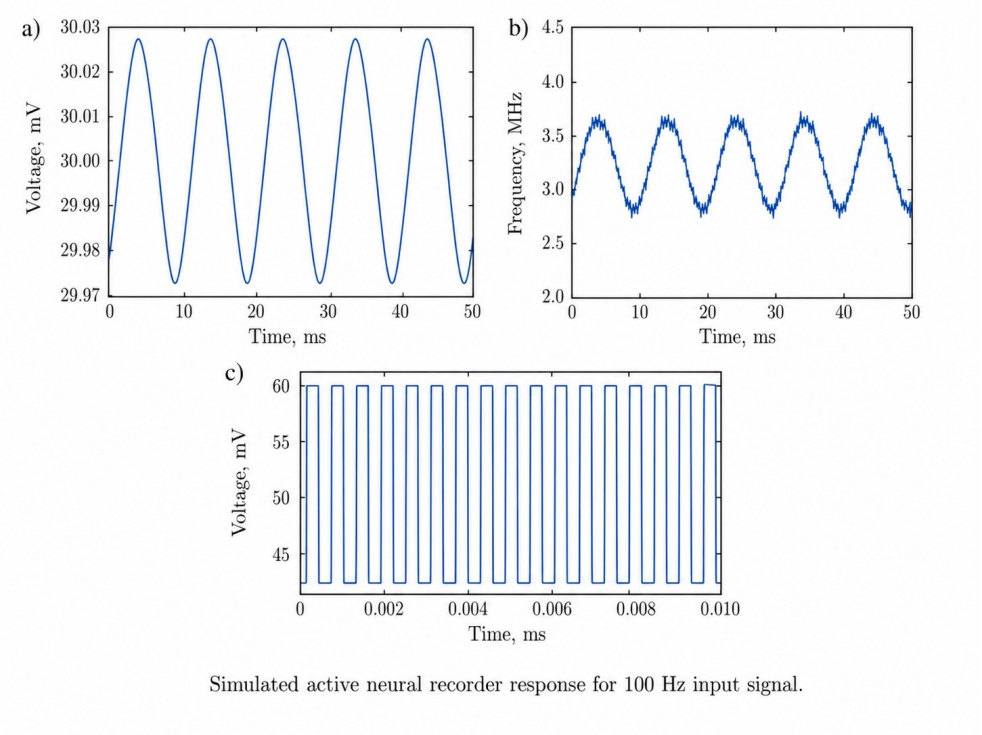

Figure 7.5: Active neural detector response for a 100 Hz sinusoidal input.

7.4.3 High-frequency input conditions

Figures 7.6, 7.7 and 7.8 show the detector response when the input frequency is increased to 1 kHz, 10 kHz and 100 kHz. These conditions are above the main passband of the neural front-end and are included to evaluate the rejection of high-frequency input components.

At 1 kHz, the input waveform is still visible at the observation window, but the frequency modulation is already reduced compared with the neural-band case. At 10 kHz, the frequency trace becomes nearly constant around the carrier frequency, showing that the OTA has reduced the effective control-voltage swing. At 100 kHz, the modulation is further suppressed, and the detector output behaves mainly as a stable high-frequency carrier. This behavior is consistent with the finite bandwidth of the OTA and with the intended operation of the system, which is not designed to preserve high-frequency input disturbances.

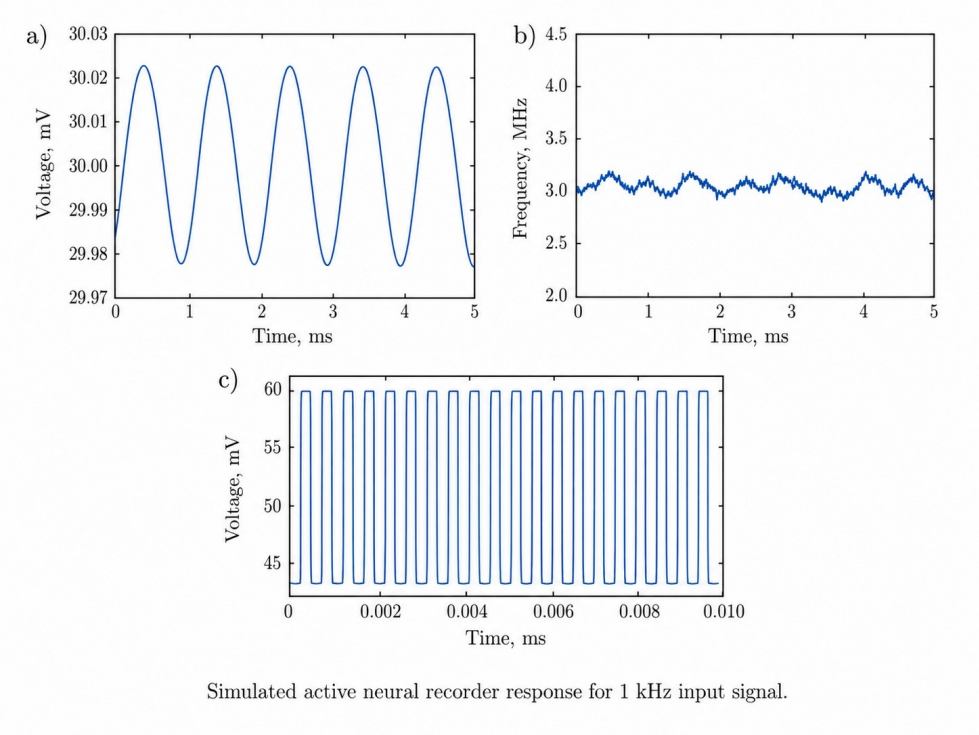

Figure 7.6: Active neural detector response for a 1 kHz sinusoidal input.

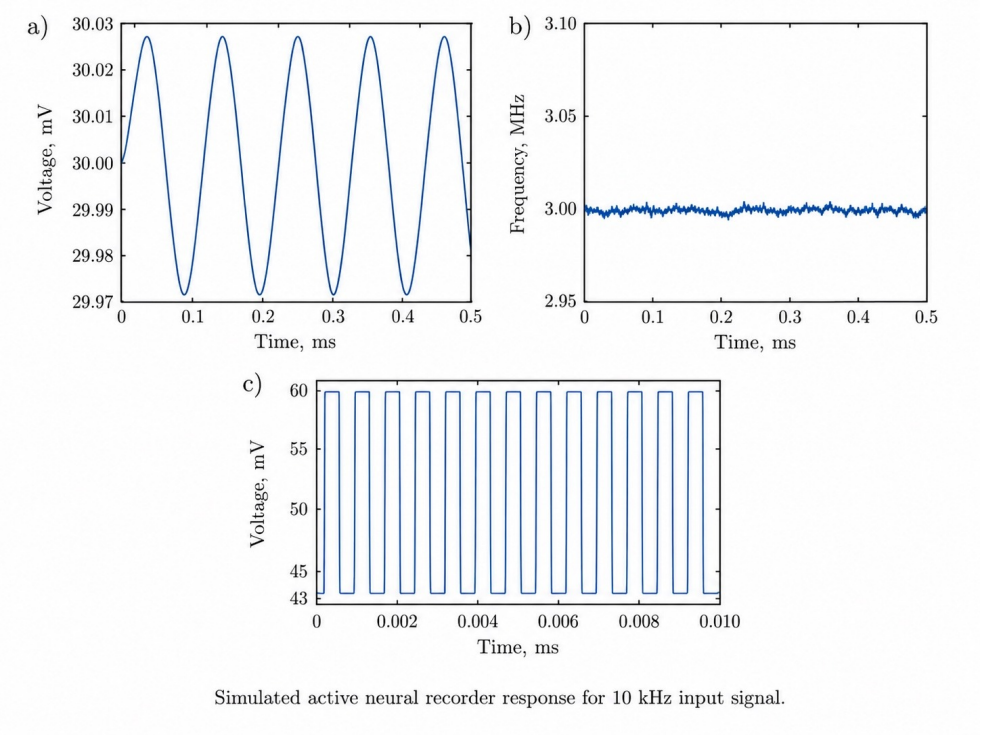

Figure 7.7: Active neural detector response for a 10 kHz sinusoidal input.

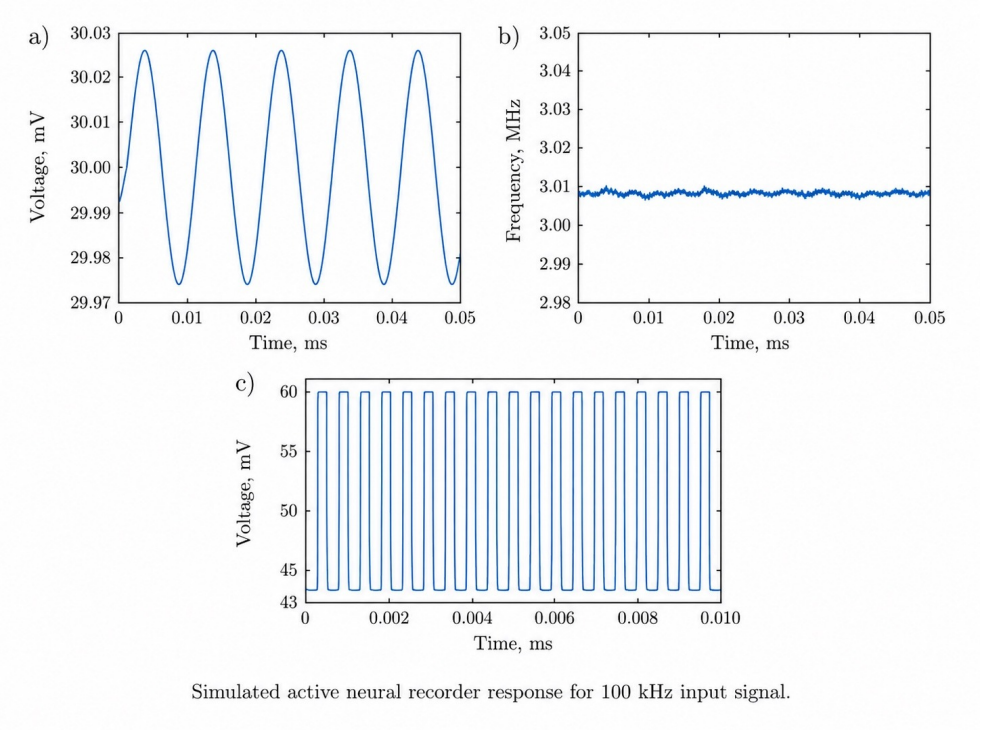

Figure 7.8: Active neural detector response for a 100 kHz sinusoidal input.

7.5 Discussion

The additional simulations confirm the expected frequency-dependent behavior of the complete active neural detector. In the useful low-frequency neural range, the OTA produces a control signal that modulates the VCO frequency, and the output buffer preserves the square-wave waveform needed for timing-domain readout. The neural information is therefore transferred from the voltage domain to the frequency domain rather than being transmitted as a small low-frequency voltage.

The response is not identical for all input frequencies because the OTA is not a broadband voltage amplifier. The offset-cancellation path reduces very slow variations, while the finite high-frequency bandwidth attenuates signals outside the target neural band. As a result, the strongest and most interpretable frequency modulation is obtained for input frequencies close to the useful neural band. At high input frequencies, the VCO continues to oscillate, but the frequency deviation becomes small because the OTA output no longer provides a large control-voltage swing.

Table 7.2 summarizes the qualitative interpretation of the representative simulations. The result confirms that the active detector can operate as a frequency-encoding front-end for weak sinusoidal neural inputs while rejecting components outside the intended operating range.

Table 7.2: Qualitative summary of the sinusoidal-input validation.

Input condition	Observed behavior	Interpretation
1 Hz	Slow frequency variation with compensation influence	Low-frequency response affected by offset control
10 Hz	Clear modulation of the oscillation frequency	Useful low-frequency operation
30 Hz	Nominal validation case with visible frequency encoding	Reference neural-input condition
100 Hz	Frequency tracking remains visible	Upper neural-band validation
1 kHz	Reduced modulation depth	Beginning of high-frequency attenuation
10 kHz	Carrier frequency becomes nearly constant	High-frequency input suppression
100 kHz	Very small frequency deviation around the carrier	Out-of-band rejection

7.6 Chapter summary

This chapter adds the required end-to-end functional validation with sinusoidal neural inputs. The simulations show that the complete OTA–VCO–buffer chain converts microvolt-level sinusoidal input variations into a measurable frequency modulation around a megahertz carrier. The effect is most evident in the neural frequency band, while very slow and high-frequency components are attenuated by the front-end dynamics. These results support the use of the active neural detector as a compact frequency-encoding interface for low-amplitude neural signals.

Appendix A

Supplementary Design Notes

A.1 Approximate neural-to-frequency conversion calculation

Assume a sinusoidal neural input with amplitude $V_n = 30\,\mu\text{V}$ and first-stage gain $A_v = 49\,\text{dB}$. The linear gain is

$$A_{v,lin} = 10^{49/20} \approx 282. \tag{A.1}$$

The corresponding OTA-output amplitude is approximately

$$V_{OTA} \approx A_{v,lin} V_n \approx 8.5\,\text{mV}. \tag{A.2}$$

Using the TT value $K_{VCO} \approx 45\,\text{MHz/V}$, the estimated frequency deviation is

$$\Delta f \approx K_{VCO} V_{OTA} \approx 0.38\,\text{MHz}. \tag{A.3}$$

This simple estimate is consistent with the simulated behavior of the complete system, where the frequency variation is on the order of hundreds of kilohertz to around one megahertz depending on operating point and process corner.

A.2 Dynamic power of the output buffer

For a capacitive load C_L, the switching power can be approximated by

$$P_{dyn} = \alpha C_L V_{DD}^2 f_{osc}. \tag{A.4}$$

This expression explains why the buffer must be sized carefully. A larger final inverter reduces edge degradation but increases the capacitance driven by the previous stage. A larger load also increases power linearly. Since the VCO works in the megahertz range, even small capacitive loads can consume a noticeable fraction of the system power.

A.3 Notes on figures

The architecture, layout and simulation graphics are formatted with a consistent technical style. The plots use aligned labels, compact captions and comparable sizing so that the block functions and verification results can be read without distracting formatting differences.

Bibliography

[1] A. Bontempi, *Wireless Power Transfer for Miniaturized Neural Implantable Medical Devices: CEREBRO Project*, Ph.D. thesis excerpt, Politecnico di Torino, 2025.

[2] R. R. Harrison and C. Charles, "A low-power low-noise CMOS amplifier for neural recording applications," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 6, pp. 958–965, 2003.

[3] W. Wattanapanitch, M. Fee, and R. Sarpeshkar, "An energy-efficient micropower neural recording amplifier," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 1, no. 2, pp. 136–147, 2007.

[4] T. Jochum, T. Denison, and P. Wolf, "Integrated circuit amplifiers for multi-electrode intracortical recording," *Journal of Neural Engineering*, vol. 6, no. 1, p. 012001, 2009.

[5] B. Gosselin, "Recent advances in neural recording microsystems," *Sensors*, vol. 11, no. 5, pp. 4572–4597, 2011.

[6] R. Muller, S. Gambini, and J. M. Rabaey, "A 0.013 mm^2, 5 μW, DC-coupled neural signal acquisition IC with 0.5 V supply," *IEEE Journal of Solid-State Circuits*, vol. 47, no. 1, pp. 232–243, 2012.

[7] X. Qian, Y. P. Xu, and T. H. Lee, "A micropower low-noise neural recording front-end circuit for epileptic seizure detection," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 6, pp. 1392–1405, 2011.

[8] R. F. Yazicioglu, P. Merken, R. Puers, and C. Van Hoof, "A 60 μW 60 $nV/\sqrt{Hz}$ readout front-end for portable biopotential acquisition systems," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 5, pp. 1100–1110, 2007.

[9] C. C. Enz, F. Krummenacher, and E. A. Vittoz, "An analytical MOS transistor model valid in all regions of operation and dedicated to low-voltage and low-current applications," *Analog Integrated Circuits and Signal Processing*, vol. 8, pp. 83–114, 1995.

[10] P. G. A. Jespers and B. Murmann, *Systematic Design of Analog CMOS Circuits: Using Pre-Computed Lookup Tables*. Cambridge, U.K.: Cambridge University Press, 2017.

[11] B. Razavi, *Design of Analog CMOS Integrated Circuits*. New York, NY, USA: McGraw-Hill, 2001.

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

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

[14] R. J. Baker, *CMOS: Circuit Design, Layout, and Simulation*, 3rd ed. Hoboken, NJ, USA: Wiley-IEEE Press, 2010.

[15] A. Hajimiri and T. H. Lee, "A general theory of phase noise in electrical oscillators," *IEEE Journal of Solid-State Circuits*, vol. 33, no. 2, pp. 179–194, 1998.

[16] A. A. Abidi, "Phase noise and jitter in CMOS ring oscillators," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 8, pp. 1803–1816, 2006.

[17] T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*, 2nd ed. Cambridge, U.K.: Cambridge University Press, 2004.

[18] C. Gabriel, S. Gabriel, and E. Corthout, "The dielectric properties of biological tissues: I. Literature survey," *Physics in Medicine and Biology*, vol. 41, no. 11, pp. 2231–2249, 1996.

[19] A. Kiourti and K. S. Nikita, "A review of implantable patch antennas for biomedical telemetry: Challenges and solutions," *IEEE Antennas and Propagation Magazine*, vol. 54, no. 3, pp. 210–228, 2012.

[20] R. Muller *et al.*, "A minimally invasive 64-channel wireless μECoG implant," *IEEE Journal of Solid-State Circuits*, vol. 50, no. 1, pp. 344–359, 2015.

[21] S. Ha *et al.*, "Silicon-integrated high-density electrocortical interfaces," *Proceedings of the IEEE*, vol. 105, no. 1, pp. 11–33, 2017.