POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Ultra-Low-Power Analog Front-End for Audio Scene Classification Based on SNNs

Tesi di Laurea Magistrale in
Electronics Engineering - Ingegneria Elettronica

Author: **Stefano Colombo**

Student ID: 968781
Advisor: Prof. Giorgio Ferrari
Co-advisors: Michele Mastella
Academic Year: 2025-26

Abstract

The increasing demand for always-on edge intelligence has driven the development of ultra-low-power sensing architectures capable of extracting meaningful information while operating under extremely stringent energy constraints. In acoustic inference systems, conventional digital approaches rely on high-resolution Analog-to-Digital Converters (ADCs) followed by computationally intensive feature extraction algorithms, resulting in power consumption that is often incompatible with long-term autonomous operation.

This thesis presents the design and analysis of an ultra-low-power analog front-end intended for audio scene classification through Spiking Neural Networks (SNNs). The proposed architecture performs analog preprocessing before digitization, significantly reducing the computational burden of subsequent processing stages. The system consists of three main building blocks: a low-noise amplifier providing programmable closed-loop gain while minimizing power consumption, a bank of tunable band-pass filters that decomposes the input signal into sixteen frequency channels, and a voltage-to-current converter that performs signal rectification while generating output currents suitable for direct interfacing with analog neuromorphic hardware.

Particular attention is devoted to the analytical study of each circuit block. Small-signal and large-signal models are derived to guide the design process, while stability, frequency response, power consumption and robustness against process and mismatch variations are extensively investigated through circuit simulations. The implementation employs tunable CMOS pseudo-resistors to realize the large time constants required for low-frequency operation within a compact silicon area.

The obtained results demonstrate the feasibility of the proposed analog processing chain as an energy-efficient interface for neuromorphic acoustic sensing systems, providing a suitable foundation for future fully integrated always-on audio classification platforms.

Keywords: ultra-low-power analog design, audio scene classification, spiking neural networks, analog front-end

Abstract in lingua italiana

La crescente diffusione dell'intelligenza artificiale ai margini della rete (edge AI) ha favorito lo sviluppo di architetture di acquisizione caratterizzate da consumi estremamente ridotti, capaci di estrarre informazioni significative operando sotto severi vincoli energetici. Nei sistemi di inferenza acustica, gli approcci digitali convenzionali si basano su convertitori Analogico-Digitali (ADC) ad alta risoluzione seguiti da algoritmi di estrazione delle caratteristiche computazionalmente onerosi, con consumi di potenza spesso incompatibili con dispositivi always-on destinati a un funzionamento autonomo di lunga durata.

Questa tesi presenta la progettazione e l'analisi di un front-end analogico a bassissimo consumo destinato alla classificazione di scene acustiche mediante Spiking Neural Networks (SNN). L'architettura proposta esegue il pre-processing del segnale direttamente nel dominio analogico, riducendo significativamente il carico computazionale degli stadi di elaborazione successivi. Il sistema è costituito da tre blocchi principali: un amplificatore a basso rumore con guadagno programmabile, una banca di filtri passa-banda sintonizzabili che suddivide il segnale in sedici canali di frequenza e un convertitore tensione-corrente che realizza la rettificazione del segnale generando correnti direttamente compatibili con l'interfacciamento verso neuroni analogici neuromorfici.

Particolare attenzione è dedicata allo studio analitico di ciascun blocco circuitale. Modelli a piccolo e grande segnale vengono derivati per guidare il processo di progetto, mentre stabilità, risposta in frequenza, consumo di potenza e robustezza rispetto alle variazioni di processo e mismatch sono valutati mediante estese simulazioni circuitali. L'implementazione impiega pseudo-resistori CMOS sintonizzabili per realizzare le elevate costanti di tempo richieste dal funzionamento alle basse frequenze mantenendo al contempo un'occupazione di area contenuta.

I risultati ottenuti dimostrano la fattibilità della catena di elaborazione analogica proposta come interfaccia ad elevata efficienza energetica per sistemi neuromorfici di inferenza acustica, ponendo le basi per il futuro sviluppo di piattaforme completamente integrate per la classificazione audio always-on.

Parole chiave: progettazione analogica a bassissimo consumo, classificazione di scene acustiche, reti neurali a spike, front-end analogico

Contents

1 | Introduction

1.1. Introduction

The recent revolution in artificial intelligence (AI), driven by deep learning advancements, has delivered remarkable breakthroughs in areas like computer vision and natural language processing. However, this progress has come at a severe computational cost, with state-of-the-art models demanding massive processing power and unsustainable energy consumption.
This work focuses on designing an ultra-low-power audio processing chain tailored to interface with recently developed spiking neural networks (SNNs), which provide a great reduction in power consumption when compared to traditional neural networks.

1.1.1. Spiking Neural Networks Origin and Working Principles

The massive energy requirements of modern deep learning highlight a fundamental architectural bottleneck, commonly referred to as the "memory wall" in traditional Von Neumann computing systems.
In these conventional configurations, the processing unit (CPU or GPU) is physically separated from the memory unit: data must be therefore continuously transferred back and forth between units during computation, which introduces high latencies and dominates the overall energy dissipation profile.
This bottleneck poses a critical constraint when deploying powerful AI models in real-time edge computing, embedded systems, and other power-constrained environments.
To address these limitations, neuromorphic computing emerged as a promising non-Von Neumann alternative. Heavily inspired by the structural and energy-efficient properties of the human brain, neuromorphic systems completely merge computation and memory while processing information asynchronously through event-driven paradigms.
At the core of this hardware approach sit Spiking Neural Networks (SNNs), which act as a technological convergence point bridging biological realism and computational efficiency. Early milestones leading toward this paradigm include foundational discoveries in

bioelectricity by Luigi Galvani [9], who demonstrated that biological tissues can generate and respond to electrical stimulation, and the anatomical characterization of the biological neuron by Santiago Ramón y Cajal [1], who described the human nervous system as a network made of cellular elements (neurons) and contact zones (synapses), in contrast with the idea of Golgi that professed a continuous network.
Spiking Neural Networks are often referred to as the third generation of neural network models, adopting a communication scheme that closely mimics biological neural processing.
To this purpose, SNNs convey information via very short electrical pulses called spikes. Information can then be spike-encoded and read by the neuron in different possible ways, such as: spiking rate, interspike time, spike counting and many others [2].
The basic computational unit in an SNN is typically based on a Leaky Integrate-and-Fire (LIF) model [5], described by the following equation:

$$\tau_m \frac{dV(t)}{dt} = -(V(t) - V_{rest}) + R_{memb} I(t) \tag{1.1}$$

The input current signal $I(t)$ signal is integrated into a membrane capacitance C_{memb}, so that its voltage rises above the resting potential V_{rest}, eventually triggering a comparator that fires a spike at the output of the neuron.
The resistance of the membrane provides a "leaky" discharge of the membrane capacitance, that follows an exponential decay with a time constant $\tau_m = R_{memb} C_{memb}$.
The membrane capacitance, therefore, provides a memory for past inputs, but this memory eventually fades over time.

1.2. Custom Analog LIF Neuron Implementation

To successfully interface the ultra-low-power audio processing chain developed in this work, it is fundamental to thoroughly analyze the silicon implementation of the target neuromorphic hardware. The reference architecture employs a highly optimized, fully analog Leaky Integrate-and-Fire (LIF) neuron, specifically designed to minimize power dissipation and footprint area while providing robust integration dynamics [4, 8]. As shown in the block diagram of Figure 1.1, the empirical circuit design discards complex biomimetic channel equations, retaining only the strictly essential neural properties: charge integration, programmable leakage, and event-driven spike generation.
Unlike biologically plausible models (such as the Hodgkin-Huxley model) that require complex, area-consuming analog circuitry to replicate precise ion channel dynamics, this implementation biases all MOSFETs within the core neural circuit in the weak inversion

(subthreshold) regime. In this operating region, the drain current exhibits an exponential dependence on the gate-to-source voltage, allowing the neuron to operate with resting power consumption in the order of mere femtowatts, at the cost of increased sensitivity to process, voltage, and temperature (PVT) variations.

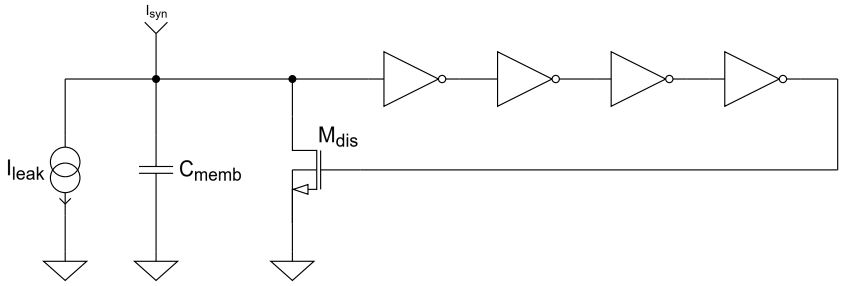

Figure 1.1: Circuit schematic of the custom core ultra-low-power analog LIF neuron implementation displaying the basic integration and leakage node connected to the subthreshold inverter chain.

1.2.1. Subthreshold Integration and Spike Generation

The core computational mechanism of the custom neuron relies on a 100 fF integrated membrane capacitor (C_{memb}), as illustrated in Figure 1.1. When an input event occurs, a generic synaptic current I_{syn} is applied to the membrane node. If the injected current exceeds the neuron's programmed leakage current (I_{leak}), the net charge is integrated onto C_{memb}, causing the membrane potential (V_{memb}) to rise.

Traditional integrated integrate-and-fire neurons often employ power-hungry analog comparators or complex positive-feedback loops to detect when V_{memb} crosses the firing threshold. In stark contrast, this ultra-low-power implementation utilizes a much simpler, highly scalable approach: a chain of four cascaded subthreshold inverters.

The first inverter in the chain acts as the primary threshold detector. Its physical switching threshold acts as the biological firing threshold (V_{th}). As the membrane potential slowly ramps up and crosses this critical voltage level, the first inverter begins to transition. The subsequent cascaded inverters act as high-gain amplifier stages, rapidly squaring up the slow analog transition into a sharp, full-swing digital square wave (0 to V_{dd}). This design choice completely eliminates the static power dissipation typically associated with continuously biased analog comparators and ensures the generation of a clean digital spike, which is crucial for reliable asynchronous routing to subsequent digital layers or external monitoring interfaces.

1.2.2. The Reset Mechanism and Body Effect Exploitation

Once the digital spike is generated at the output of the inverter chain, the circuit must reset the membrane potential back to its resting state to allow for subsequent integrations. This is achieved through a dedicated feedback path.
The output spike is fed back to the gate of an NMOS reset transistor (M_{dis}), connected in parallel with the membrane capacitor as shown in Figure 1.1. During spike generation, this transistor is strongly turned on, rapidly discharging C_{memb} toward the lower supply rail ($V_{ss} = 0.2\,\mathrm{V}$). However, the most critical design challenge in this subthreshold topology is ensuring that the reset device does not leak charge during the slow integration phase. Even when $V_{gs} \approx 0\,\mathrm{V}$, a standard subthreshold NMOS would exhibit an off-state leakage current that could easily overshadow the extremely small (fA to pA range) synaptic and leakage currents, completely disrupting the neuron's temporal dynamics.
To counteract this, the design ingeniously exploits the body effect. By carefully managing the bulk-to-source voltage of the reset transistor, its threshold voltage is artificially increased during the resting and integration phases. This effectively suppresses its subthreshold leakage to near-zero levels, guaranteeing robust integration even across an extended temperature range from $0\,^{\circ}\mathrm{C}$ to $60\,^{\circ}\mathrm{C}$.
Furthermore, physical characterization of the circuit reveals that the reset mechanism is not instantaneous. The finite propagation delay through the four cascaded inverters introduces a brief, voltage-dependent reset delay [7]. During this short window, the membrane potential is temporarily held high before being pulled down. This non-ideal physical behavior inadvertently introduces a short absolute refractory period, which limits the maximum firing rate of the neuron and significantly influences its response to rapid bursts of incoming spikes: a crucial factor to consider when tuning the preceding audio processing chain.

1.2.3. Programmability and Temperature Compensation

A fundamental enhancement in this specific design, differentiating it from basic analog LIF implementations, is the substitution of a static passive membrane resistor with a highly tunable, constant leakage-current generator (I_{leak}). In the context of audio processing, input signals exhibit temporal dependencies across vastly different timescales. Consequently, the neuron must be capable of adapting its integration time window.
Implementing a large time constant (τ_m) using a standard integrated polysilicon resistor would require an unfeasible amount of silicon area. Instead, this architecture employs a programmable current sink utilizing floating-gate (FG) transistors. By precisely injecting

or tunneling electrons onto the isolated floating gate, the threshold voltage of these devices can be permanently altered, allowing for non-volatile programming of the leakage current. This mechanism enables I_{leak} to be finely tuned over three orders of magnitude, ranging from 100 fF to 100 pA, granting direct and precise control over the neuron's "forgetting" rate.

Finally, to address the severe temperature dependence inherent to subthreshold current generators, the leakage circuit implements a differential current extraction topology. By deriving the final leakage current from the difference between two closely matched subthreshold branches, the circuit achieves first-order cancellation of temperature variations. This ensures that the carefully programmed temporal dynamics of the neuron remain stable and reliable, regardless of external environmental fluctuations.

1.2.4. Interfacing with the Neuron

A critical aspect of the ultra-low-power architecture developed in this work is the physical and electrical interfacing between the output of the proposed audio processing chain and the input stage of the custom analog LIF neuron. To preserve the extreme energy efficiency of the neuromorphic substrate and minimize hardware overhead, the interface avoids complex intermediate buffer topologies, relying instead on the direct injection of signal currents into the neural membrane.

As illustrated in the detailed interface schematic of Figure 1.2, the generic input path I_{syn} is physically implemented via two parallel PMOS injector transistors, denoted as M_1 and M_2 [4, 8]. These transistors have their sources connected to the upper supply rail (V_{DD}) and their drains tied directly to the membrane capacitor C_{memb}, operating as voltage-controlled current sources.

The injection mechanism is dynamically driven by two control voltages, V_{gp1} and V_{gp2}, applied directly to the gates of M_1 and M_2 respectively. These two signals carry the processed acoustic information generated by the preceding stages of the audio chain. Depending on the specific operating state of the driving network, M_1 and M_2 modulate their channel conductances to deliver target injection currents directly onto the membrane node.

The overall charge-domain behavior governing the membrane potential (V_{memb}) during the stimulation phase can be expressed by the following differential equation:

$$C_{memb}\frac{dV_{memb}(t)}{dt} = I_{M1}(V_{gp1}) + I_{M2}(V_{gp2}) - I_{leak} \tag{1.2}$$

where I_{M1} and I_{M2} represent the currents sourced by the two PMOS injectors, and I_{leak} is the constant, temperature-compensated leakage current sinking charge from the node [4].
This dual-injector direct-coupling topology provides several distinct advantages for our system:

- **Compact Footprint**: By connecting the drains of M_1 and M_2 directly to C_{memb}, the design completely eliminates the need for intermediate transconductance stages or multi-stage current mirrors, drastically minimizing the silicon area per neural channel [8].
- **Subthreshold Compatibility**: When V_{gp1} and V_{gp2} operate near the supply rail, M_1 and M_2 function strictly within the weak inversion regime [4]. This ensures that the injected currents remain naturally constrained to the femto-to-picoampere range, allowing for slow, energy-efficient integration on a compact 100 fF membrane capacitance [4, 8].
- **Superposition of Features**: The parallel configuration allows the membrane node to act as a physical analog summation point. The neuron natively performs the spatiotemporal integration of the currents modulated by both V_{gp1} and V_{gp2}.

Whenever the combined input currents exceed the baseline leakage ($I_{M1} + I_{M2} > I_{leak}$), the net positive current deposits charge onto C_{memb} [4]. The membrane potential V_{memb} ramps up until it crosses the threshold of the first inverter in the four-stage subthreshold chain, triggering a full-swing digital spike [4, 8]. Immediately following the event, the feedback path drives the gate of the NMOS discharge transistor M_{dis}, introducing a brief, voltage-dependent propagation delay before completely resetting V_{memb} to ground and preparing the node for the subsequent integration cycle [7].

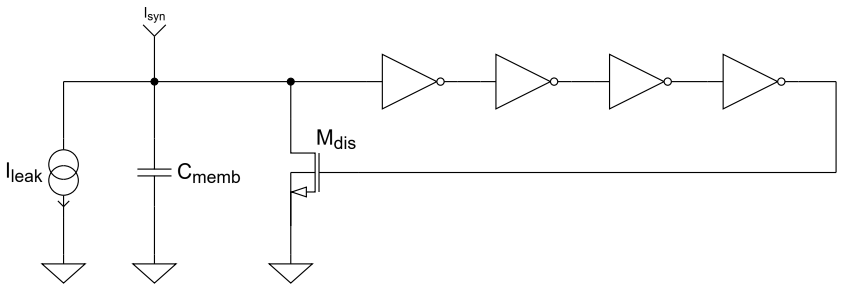

Figure 1.2: Detailed circuit schematic of the analog LIF neuron showcasing the direct current injection interface through the two PMOS transistors M_1 and M_2 mapped onto the membrane capacitor.

2 | Processing Chain Description

2.1. Introduction and State-of-the-Art

The growing diffusion of always-on Internet-of-Things (IoT) nodes, wearable electronics and battery-powered intelligent sensors has shifted the design focus from maximizing computational performance to minimizing energy consumption. In these systems, the sensing front-end continuously monitors the environment while operating under extremely stringent power constraints, often relying on energy harvesting or small-capacity batteries. Consequently, reducing the power consumption of the always-on sensing chain has become one of the main challenges in modern integrated circuit design [11, 14].

In the context of acoustic scene recognition and Voice Activity Detection (VAD), the feature extraction stage represents a significant fraction of the overall computational cost. Conventional digital-intensive systems first digitize the microphone signal using high-resolution Analog-to-Digital Converters (ADCs) and subsequently extract acoustic features through computationally expensive digital processing, typically based on Fast Fourier Transform (FFT) engines and Mel-Frequency Cepstral Coefficients (MFCCs) [11, 12]. Although this approach guarantees excellent flexibility and robustness, the ADC and FFT stages often dominate both area occupation and power consumption, making them unsuitable for sub-μW always-on applications.

To overcome these limitations, increasing research efforts have focused on moving feature extraction before digitization. Rather than reconstructing the input waveform with high fidelity, the analog front-end directly extracts information relevant for the classification task, significantly reducing the amount of data that must be digitized and processed. Recent reviews on sub-μW analog computing architectures demonstrate that this paradigm enables energy reductions of several orders of magnitude, with complete feature extraction front-ends operating at only a few tens of nanowatts [14].

Among the different analog computing paradigms proposed in literature—including voltage-domain, charge-domain and time-domain implementations, as well as level-crossing ADC approaches—the voltage-domain continuous-time architecture has emerged as one of the

most effective compromises between implementation complexity, robustness and power efficiency for acoustic inference systems [14]. The work presented in this thesis follows this design philosophy.

2.2. Digital (MFCC) vs Analog Approaches

The architectural transition from digital to analog feature extraction is motivated by the observation that acoustic inference systems fundamentally differ from conventional signal acquisition systems. While communication and audio recording applications require accurate waveform reconstruction, inference tasks only require features capable of preserving the information relevant for classification.

Theoretical analyses and architectural simulations presented by Yang *et al.* demonstrate that the relatively low-to-medium Signal-to-Noise Ratio (SNR) of everyday acoustic environments considerably relaxes the linearity requirements traditionally imposed on analog circuits without significantly degrading classification accuracy [12]. As a consequence, several design constraints that normally dominate analog integrated circuit design can be traded for substantial power savings.

Although continuous CMOS scaling has dramatically reduced the energy consumption of digital logic, analog circuits benefit much less from technology scaling and remain fundamentally constrained by noise, bandwidth and linearity requirements. Conventional analog design therefore attempts to maximize linearity, often requiring large bias currents and high supply voltages. Conversely, the proposed analog feature extraction paradigm intentionally exploits the intrinsic nonlinear behavior of MOS devices operating in weak and moderate inversion. Circuit blocks such as clipping amplifiers and Class-B/Class-AB rectifiers become advantageous rather than detrimental, allowing complex mathematical operations to be naturally implemented by transistor physics while minimizing static current consumption [12].

This change in design philosophy enables analog feature extraction to achieve substantially higher energy efficiency than conventional digital-intensive solutions while maintaining comparable inference performance. Rather than considering circuit nonlinearities as a limitation, they become an additional design degree of freedom that can be exploited to reduce power consumption.

2.3. Specifications and Frequency Channels

The architecture adopted in this work is inspired by the nanowatt acoustic inference system proposed by Yang *et al.* [12]. The acoustic spectrum is decomposed into multiple parallel frequency channels, each independently extracting the signal energy within a limited frequency band before the information is transferred to the classifier.

The objective of the proposed processing chain is to generate a compact analog representation of the spectral content of the input audio signal before it reaches the neuromorphic classifier. To this end, the microphone signal is first amplified by a low-noise amplifier and subsequently decomposed into multiple frequency bands through a parallel bank of band-pass filters. The output of each filter is then rectified, extracting the envelope (i.e., the modulus) of the corresponding frequency component. Finally, these analog quantities are converted into currents that directly drive the input stage of the target Leaky Integrate-and-Fire (LIF) neuron. Owing to the intrinsic temporal integration performed by the membrane capacitance, the neuron naturally computes a time-average of the envelope associated with each frequency band. Collectively, the outputs of the sixteen parallel channels provide a coarse analog approximation of the spectral distribution of the input signal, preserving the information required for subsequent audio classification while avoiding the computational cost of a conventional digital Fourier transform.

Following the system-level optimization reported in the literature, sixteen frequency channels were selected as an effective compromise between classification accuracy, implementation complexity and overall power consumption [11, 12]. Increasing the number of channels provides only marginal improvements in inference performance while significantly increasing silicon area and static power consumption.

The channel center frequencies are geometrically distributed between approximately 100 Hz and 6 kHz, allowing an efficient representation of the frequency components that contain most of the information required for speech-related inference tasks. Such logarithmic spacing also resembles the frequency resolution of the human auditory system, enabling an efficient allocation of computational resources across the acoustic spectrum [11].

The analog front-end is therefore designed to extract a low-power spectral representation of the input signal rather than faithfully reconstructing its waveform. Although the implemented processing chain does not explicitly compute a Fourier Transform, the decomposition of the signal into sixteen frequency bands, followed by envelope extraction and temporal integration, provides a compact representation of the signal energy distribution across the acoustic spectrum. This representation contains the information required

by the subsequent Spiking Neural Network while considerably reducing the amount of data that must be processed compared to conventional digital approaches based on FFT and MFCC extraction [11, 12].

2.4. Need for Tunable Pseudo-Resistors

The implementation of the proposed analog front-end requires the realization of extremely large time constants in order to place high-pass corners below 100 Hz while maintaining compact integrated capacitors. Since integrated passive resistors with values of several hundreds of megaohms or gigaohms would occupy prohibitive silicon area, transistor-based pseudo-resistors have become the preferred solution for ultra-low-frequency analog integrated circuits [6].

Besides the excessive layout occupation, conventional passive resistors introduce distributed parasitic capacitances that degrade the frequency response and increase high-frequency noise. Pseudo-resistors overcome these limitations by exploiting the exponential current-voltage characteristics of MOS transistors operating in deep subthreshold, achieving equivalent resistance values several orders of magnitude larger than physically realizable passive devices while occupying only a small silicon area [6].

Among the numerous topologies proposed in literature, this work adopts the tunable pseudo-resistor architecture presented by Guglielmi *et al.* [6]. The selected implementation combines symmetrical current-voltage characteristics, wide resistance tuning capability and reduced sensitivity to common-mode voltage, making it particularly suitable for precision analog filtering. By exploiting floating voltage generators and MOS transistors operating in deep subthreshold, equivalent resistance values ranging from several tens of MΩ up to multiple GΩ can be obtained within a very limited silicon area [6].

These characteristics make the proposed topology particularly attractive for the realization of multi-channel analog filter banks, where closely matched low-frequency poles and robustness against Process-Voltage-Temperature (PVT) variations are essential to ensure repeatable feature extraction across all channels. Consequently, the adopted pseudo-resistor architecture represents a key enabling component of the proposed analog front-end.

3 | Low-Noise Amplifier

3.1. Topology Introduction and Requirements

The LNA topology adopted in this work is the one proposed in [12], as shown in Figure 3.1: it is composed of a first inverter-based amplifying stage (M_1 and M_2) and a source-follower second stage (M_3), both biased in weak inversion with currents in the order of tens of nanoamperes at most.

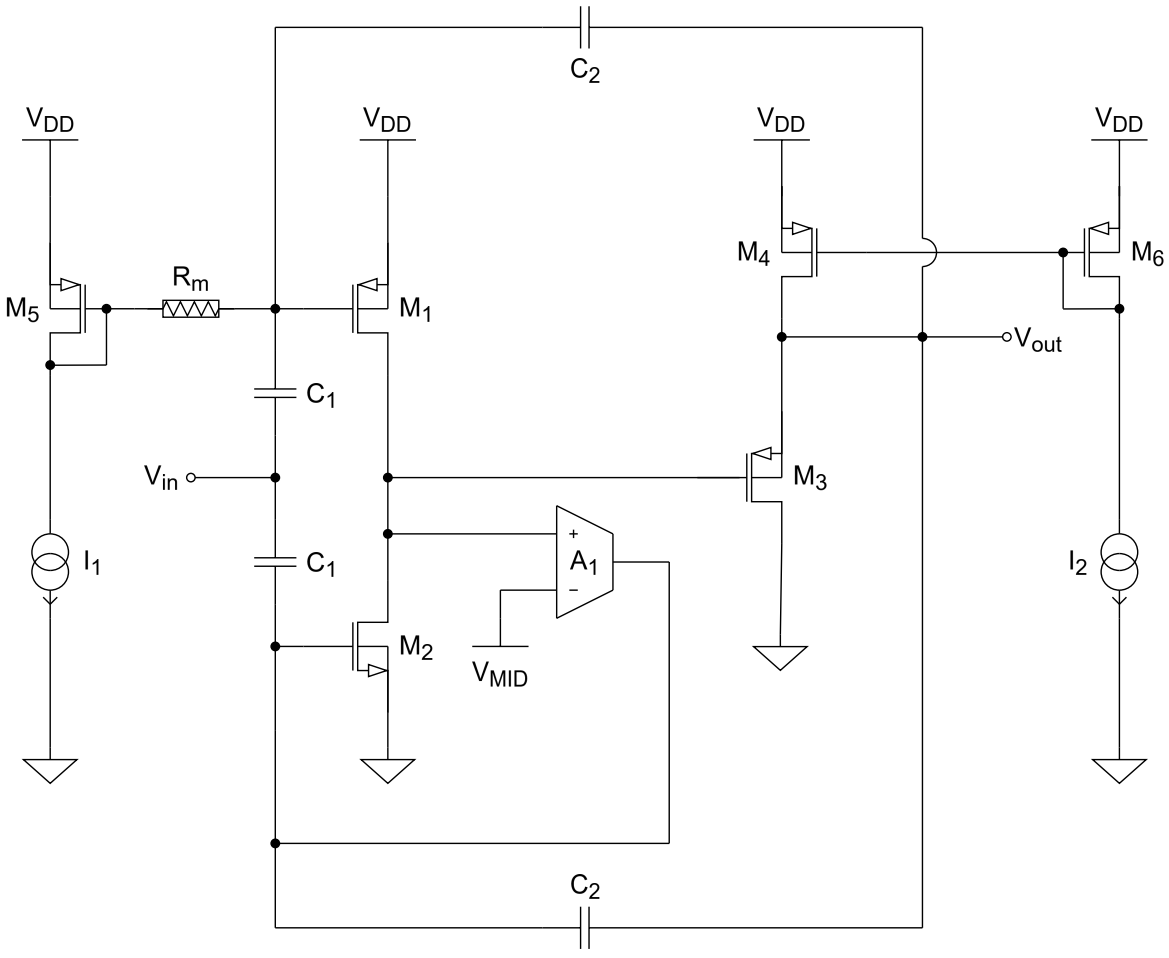

Figure 3.1: Schematic of the proposed LNA.

The reason behind this topology choice is the requirement for ultra-low power consump-

tion, good noise performance and the ability to drive a non-negligible capacitive load at the output.
The power consumption of $\sim$ 15 nW obtained in [12] was a promising figure of merit, even considering the technological difference between the 65 nm node used in [12] and the 180 nm technology of our implementation, which is also biased at a higher supply voltage of 1.8 V.
The capacitive load at the output, due to the 16 band-pass filters next in the processing chain, is driven by a source-follower-based second stage that provides low impedance at the output. As a last consideration, the use of an inverter-based first stage leads to an improvement of the noise efficiency factor (NEF, introduced in [10]), thanks to a transconductance twice as large, as both pMOS and nMOS contribute to signal amplification.
The NEF represents how much noise for a given bias current is introduced by an amplifying stage operating in weak inversion. It is given by:

$$NEF = v_{n,rms}\sqrt{\frac{2\,I_{tot}}{4\,k\,T\,\pi\,V_T\,BW}} \propto \sqrt{\frac{1}{g_m}} \tag{3.1}$$

where $v_{n,rms}$ is the equivalent input rms noise, I_{tot} is the total current consumption of the amplifier, k is the Boltzmann constant, T is the temperature, V_T is the thermal voltage and BW is the bandwidth of the amplifier.
As the above formula shows, an improvement by a factor 2 of the transconductance g_m, for the same bias current, leads therefore to an improvement by a factor $\sqrt{2}$ of the NEF.
For a MOSFET biased in weak inversion the transconductance formula is:

$$g_m = \frac{\partial I_D}{\partial V_G} \simeq \frac{I_D}{n\,V_T} \tag{3.2}$$

supposed the source is set at a fixed voltage, where I_D is the drain current, n is the subthreshold coefficient (assumed ~ 1.3) and $V_T = \frac{kT}{q} \approx 25$ mV is the thermal voltage.
The drain current in weak inversion is given by:

$$I_D = I_0 \exp\left(\frac{V_G - n\,V_S - V_{TH}}{n\,V_T}\right)\left(1 - \exp\left(-\frac{V_{DS}}{V_T}\right)\right) \tag{3.3}$$

where I_0 is the process-dependent current factor, V_{TH} is the threshold voltage of the transistor, V_G is the gate voltage, V_S is the source voltage and V_{DS} is the drain-to-source voltage.

3.1.1. Ideal Gain Analysis

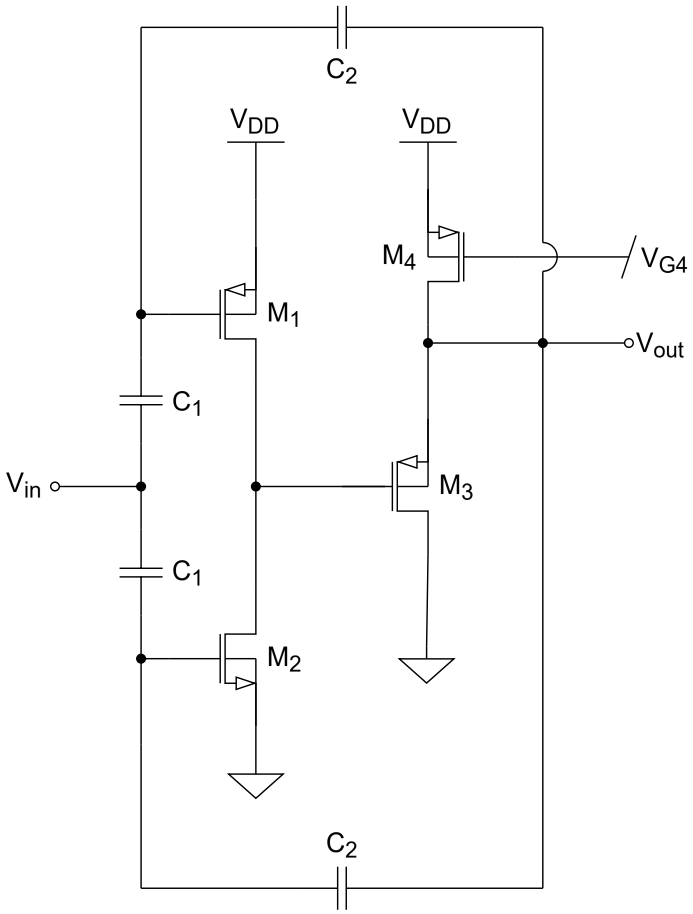

Figure 3.2: Simplified schematic of the LNA for ideal gain analysis.

The first step to understand how this kind of amplifier works is to find its ideal gain, which is the gain of the circuit if the loop closed by C_2 has an infinite gain, constant in frequency.

To this purpose, it is convenient to consider the simplified schematic shown in Figure 3.2, where the biasing network is neglected, including the pseudoresistor R_m (whose purpose will be explained later in this section) and A_1.

As already stated, the first stage consists of an inverter-based amplifier, while the second stage is implemented as a source follower.

The overall feedback structure can be interpreted as an inverting amplifier, in which the two resistors that typically set the gain are replaced by the capacitors C_1 and C_2.

Note that the corresponding upper and lower capacitors are designed of the same value on purpose, to make the two amplifying paths symmetrical.

The resulting in-band ideal closed-loop gain is therefore:

$$G = -\frac{C_1}{C_2}\,. \tag{3.4}$$

The reintroduction of the bias network, as in Figure 3.3, leads to poles and zeros in the transfer function, due to the interaction between capacitors C_1, C_2 and C_{A_1} and the resistors R_m and R_{A_1}.
Note that C_{A_1} and R_{A_1} are respectively the output capacitance and resistance of A_1.
The singularities introduced by the two MOSFET stages will be later analyzed in Sections 3.3 and 3.4, so they are here neglected.
Being that this analysis is aimed at finding the ideal gain of the circuit, the output resistances of M_1 and M_2 are assumed to be infinite, while the output node of M_1 and M_2 is assumed to act as a virtual ground: for this to happen, the drain currents of M_1 and M_2 have to be opposite.
Assuming for simplicity that $g_{m1} = g_{m2}$, the following condition holds:

$$g_{m1}\, v_x = -g_{m2}\, v_y \tag{3.5}$$

which means that the nodes v_x and v_y, corresponding to the gate nodes of M_1 and M_2, move with opposite sign with respect to their DC bias point.
To clarify, the condition above is verified for the main loop (the one formed by the two amplifying stages and the capacitive partition made by C_1 and C_2) at medium frequency, but also for the loop formed by A_1 and M_2, which is active at very-low frequency. The one assumed ideal is the former.

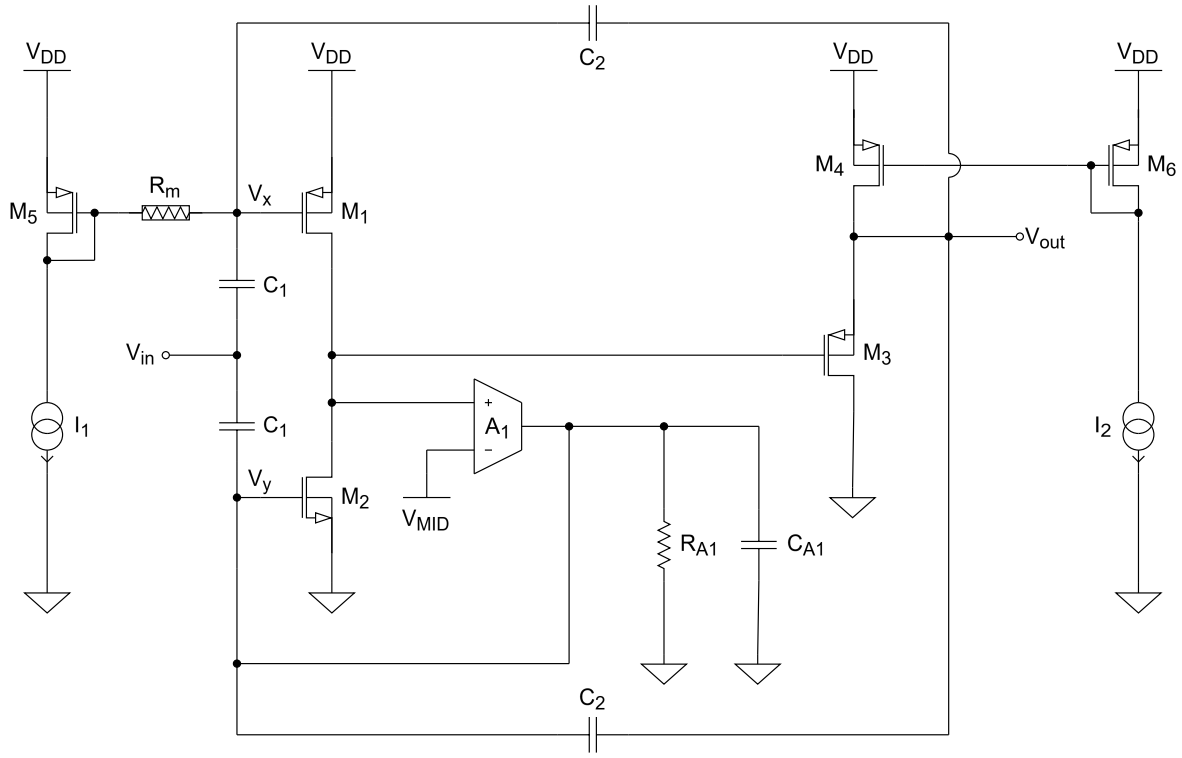

Figure 3.3: Schematic of the LNA for the analysis of the full equation of the ideal gain.

The following equations apply:

$$\begin{cases} v_x = -v_y \\ (v_{in} - v_x)\, sC_1 = \dfrac{v_x}{R_m} + (v_x - v_{out})\, sC_2 \\ (v_y - v_{out})\, sC_2 + v_y \left(sC_{A_1} + \dfrac{1}{R_{A_1}} \right) = (v_{in} - v_y)\, sC_1 + v_{out}\, g_{m,A_1} \end{cases} \tag{3.6}$$

where g_{m,A_1} is the transconductance of A_1, the transfer function of the source-follower stage was assumed ideal and therefore unitary and R_m was assumed much larger than the transconductance of M_5.
The transfer function $H(s) = v_{out}(s) \,/\, v_{in}(s)$ can be then derived as:

$$H(s) = -\frac{sC_1\,R_m}{g_{m,A_1}\;(R_{A_1} \parallel R_m)} \frac{1 + s\,(2\,C_1 + 2\,C_2 + C_{A_1})\;(R_{A_1} \parallel R_m)}{1 + s\,[R_m\,(C_1 + C_2) + (1 + R_m\,/\,R_{A_1})\;C_2\,/\,g_{m,A_1}] + s^2\,(2\,C_1 + 2\,C_2 + C_{A_1})\,C_2\,R_m\,/\,g_{m,A_1}}\,. \tag{3.7}$$

Taking into account that $C_{A_1} \ll 2\,(C_1 + C_2)$ and that by design $C_1 = 10 \cdot C_2 \gg C_2$, to provide an in-band gain equal to 10, $H(s)$ can be simplified as:

$$H(s) \simeq -\frac{sC_1\,R_m}{g_{m,A_1}\;(R_{A_1} \parallel R_m)} \frac{1 + 2\,sC_1\,(R_{A_1} \parallel R_m)}{1 + s\,[R_m\,C_1 + (1 + R_m\,/\,R_{A_1})\,C_2\,/\,g_{m,A_1}] + 2\,s^2\,C_1\,C_2\,R_m\,/\,g_{m,A_1}}\,. \tag{3.8}$$

The ideal gain exhibits therefore two zeros: one in the origin and one at

$$f_{z2} = \frac{1}{4\pi\,(R_{A_1} \parallel R_m)\,(C_1 + C_2)}\,. \tag{3.9}$$

The two poles given by the denominator can be either real if $Q \leq 0.5$ or complex-conjugate if $Q > 0.5$, where Q is the quality factor.
The equation that leads to the two poles can be written as:

$$1 + \frac{s}{Q\,\omega_n} + \frac{s^2}{{\omega_n}^2} = 1 + s\left[R_m\,C_1 + \left(1 + \frac{R_m}{R_{A_1}}\right)\frac{C_2}{g_{m,A_1}}\right] + 2\,s^2\,C_1\,C_2\,\frac{R_m}{g_{m,A_1}} = 0 \tag{3.10}$$

from which we can gather:

$$\begin{cases} \omega_n = \sqrt{\dfrac{g_{m,A_1}}{2\,R_m\,C_1\,C_2}} \\ Q = \dfrac{1}{\omega_n}\left[R_m\,C_1 + \left(1 + \dfrac{R_m}{R_{A_1}}\right)\dfrac{C_2}{g_{m,A_1}}\right]^{-1} = \dfrac{\sqrt{2\,g_{m,A_1}\,R_m\,C_1\,C_2}}{g_{m,A_1}\,R_m\,C_1 + C_2\,(1 + R_m\,/\,R_{A_1})}\,. \end{cases} \tag{3.11}$$

Knowing the formula of Q we can now find the critical value of R_m for which the two poles are real and coincident by setting $Q = 0.5$, which yields:

$$\left(g_{m,A_1} C_1 + \frac{C_2}{R_{A_1}}\right)^2 {R_m}^2 + \left[2 C_2 \left(g_{m,A_1} C_1 + \frac{C_2}{R_{A_1}}\right) - 8 g_{m,A_1} C_1 C_2\right] R_m + {C_2}^2 = 0 \tag{3.12}$$

leading to:

$$R_{m,crit.} = \frac{3 g_{m,A_1} C_1 C_2 - {C_2}^2/R_{A_1} \pm 2 C_2 \sqrt{2 g_{m,A_1} C_1 (g_{m,A_1} C_1 - C_2/R_{A_1})}}{(g_{m,A_1} C_1 + C_2/R_{A_1})^2} . \tag{3.13}$$

A simplified formula can be found assuming $C_1 \gg C_2$ and considering that, for the OTA A_1 to have some gain, it is always true that $g_{m,A_1} \cdot R_{A_1} >> 1$ (which means $g_{m,A_1} >> {R_{A_1}}^{-1}$) and:

$$R^{\pm}_{m,crit.} \simeq \left(3 \pm 2\sqrt{2}\right) \frac{1}{g_{m,A_1}} \frac{C_2}{C_1} . \tag{3.14}$$

The explicit condition to have real poles is:

$$R_{m,crit.} \leq R^{-}_{m,crit.} \vee R_{m,crit.} \geq R^{+}_{m,crit.} \tag{3.15}$$

If we choose the smaller solution, such that $R_m < R^{-}_{m,crit.}$, and compute the two separate poles with the dominant pole approximation, we get:

$$\begin{cases} f_{p1} \simeq \dfrac{1}{2\pi\, b} = \dfrac{1}{2\pi} \dfrac{g_{m,A_1}}{C_2} \\ f_{p2} \simeq \dfrac{b}{2\pi\, a} = \dfrac{1}{4\pi} \dfrac{1}{R_m(C_1 + C_2)} . \end{cases} \tag{3.16}$$

where a and b are the coefficients of the quadratic and linear term, respectively.
Note that the second pole almost cancels out with the zero, as typically $R_{A_1} \gg R_m$.
With typical values for C_2 in the order of tens of fF, even an extremely small value for g_{m,A_1} in the order of tens of pS yields a high-pass cutoff frequency in the order of hundreds of hertz, which is unacceptable for the type of audio application targeted in this work.

If we choose instead $R_m > R^+_{m,crit.}$ the two poles are given by:

$$\begin{cases} f_{p1} \simeq \dfrac{1}{2\pi} \dfrac{1}{R_{A_1}\left(g_{m,A_1} R_m C_1 + C_2\right)} \\ f_{p2} \simeq \dfrac{1}{4\pi}\left(\dfrac{g_{m,A_1}}{C_2} + \dfrac{1}{R_m C_1}\right) \end{cases} \tag{3.17}$$

which is a more acceptable result.
Lastly, substituting $R_m = R^{\pm}_{m,crit.}$ into equation (3.12) we get:

$$\begin{aligned} 1 + s\frac{C_2}{g_{m,A_1}} 2\left(2 \pm \sqrt{2}\right) + s^2 \left(\frac{C_2}{g_{m,A_1}}\right)^2 2\left(3 \pm 2\sqrt{2}\right) = 0 \\ \left(1 + s\frac{C_2}{g_{m,A_1}}\left(2 \pm \sqrt{2}\right)\right)^2 = 0 \end{aligned} \tag{3.18}$$

that leads to the two coincident poles being:

$$f_{p1} \equiv f_{p2} \simeq \frac{1}{2\pi} \frac{2 \mp \sqrt{2}}{2} \frac{g_{m,A_1}}{C_2} \,. \tag{3.19}$$

To give some numbers, in our implementation we chose $C_1 = 300\,\text{fF}$ and $C_2 = 30\,\text{fF}$. As will be detailed in Section 3.2, the parameters of the auxiliary OTA A_1 are set to $g_{m,A_1} \approx 16\,\text{pS}$ and $R_{A_1} \approx 20\,\text{T}\Omega$.
Knowing these circuit parameters, we can compute $R^+_{m,crit.} \approx 36\,\text{G}\Omega$. By tuning the pseudoresistor R_m to get that value, the two poles end up coincident at $f_{p1} \equiv f_{p2} \approx 25\,\text{Hz}$. The frequency of the second zero can also now be computed from equation (3.9) and keeping in mind that $R_m \ll R_{A_1}$:

$$f_{z2} \simeq \frac{1}{4\pi R_m \left(C_1 + C_2\right)} \approx 7\,\text{Hz}\,. \tag{3.20}$$

3.2. Circuit Biasing

To correctly operate, the output node of this first stage must be set to a reference DC voltage V_{MID}, equal to half the supply voltage V_{DD}. This is achieved by means of a secondary control loop governed by the OTA A_1, which senses the difference between the intermediate node and V_{MID} and modulates the gate of M_2 to inject or extract the necessary current to stabilize the operating point.
It must be noted that this secondary loop is effective only at very low frequencies. To

properly evaluate the frequency range in which A_1 forces the intermediate node to V_{MID}, it is necessary to compute the bandwidth of this auxiliary loop.
Cutting the loop at the negative input of A_1, the open-loop DC gain of the auxiliary control system can be derived as:

$$G_{loop,A1}(0) = g_{m,A_1}\, R_{A_1}\, g_{m2}\, (r_{out,M_1} \parallel r_{out,M_2}) = g_{m,A_1}\, R_{A_1}\, g_{m2}\, r_{out} \tag{3.21}$$

where r_{out} is the equivalent output resistance of the first inverter-based stage.
The loop presents two main poles.
The first one is given by the output node of A_1, dictated by its equivalent output resistance R_{A_1} and the parasitic capacitance C_{out,A_1} connected to that node, given by $C_{out,A_1} = C_{A_1} + C_{gs2} + C_{gb2}$. It results:

$$f_{p1,LF-loop} = \frac{1}{2\pi\, R_{A_1}\, C_{out,A_1}}\,. \tag{3.22}$$

The second pole is given by the intermediate node of the two amplifying stages of the LNA: it is determined by the equivalent output resistance of the first stage r_{out} and the equivalent capacitance $C_{out,stage1}$ at that node (which includes the input capacitance of the second stage, the output capacitance of the first stage and the input capacitance of A_1):

$$f_{p2,LF-loop} = \frac{1}{2\pi\, r_{out}\, C_{out,stage1}}\,. \tag{3.23}$$

Given the extremely large value of R_{A_1} and the high transconductance g_{m2} with respect to g_{m,A_1}, $f_{p1,LF-loop}$ becomes the dominant pole of the auxiliary loop.
In our implementation, setting a tail bias current of 1 pA for A_1 results in a transconductance $g_{m,A_1} \approx 16\,\text{pS}$ and an output resistance $R_{A_1} \approx 20\,\text{T}\Omega$. Knowing that $g_{m2} \approx 50\,\text{nS}$ and $r_{out} \approx 21\,\text{G}\Omega$, the resulting DC loop gain is:

$$G_{loop,LF}(0) \approx 16\,\text{pS} \cdot 20\,\text{T}\Omega \cdot 50\,\text{nS} \cdot 21\,\text{G}\Omega \approx 336 \approx 50.5\,\text{dB}\,. \tag{3.24}$$

Considering the parasitic capacitances $C_{out,A_1} \approx 30\,\text{fF}$ and $C_{out,stage1} \approx 32\,\text{fF}$, the frequencies of the two poles can be computed as:

$$\begin{cases} f_{p1,LF-loop} = \dfrac{1}{2\pi \cdot 20\,\mathrm{T\Omega} \cdot 30\,\mathrm{fF}} \approx 0.26\,\mathrm{mHz} \\ f_{p2,LF-loop} = \dfrac{1}{2\pi \cdot 21\,\mathrm{G\Omega} \cdot 32\,\mathrm{fF}} \approx 237\,\mathrm{Hz}\,. \end{cases} \tag{3.25}$$

The frequency at which the auxiliary loop gain crosses the 0 dB axis, identifying the point where A_1 is no longer capable of imposing V_{MID} at the intermediate node (and where the loop can no longer be considered ideal), is the unity-gain frequency $f_{0dB,LF-loop}$. Since $f_{p1,LF-loop}$ is extremely dominant, $f_{0dB,LF-loop}$ can be safely estimated as the gain-bandwidth product of the auxiliary loop:

$$f_{0dB,LF-loop} \simeq G_{loop,LF}(0) \cdot f_{p1,LF-loop} \approx 336 \cdot 0.26\,\mathrm{mHz} \approx 87\,\mathrm{mHz}\,. \tag{3.26}$$

Beyond this sub-hertz frequency limit, the auxiliary feedback mechanism is completely negligible for the AC signal processing, allowing the main AC capacitive feedback (introduced in Section 3.1.1) to dictate the LNA response.

3.3. Open Loop Circuit Analysis

3.3.1. First Stage

As already described, the first stage is inverter-based. Its gain is therefore given by:

$$G_I = -\left[g_{m1}\left(r_{out,M_1} \parallel r_{out,M_2}\right) + g_{m2}\left(r_{out,M_1} \parallel r_{out,M_2}\right)\right] \tag{3.27}$$

where g_{m1} and g_{m2} are the transconductances of transistors M_1 and M_2, while r_{out,M_1} and r_{out,M_2} are their output resistances.
Note that, differently from Section 3.1.1, here it is assumed $v_x = v_y$. The output resistance of a MOSFET in saturation regime is computed as:

$$r_{out} = \left(\frac{\partial I_D}{\partial V_{DS}}\right)^{-1}\Bigg|_{V_{GS}=const.} = \left[I_{D0}\,\frac{\partial(1+\lambda\,V_{DS})}{\partial V_{DS}}\right]^{-1} = \frac{1}{\lambda\,I_{D0}} = \frac{V_A}{I_{D0}} \tag{3.28}$$

where V_A is the Early voltage, proportional to the transistor length, and I_{D0} is given by:

$$I_{D0} = I_0 \exp\left(\frac{V_G - n\,V_S - V_{TH}}{n\,V_T}\right) \tag{3.29}$$

with I_0 being the characteristic current of the device.
Supposed, as a simplification:

$$\begin{cases} g_{m1} = g_{m2} = g_m \\ r_{out,M_1} = r_{out,M_2} = r_{out} \end{cases} \tag{3.30}$$

the gain ends up being:

$$G_I = -\, g_m\, r_{out}\,. \tag{3.31}$$

The stray capacitance at the output node, already analyzed in Section 3.2 as $C_{out,stage1}$, introduces a first pole in the open loop transfer function, that will be named f_{p3} to avoid confusion with the poles found in Section 3.1.1.
The frequency of this pole is given by:

$$f_{p3} = \frac{1}{2\pi\,(r_{out}\,/\,2)\,C_{out,stage1}} = \frac{1}{\pi\, r_{out}\, C_{out,stage1}}\,. \tag{3.32}$$

Recalling the numerical values of transconductance ($g_m \approx 50\,\text{nS}$) and output resistance ($r_{out} \approx 21\,\text{G}\Omega$) introduced in Section 3.2, the resulting gain is:

$$G_I = -\, g_m\, r_{out} \approx -1050\,. \tag{3.33}$$

The frequency of the pole can be computed from equation (3.32) as:

$$f_{p3} = \frac{1}{\pi\, r_{out}\, C_{out,stage1}} \approx 237\,\text{Hz}\,. \tag{3.34}$$

and the gain-bandwidth-product is:

$$\text{GBWP} = |G_I| \cdot f_{p3} \approx 250\,\text{kHz}\,. \tag{3.35}$$

3.3.2. Second Stage

The second stage is a source-follower buffer, whose gain is:

$$G_{II} = \frac{g_{m3}\, r_{out,M_4}}{1 + g_{m3}\, r_{out,M_4}} \approx 1 \tag{3.36}$$

as we can consider $g_{m3}\, r_{out,M_4} \gg 1$.
The large capacitance at the output node, mainly due to the input transistor of the 16 BPFs that follow in the processing chain, leads to a second pole in the open loop frequency response of the LNA, named f_{p4}.
As ${g_{m3}}^{-1} \ll r_{out,M_4}$, the output resistance of the stage is:

$$r_{out} = {g_{m3}}^{-1} \parallel r_{out,M_4} \simeq {g_{m3}}^{-1}\,. \tag{3.37}$$

The output capacitance is given by:

$$C_{out} \simeq 16 \cdot C_{g,in,BPF} \tag{3.38}$$

which results $\sim 240\,\text{fF}$ in our implementation, being $C_{g,in,BPF} \approx 15\text{fF}$.
The bias current for this stage was chosen to be 10 nA, resulting in a transconductance of M_3 of $\sim 216\,\text{nS}$. This specific value was selected to push the second pole frequency high enough to guarantee a good phase margin, considering the chosen closed-loop ideal gain of 10 and the GBWP of 250 kHz (a detailed stability evaluation will be provided in Section 3.4).
The frequency of this second pole ends up being:

$$f_{p4} = \frac{1}{2\pi\, r_{out}\, C_{out}} \simeq \frac{g_{m3}}{2\pi\, C_{out}} = \frac{216\,\text{nS}}{2\pi\, 240\,\text{fF}} \approx 143\,\text{kHz}\,. \tag{3.39}$$

3.4. Closed Loop Analysis

We know from the control theory that the closed loop transfer function of a system is given by:

$$G_{real}(s) = \frac{F(s)}{1 - G_{loop}(s)} = \frac{F(s)}{1 - \beta(s)\, A(s)} \tag{3.40}$$

where $A(s)$ is the open loop transfer function of the amplifier, $F(s)$ is the transfer function of the forward gain and $\beta(s)$ is the transfer function of the feedback signal, from the output to the summing node v_x, shown in Figure 3.4.

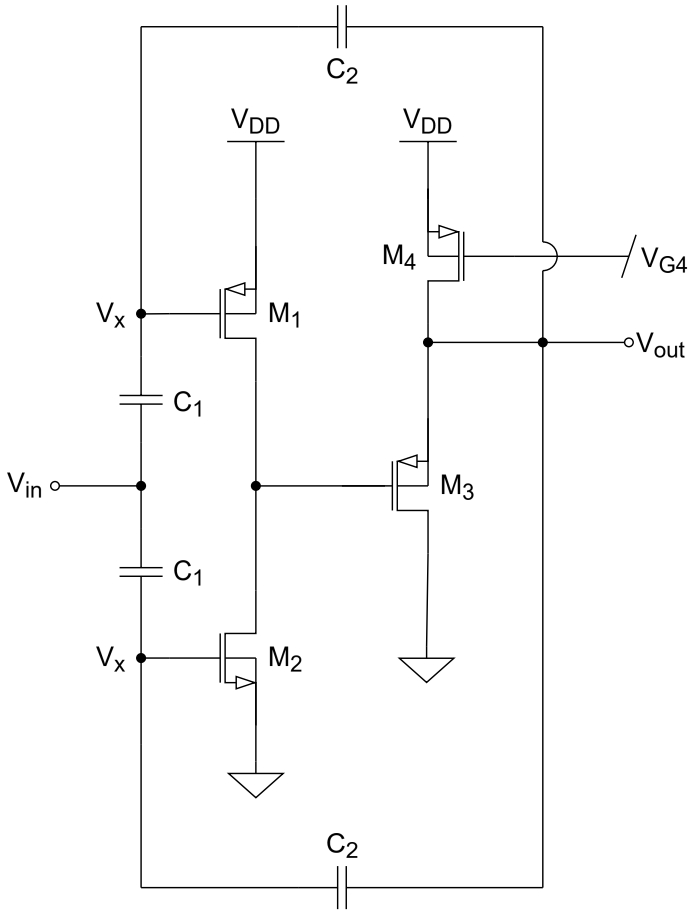

Figure 3.4: Simplified schematic of the LNA for control theory analysis.

If $|G_{loop}(s)| \gg 1$:

$$G_{real}(s) = -\frac{F(s)}{\beta(s)\,A(s)} = G_{id}(s) \tag{3.41}$$

which is called ideal gain of the closed loop system.
Once $|G_{loop}(s)| < 1 \;\Rightarrow\; \beta^{-1}(s) > A(s)$, the closed loop gain becomes equal to the open loop gain and follows it.
$F(s)$ is given by $A(s)$ reduced of the capacitive partition formed by C_1 and C_2, so that:

$$F(s) = A(s)\,\frac{1\,/\,(sC_2)}{1\,/\,(sC_1) + 1\,/\,(sC_2)} = A(s)\,\frac{C_1}{C_1 + C_2}\,. \tag{3.42}$$

The transfer function of the control system is then:

$$G_{real}(s) = \frac{F(s)}{1 - \beta(s)\,A(s)} = \frac{C_1}{C_1 + C_2}\,\frac{A(s)}{1 - \beta(s)\,A(s)} \tag{3.43}$$

The feedback factor $\beta(s)$ is given by:

$$\beta(s) = \left(\frac{1}{sC_1} + \frac{1}{sC_2}\right)^{-1} \frac{1}{sC_1} = \frac{C_2}{C_1 + C_2}\,. \tag{3.44}$$

We already know $A(s)$ from equations (3.31), (3.32) and (3.39):

$$A(s) = -\frac{|G_I|}{(1 + s\tau_3)(1 + s\tau_4)} \tag{3.45}$$

where τ_1 and τ_2 are the time constants of the two open loop poles.
The ideal gain is then, as expected:

$$G_{id}(s) \simeq G_{real}\,(A(s) \to \infty) = -\frac{C_1}{C_1 + C_2}\,\frac{1}{\beta(s)} = -\frac{C_1}{C_1 + C_2}\,\frac{C_1 + C_2}{C_2} = -\frac{C_1}{C_2}\,. \tag{3.46}$$

We can find the first pole of the real gain, the closed loop pole f'_{p3}, knowing that it is located at the frequency at which $A(s) = \beta^{-1}(s)$.
In our implementation $G_{id}(s \to \infty) = 10 = 20\,\text{dB}$ was chosen, with $C_1 = 300\,\text{fF}$ and $C_2 = 30\,\text{fF}$, so that:

$$\beta^{-1}(s) = \frac{C_1 + C_2}{C_2} = 11 \tag{3.47}$$

and:

$$f'_{p3} \simeq \frac{\text{GBWP}}{\beta^{-1}(s)} = \frac{250\,\text{kHz}}{11} \approx 23\,\text{kHz}. \tag{3.48}$$

The phase margin of the closed loop system is given by:

$$\phi_M = 180° - 90° - tan^{-1}\left(\frac{f'_{p3}}{f_{p4}}\right) = 180° - 90° - tan^{-1}\left(\frac{23\,\text{kHz}}{143\,\text{kHz}}\right) \approx 81° \tag{3.49}$$

which guarantees the stability of the system.
To better complete the analysis, it is worth noting that the closed loop transfer function exhibits two feedforward zeros.
To find their frequency the following equations must be verified:

$$\begin{cases} v_{out} = 0 \\ 2\, g_{m1}\, v_{in} + v_{out,stage1}\, sC_{out,stage1} \simeq 0 \\ 2\,(v_{in} - v_{out})\, sC_2 - g_{m3}\, v_{out,stage1} = 0 \end{cases} \tag{3.50}$$

where $v_{out,stage1}$ is the output voltage of the first stage. It results:

$$s_z^2 = -\frac{g_{m1}\, g_{m3}}{C_2\, C_{out,stage1}} \Rightarrow s_z = \pm\sqrt{\frac{g_{m1}\, g_{m3}}{C_2\, C_{out,stage1}}} \Rightarrow f_{z3} = f_{z4} \approx 380\,\text{kHz} \tag{3.51}$$

which means that they are two coincident zeros, one positive (right-half-plane) and one negative (left-half-plane).
The Bode plot of $A(s)$ and $G_{real}(s)$, included the AC coupling discussed in Section 3.1.1, is shown in Figure 3.5.

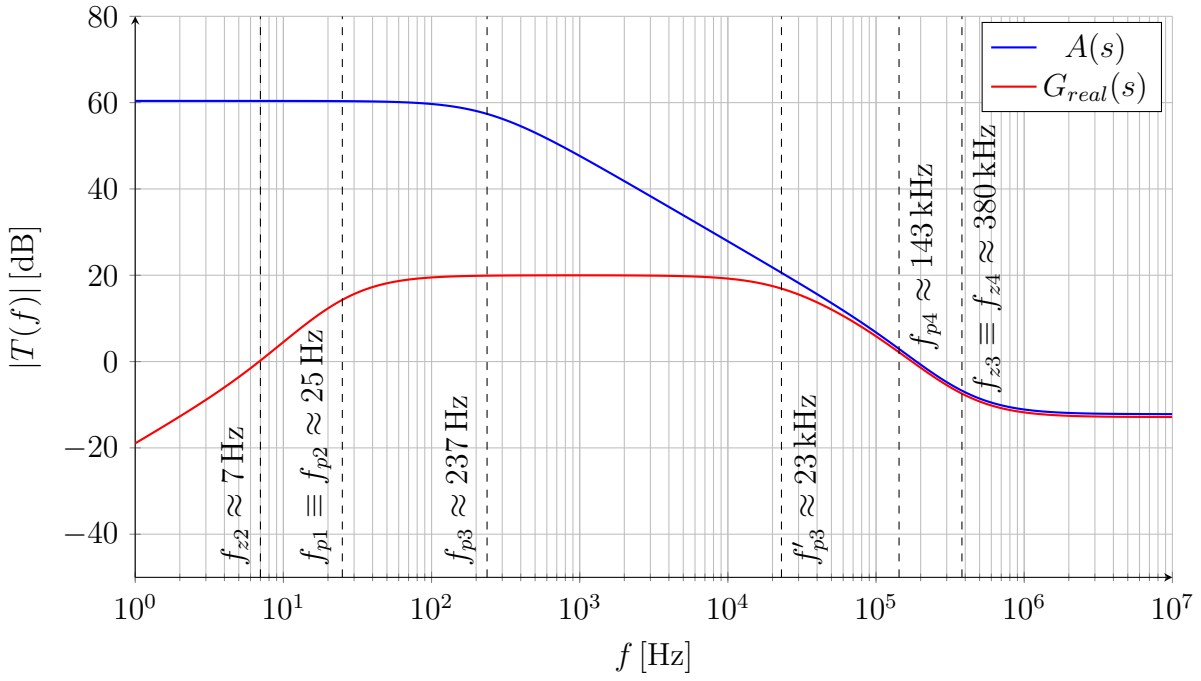

Figure 3.5: Bode plot of forward and closed loop gain of the LNA.

The simulation results of the real AC transfer function are in substantial agreement with the theoretical model, capturing the overall frequency behavior despite minor non-ideal deviations, as shown in Figure 3.6.

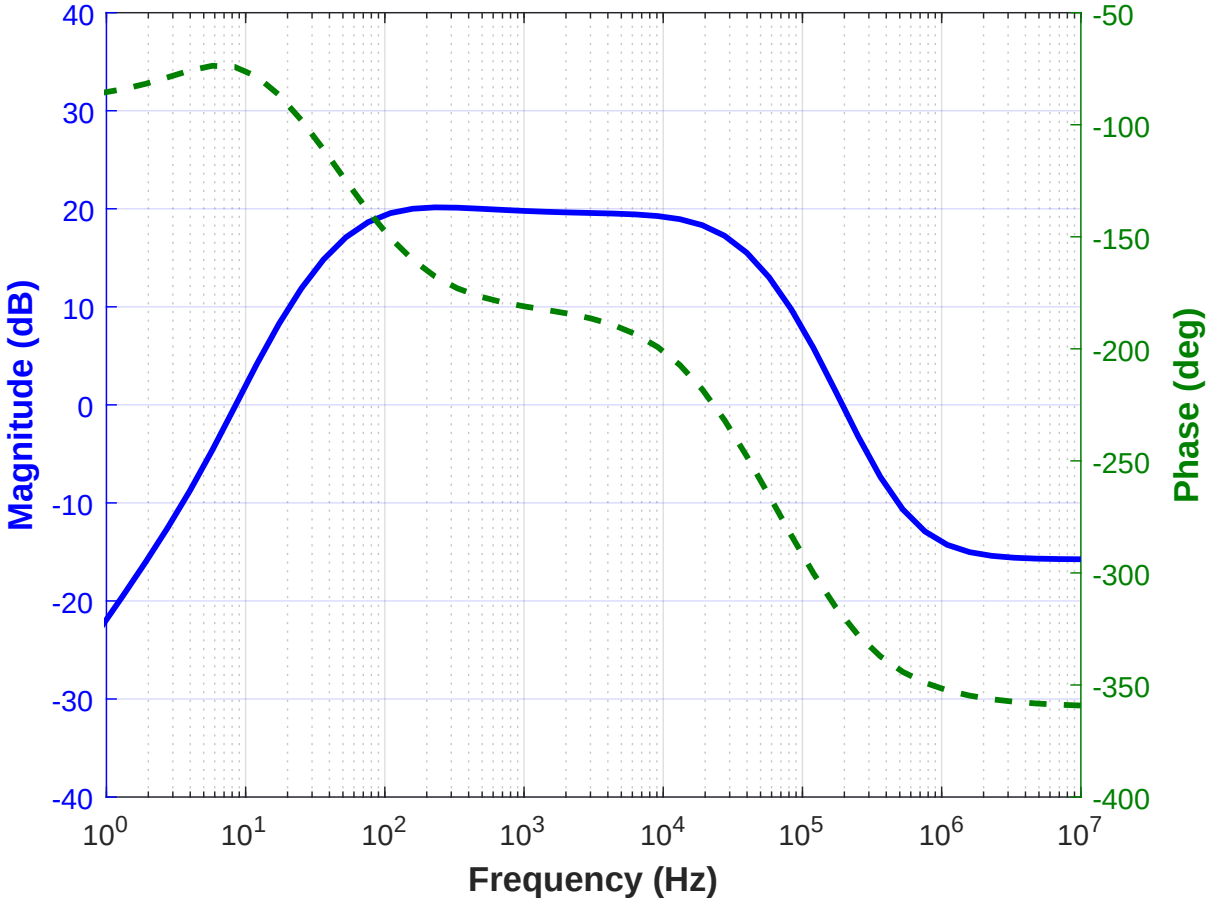

Figure 3.6: Simulated AC transfer function of the implemented LNA.

3.5. Performance Analysis

3.5.1. Power Consumption

The implementation of the LNA focuses strictly on meeting the ultra-low-power requirements necessary for always-on acoustic edge devices. Biased in the weak inversion regime, the circuit exhibits a total current consumption of only 24 nA from the 1.8 V supply rail. This yields an overall power dissipation of 44 nW. Within this power profile, the first inverter-based amplifying stage and the second source-follower stage account for the vast majority of the current budget, while the high-value pseudoresistor R_m contributes a negligible power overhead, consuming a resting current of approximately 30 pA. This exceptionally low power figure demonstrates the viability of the chosen 180 nm subthreshold design for energy-constrained processing chains.

3.5.2. Noise

A key figure of merit for the acoustic front-end is its capability to amplify micro-volt level signals coming from the microphone without introducing excessive thermal or flicker noise components. Through integrated circuit simulations, the equivalent input-referred RMS noise of the LNA was evaluated to be 43 μV, integrated across the target acoustic bandwidth. This performance is primarily achieved by leveraging the inverter-based topology of the first stage, which doubles the operational transconductance for a given bias current, and by optimizing the subthreshold pseudo-resistor R_m to minimize its electronic noise contribution well below that of a physical resistor of equivalent value.

3.5.3. Process and Mismatch Variability

To assess the robustness of the LNA against process variations and intra-die mismatch, a Monte Carlo simulation was performed across $N = 200$ runs. The simulated statistical variation of the closed-loop voltage gain response as a function of frequency is illustrated in Figure 3.7.

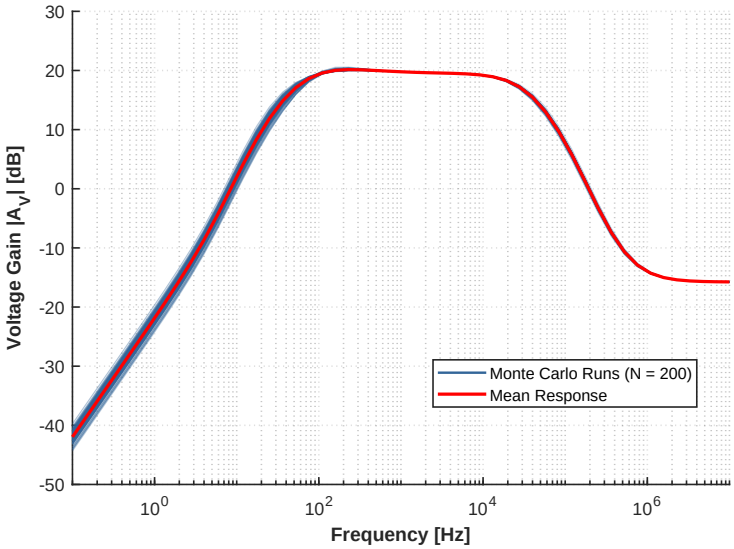

Figure 3.7: Monte Carlo simulation results ($N = 200$ runs) of the LNA closed-loop voltage gain alongside the calculated mean response, showcasing the impact of process and mismatch variability.

The simulation tracks the voltage gain $|A_V|$ from 10^{-1} Hz to over 10^7 Hz. The nominal in-band gain settles firmly at 20 dB, satisfying the $G = -C_1/C_2 = 10$ ratio requirement defined during the ideal analysis. The family of curves highlights that the mid-band gain

and the high-frequency cutoff ($f'_{p3} \approx 23\,\text{kHz}$) display negligible dispersion, indicating that the capacitive feedback network provides excellent stability against parameter variations. Conversely, a visible spread can be observed at the low-frequency high-pass cutoff region (1 Hz to 50 Hz). This dispersion is directly tied to the exponential sensitivity of the subthreshold transistors implementing the pseudoresistor R_m and the auxiliary OTA A_1, whose transconductance values fluctuate with process and mismatch variations. Nevertheless, the mean response (indicated by the red solid line) guarantees that the coincident pole frequency stays centered around 25 Hz, effectively suppressing DC offsets while keeping the entire audio spectrum unattenuated.

4 | Band-Pass Filterbank

4.1. Topology Introduction and Requirements

The structure adopted for each band-pass filter of the filterbank, shown in Figure 4.1, is the same as the one proposed by [12], which is based on a flipped voltage follower like the one described in [3].

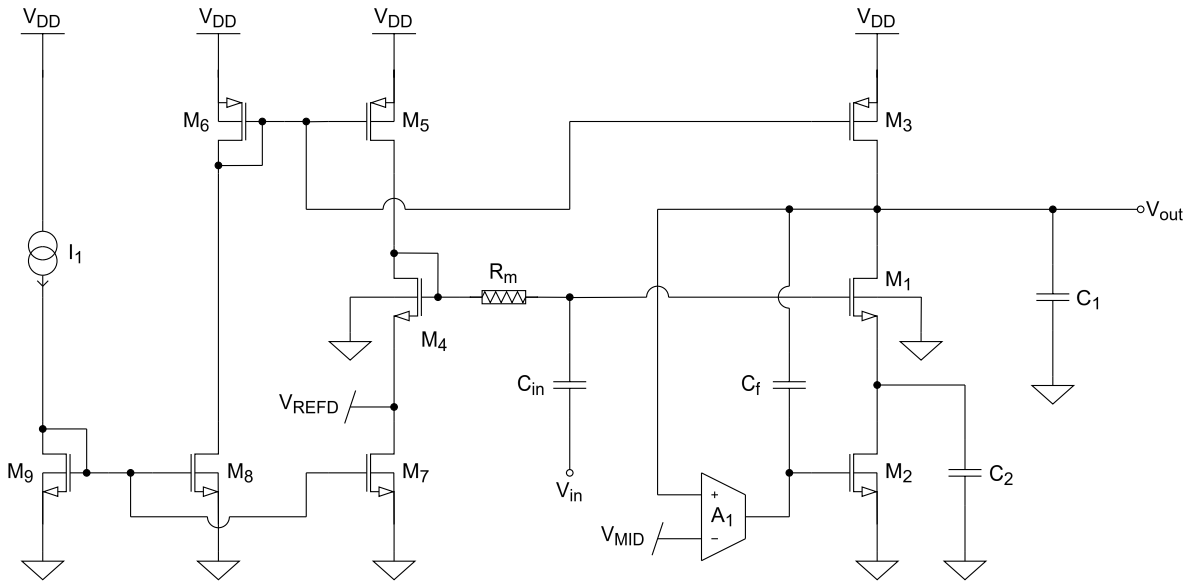

Figure 4.1: Schematic of a single BPF cell.

Transistor M_1 acts as a source-follower with an active load provided by transistor M_2, while the capacitor C_f implements a negative feedback loop that links capacitors C_1 and C_2, so that their interaction leads to a complex-conjugate pole pair, yielding the typical band-pass filter frequency response.

The circuit behavior in small-signal analysis is explained in further detail in Sections 4.1.2, 4.1.3 and 4.1.4.

This topology was considered a good choice as elementary cell of the filterbank because, once capacitors C_1 and C_2 are chosen, the resonant frequency of the filter can be adjusted tuning the current I_1, so that no changes to the physical structure are needed and, fur-

thermore, without impacting the chosen gain and quality factor.
Another reason was that, while a higher resonant frequency inherently increases power consumption, the required current I_1 remains remarkably low, reaching only $\sim 500\,\text{pA}$ for the highest-frequency band, implemented at 6 kHz.

4.1.1. DC Bias Network and AC Coupling

The bias network on the left, composed of transistors M_{4-9}, is adopted to bias transistors M_1 and M_3 with the same current I_1.
Additionally, the current of M_4 is imposed both by M_5 and M_7, so that applying a voltage V_{REFD} at the source of M_4 that same voltage will be the bias point of the source of M_1, being both its current and gate voltage fixed at the same values of M_4.
Fixing the source of M_4 to V_{REFD} means fixing the drain of M_2 too and, choosing $V_{REFD} = 150\text{mV}$, it is guaranteed that transistor M_2 will be in saturation region.
Note that the presence of M_7 is justified by the fact that V_{REFD} is a low-power voltage reference, with limited current sinking capabilities: M_7 is therefore used to subtract the bias current coming from the source of M_4, preventing this excess current from entering the V_{REFD} node and undermining its regulation stability.
The amplifier A_1 (a simple single stage OTA) implements a negative feedback loop, together with M_1 and M_2, that sets the output node DC bias at V_{MID}.
To decouple the input from the DC bias of the gate of M_1, the capacitor C_{in} implements an RC high-pass filter, that introduces a zero in the origin and a pole at the cut-off frequency.
To get an AC cut-off frequency sufficiently lower than the first band of the filterbank, placed at 100 Hz, a pseudoresistor R_m had to be introduced so that C_{in} could be kept small, to minimize silicon area occupation. The final values for those components were chosen to be 50 GΩ and 150 fF respectively, so that:

$$f_{p,ac} \simeq \frac{1}{2\pi\, R_m\, C_{in}} \approx 20\,\text{Hz}\,. \tag{4.1}$$

4.1.2. Low-Frequency Loop Analysis

The first step to understand the AC behavior of the circuit is to find the frequency at which the loop closed by A_1 stops keeping the output node at virtual ground, so that the signal can propagate to the output.
We might call this loop gain "low-frequency loop gain", to distinguish it from the "high-frequency loop gain" of the loop formed by M_1, M_2 and C_f, which will be analyzed in

Section 4.1.3. The loop starts to become ineffective when its gain is less than unitary, so we need to find the condition for which $G_{loop,LF} = 1$, which will lead to a pole in the transfer function of the circuit.
To compute the loop gain, a simplified version of the circuit can be considered, shown in Figure 4.2, where the bias network is neglected (except for A_1) together with C_1 and C_2, which become relevant at much higher frequencies.

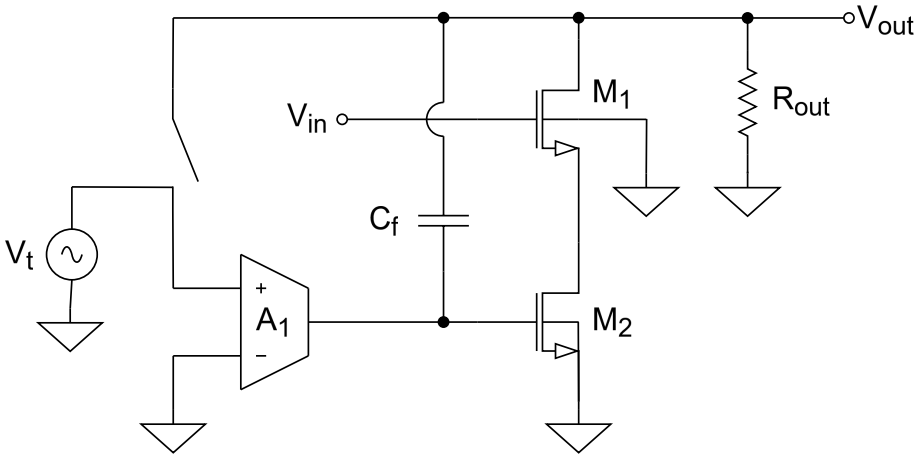

Figure 4.2: Simplified schematic of the BPF cell used to compute the low-frequency loop gain.

The loop can be cut at the positive input of A_1, where a test voltage v_t is then applied. Writing the Kirchhoff current law at the gate node of M_2 and at the output node, we get:

$$\begin{cases} i_{Cf} = g_{m,A_1}\, v_t \\ i_{Cf} = \dfrac{v_{out}}{R_{out}} + g_{m2}\, v_x \\ i_{Cf} = (v_x - v_{out})\, sC_f \end{cases} \tag{4.2}$$

from which we can derive:

$$G_{loop,LF} = -\frac{g_{m,A_1}\, R_{out}}{1 + g_{m2}\, R_{out}} \frac{1 - sC_f \,/\, g_{m2}}{sC_f \,/\, g_{m2}}\,. \tag{4.3}$$

The loop gain transfer function exhibits a pole in the origin and a positive zero placed at:

$$f_{z,loop} = \frac{g_{m,2}}{2\pi\, C_f} \tag{4.4}$$

with a high-frequency gain given by:

$$G_{loop,LF}\,(s \to \infty) = \frac{g_{m,A_1}\,R_{out}}{1 + g_{m,2}\,R_{out}}\,. \tag{4.5}$$

Note that, because of the phase reduction due to the positive zero, the high-frequency gain is required to be less than unitary to guarantee the stability of the system, which means:

$$\frac{g_{m,A_1}\,R_{out}}{1 + g_{m,2}\,R_{out}} \simeq \frac{g_{m,A_1}}{g_{m,2}} < 1 \;\Rightarrow\; g_{m,A_1} < g_{m,2} \tag{4.6}$$

where the approximation $g_{m,2}\,R_{out} \gg 1$ was considered.
The direct implication is that the bias current of the differential pair of A_1 needs therefore to be lower than the bias current of M_2, which is anyway an easy requirement.
We can now compute the frequency at which the loop gain is unitary, that is the frequency of the pole of the real gain $v_{out}\,/\,v_{in}$ of the BPF cell after which C_f can be considered as a short circuit:

$$G_{loop,LF} = 1 \;\Rightarrow\; p_1 = \frac{g_{m2}}{C_f}\,\frac{g_{m,A_1}\,R_{out}}{1 + g_{m,2}\,R_{out}} \simeq \frac{g_{m,A_1}}{C_f} \;\Rightarrow\; f_{p,fb} = \frac{g_{m,A_1}}{2\pi\,C_f}\,. \tag{4.7}$$

The frequency requirement for this pole is to be much lower than the first band of the filterbank, so that C_f can be considered a short circuit and the feedback loop between M_1 and M_2 is closed. To keep the value of C_f reasonably low to limit silicon area occupation, while meeting the requirement above, the transconductance g_{m,A_1} has to be chosen considerably low.
In this work, to get $f_{p,fb} \approx 10\,\text{Hz}$ the bias current of A_1 was chosen to be $\sim 1\,\text{pA}$, yielding:

$$\begin{cases} g_{m,A_1} \approx 16\,\text{pS} \\ C_f = \dfrac{g_{m,A_1}}{2\pi\,f_{p,fb}} \approx 250\,\text{fF}\,. \end{cases} \tag{4.8}$$

4.1.3. Band-Pass Transfer Function Analysis

To analyze the frequency response of the circuit when C_f can be considered a short circuit and the input can propagate to the output, which means above $f_{p,fb}$, the simplified version shown in Figure 4.3 can be considered, where the bias network is neglected together with A_1, as its loop is not strong enough to move the output node anymore. As a further simplification, the AC coupling capacitance C_{in} is neglected in this analysis.

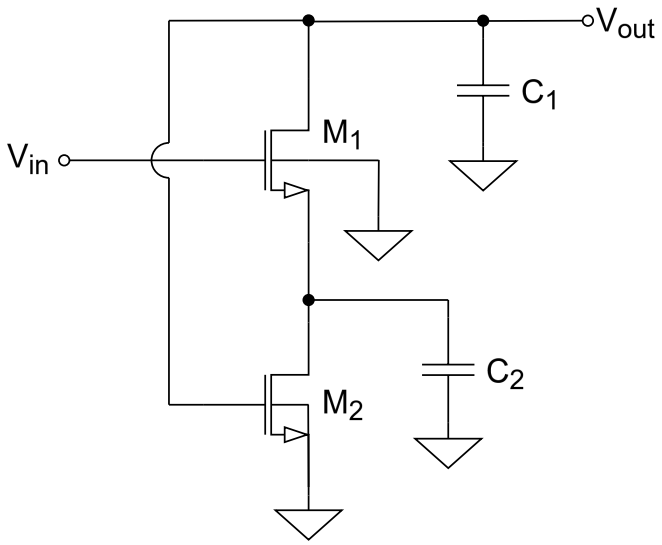

Figure 4.3: Simplified schematic for the band-pass transfer function analysis of the elementary BPF cell.

To compute the real transfer function of the circuit, we need to first compute the forward gain $A(s)$ from input to output and the gain of the high-frequency loop, $G_{loop,HF}(s)$.
To evaluate the forward gain we might neglect the current of M_2, as it is the feedback signal, and consider M_1 as a source-follower with C_2 as a load; the current of M_1 is then integrated into C_1.
The output voltage results:

$$v_{out}(s) = v_{in}(s)\left[\frac{1}{g_{m1}} + \frac{1}{sC_2}\right]^{-1}\frac{1}{sC_1} \tag{4.9}$$

that yields, after some calculations:

$$A(s) = -\frac{C_2}{C_1}\frac{1}{1 + sC_2\,/\,g_{m1}}\,. \tag{4.10}$$

The loop gain can be found by cutting the loop at the gate of M_2: a test voltage applied to the gate generates a current flowing into M_2 that is partitioned between M_1 and C_2; the fraction of current flowing into M_1 is then integrated into C_1.
The resulting loop gain is therefore:

$$G_{loop,HF}(s) = -\,g_{m2}\cdot\frac{g_{m1}}{g_{m1} + sC_2}\cdot\frac{1}{sC_1} = -\,\frac{g_{m1}}{sC_1}\,\frac{g_{m2}}{g_{m1} + sC_2}\,. \tag{4.11}$$

The real gain can now be evaluated:

$$
\begin{aligned}
G_{real}(s) = \frac{A(s)}{1 - G_{loop,HF}(s)} &= -\frac{C_2}{C_1}\,\frac{1}{1 + sC_2\,/\,g_{m1}}\left[1 + \frac{g_{m1}}{sC_1}\,\frac{g_{m2}}{g_{m1} + sC_2}\right]^{-1} = \\
&= -\frac{C_2}{C_1}\,\frac{1}{1 + sC_2\,/\,g_{m1}}\,\frac{sC_1}{g_{m2}}\,\frac{1 + sC_2\,g_{m1}}{s^2\,C_1\,C_2\,/\,(g_{m1}\,g_{m2}) + s\,C_1\,/\,g_{m2} + 1} = \\
&= -\frac{sC_2\,/\,g_{m2}}{s^2\,C_1\,C_2\,/\,(g_{m1}\,g_{m2}) + s\,C_1\,/\,g_{m2} + 1}\,.
\end{aligned}
\tag{4.12}
$$

As shown by the formula above, the real transfer function exhibits a zero in the origin and a pole pair which, if the circuit is dimensioned properly to make it complex-conjugate, results in a peaking bell-like shape typical of band-pass filters.
We know that the equation for two complex conjugate poles in canonic form is:

$$
\frac{s^2}{{\omega_0}^2} + \frac{s}{Q\,\omega_0} + 1 = 0\,. \tag{4.13}
$$

Specifically, the roots of this characteristic equation become complex-conjugate if and only if the quality factor satisfies the condition $Q > 0.5$, whereas for $Q = 0.5$ they lie on the real axis as coincident poles.
The canonic polynomial above can be recognized at the denominator of equation (4.12).

$$
\begin{cases}
f_0 = \dfrac{\omega_0}{2\pi} = \dfrac{1}{2\pi}\,\sqrt{\dfrac{g_{m1}\,g_{m2}}{C_1\,C_2}} \\
Q = \sqrt{\dfrac{g_{m1}\,C_2}{g_{m2}\,C_1}} \\
G = G_{real}(j\omega_0) = -\dfrac{C_2}{C_1}
\end{cases}
\tag{4.14}
$$

where f_0 is the resonant frequency of the filter, Q is the quality factor and G is the gain at resonance.
Considering that in weak inversion the transconductance g_m depends only on the bias current and the technology, being the current flowing into M_1 and M_2 the same we end up with:

$$
G_{real}(s) = -\frac{sC_2\,/\,g_m}{s^2\,C_1\,C_2\,/\,{g_m}^2 + s\,C_1\,/\,g_m + 1} \tag{4.15}
$$

and:

$$\begin{cases} f_0 = \dfrac{1}{2\pi}\sqrt{\dfrac{{g_m}^2}{C_1 C_2}} \\ Q = \sqrt{\dfrac{C_2}{C_1}} \\ G = -\dfrac{C_2}{C_1} \end{cases} \tag{4.16}$$

To get the desired values of $G = 4 = 12\,\text{dB}$ and $Q = 2$, C_1 and C_2 were chosen to be respectively 250 fF and 1 pF. The resonant frequency of the filter can then be adjusted by changing the bias current I_1.

4.1.4. Overall Transfer Function Analysis

Before looking at the final shape of the real transfer function of the BPF cell, an additional consideration has to be made: the overall transfer function exhibits two zeros at the origin and four poles (two being complex conjugates), presenting an inconsistency with the characteristic bell shape of a band-pass filter.
An explanation for this inconsistency can be found in a deeper analysis that considers the gate-to-drain capacitance of transistor M_2, which can be approximated to a capacitor between gate and ground amplified by the self-gain μ_{M_2} of M_2, because of the Miller effect, as shown in Figure 4.4.

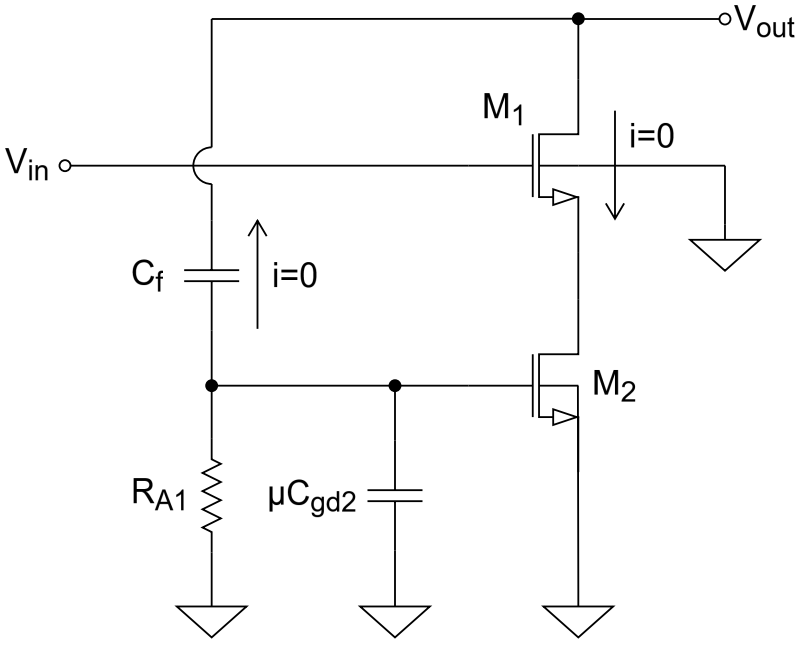

Figure 4.4: Simplified schematic of the BPF cell used to compute the feed-forward zero.

The equivalent resistor seen by such capacitor is simply given by the output resistance of the amplifier A_1, R_{A_1}, because no current can flow into C_f as the gate-to-source voltage of M_1 can be considered zero, being it a follower.
The resulting feed-forward zero is a positive (right-half-plane) zero and can be approximated as:

$$f_{z,ff} = \frac{1}{2\pi\,\mu_2\,C_{gd2}\,R_{A_1}}\,. \tag{4.17}$$

In our implementation C_{gd2} is in the order of hundreds of aF, μ_{M_2} is on the order of hundreds and $R_{A_1} \approx 500\,\mathrm{T\Omega}$ (given the previously discussed bias of $\sim 1\,\mathrm{pA}$ for A_1), which puts $f_{z,ff}$ at a very low frequency, around 1 mHz.
Combining all the obtained results, the final form of the transfer function of the BPF cell is given by:

$$H(s) = -\,\frac{sC_{in}\,R_{mir}}{1+sC_{in}\,R_{mir}}\,\frac{1+sC_{gd2}\,\mu_2\,R_{A_1}}{1+sC_f\,/\,g_{m,A_1}}\,\frac{sC_2\,/\,g_{m2}}{s^2\,C_1\,C_2\,/\,(g_{m1}\,g_{m2})+s\,C_1\,/\,g_{m2}+1}\,. \tag{4.18}$$

Figure 4.5 shows the theoretical Bode plot of the frequency response of the BPF cell centered at 100 Hz, while Figure 4.6 shows the real simulated results.
Note that for better readability both plots start from 1 Hz and therefore the feedforward zero is not visible.

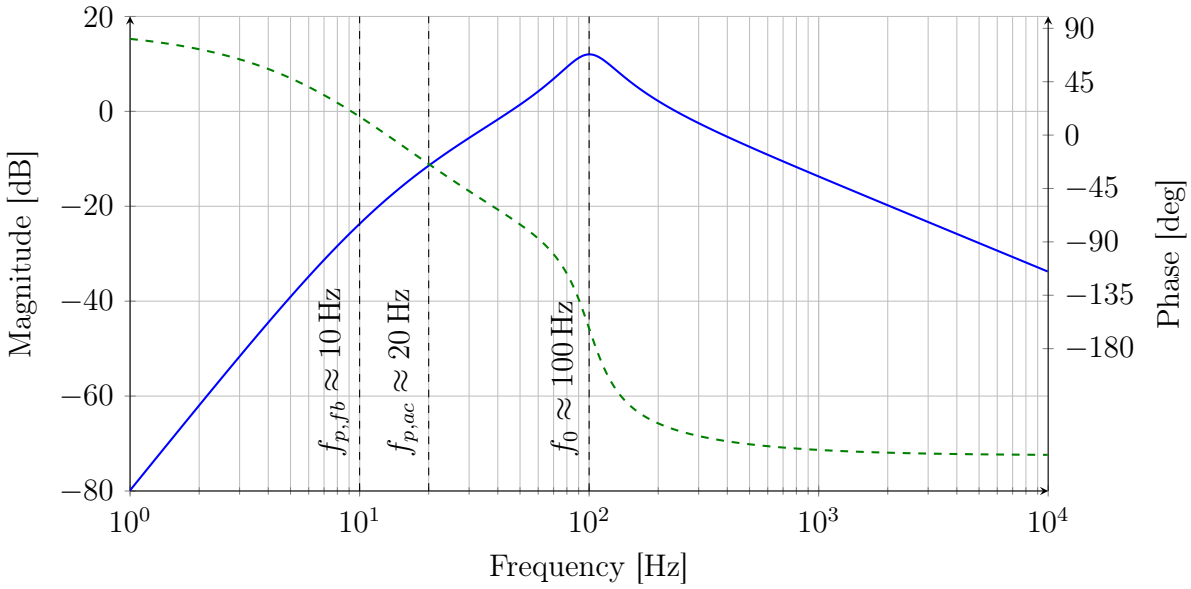

Figure 4.5: Bode plot (Magnitude and Phase) of the transfer function of the BPF cell centered at 100 Hz.

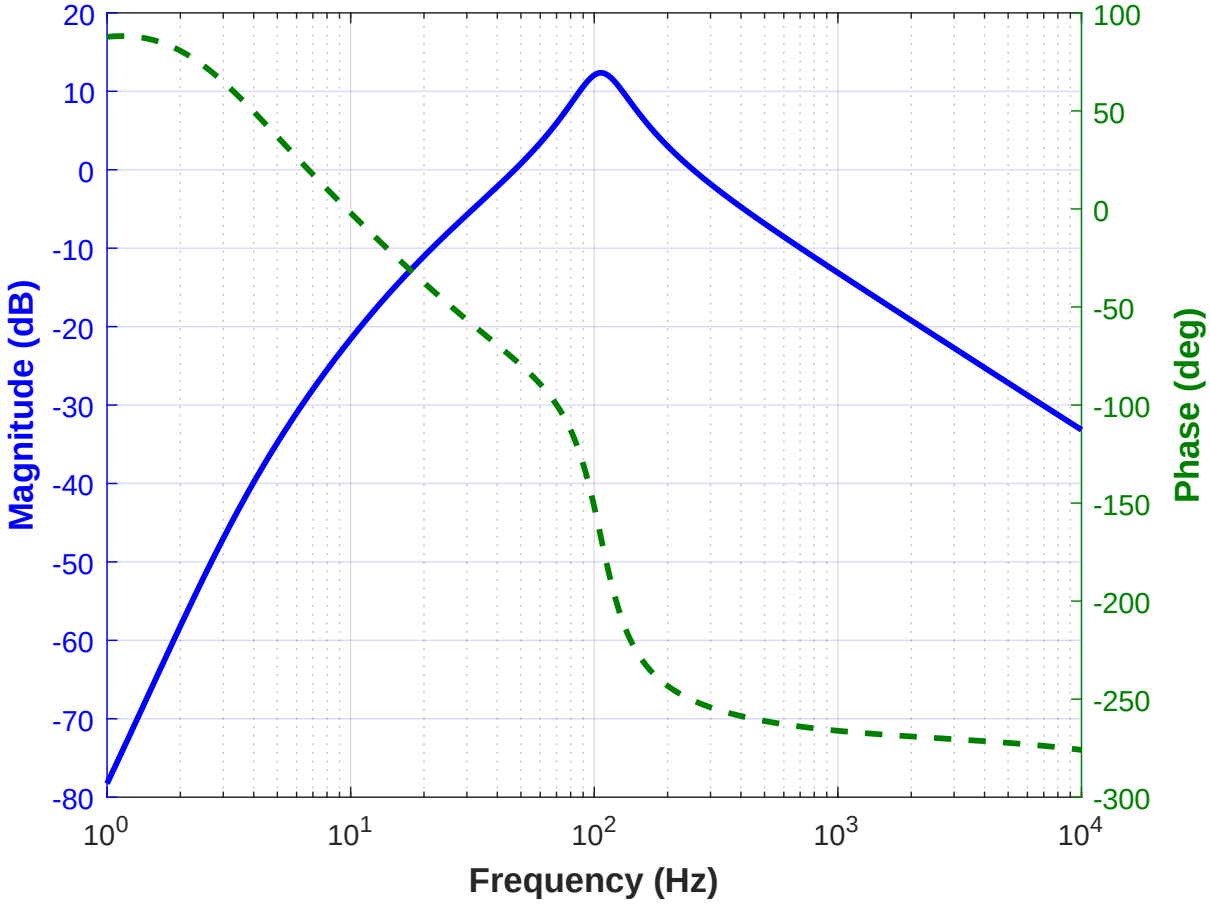

Figure 4.6: Actual simulated transfer function of the BPF cell centered at 100 Hz.

In Table 4.1 all BPF frequency bands are reported together with their required bias current I_1.

f_0 [Hz]	100	130	170	225	300	390	510	670
I_1 [pA]	10	12	15	20	26	34	45	58
f_0 [Hz]	880	1160	1520	2000	2620	3440	4500	6000
I_1 [pA]	77	100	133	175	230	300	400	540

Table 4.1: BPF filterbank frequency bands and corresponding bias current I_1.

The overall simulated frequency response of the 16-channel BPF filterbank is illustrated in Figure 4.7.

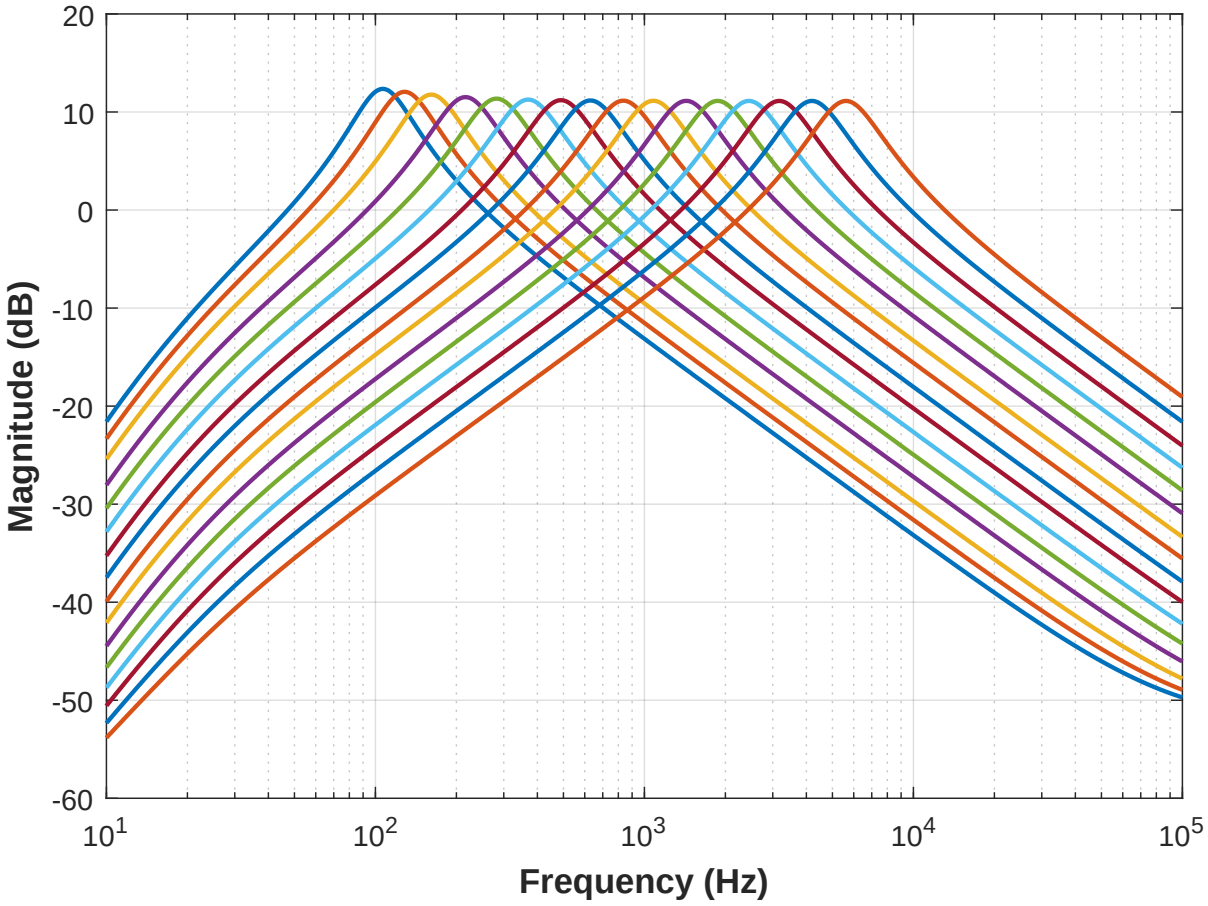

Figure 4.7: Overall simulated frequency response of the 16-channel BPF filterbank.

4.2. Performance Analysis

4.2.1. Power Consumption

The 16-channel band-pass filterbank has been engineered to implement sharp spectral decomposition with minimal energy overhead. The total current consumption across all 16 core BPF channels stands at 9 nA under nominal biasing conditions. To evaluate the true power dissipation profile of the full block, the contribution of the high-value input pseudoresistors must be added to the core budget. Each of the 16 channels includes an AC-coupling pseudoresistor R_m that draws a static current of approximately 25 pA.
The total operational current for the filterbank is therefore calculated as:

$$I_{tot,BPF} = 9\,\text{nA} + (16 \times 25\,\text{pA}) = 9.4\,\text{nA} \tag{4.19}$$

Operating from the 1.8 V supply voltage rail, the total power consumption of the 16-channel BPF filterbank amounts to exactly 16.92 nW. This remarkably low nanowatt dissipation, combined with the structural flexibility offered by tuning the subthreshold currents I_1 (ranging from 10 pA at 100 Hz up to 540 pA at 6 kHz as detailed in Table 4.1), demonstrates a highly efficient architecture suitable for always-on edge computing devices.

4.2.2. Process and Mismatch Variability

To verify the manufacturing robustness and stability of the filterbank against process variations and intra-die mismatch, an extensive Monte Carlo simulation was performed. The statistical analysis focuses on the target frequency band centered at $f_0 = 1160$ Hz, which is biased with a nominal current $I_1 = 100$ pA. The resulting probability density functions and corresponding Gaussian fits for the resonant frequency (f_0), the peak mid-band gain (G), and the quality factor (Q) are illustrated in Figure 4.8, Figure 4.9, and Figure 4.10, respectively.
The statistical distribution of the resonant frequency f_0, obtained over $N = 2904$ simulation runs, is shown in Figure 4.8.

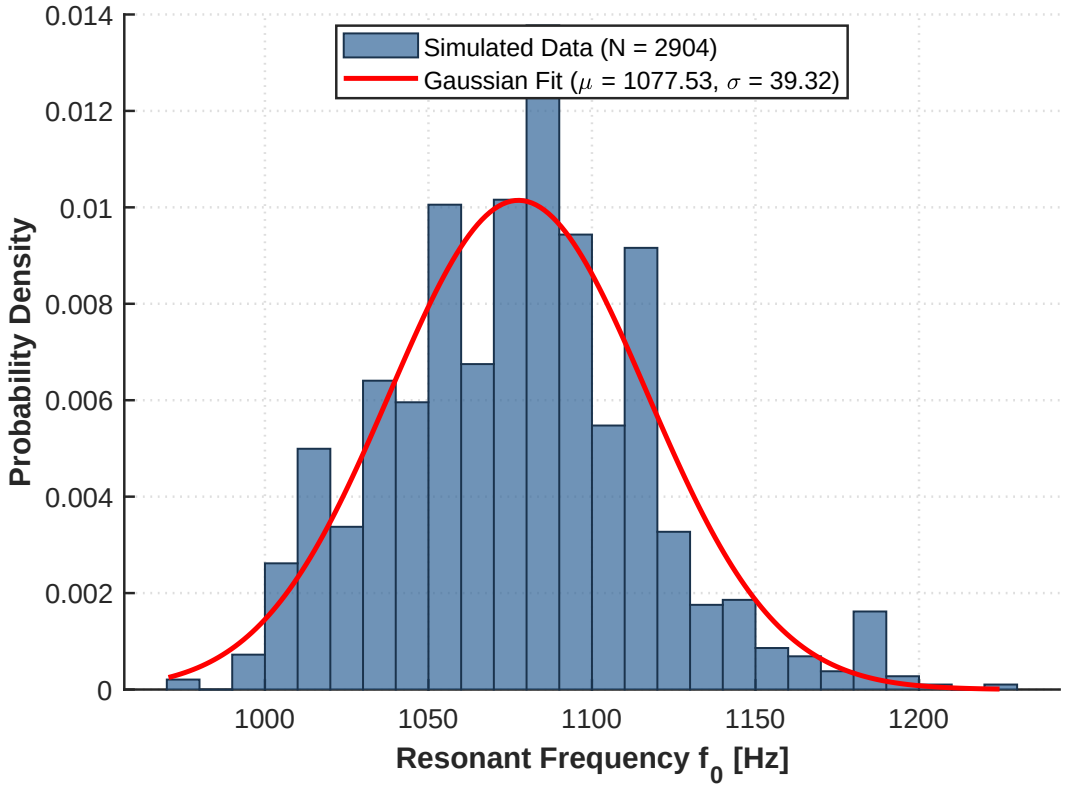

Figure 4.8: Monte Carlo statistical distribution ($N = 2904$ runs) and Gaussian fit of the resonant frequency f_0 for the channel centered at 1160 Hz.

The extracted Gaussian fit features a mean value $\mu = 1077.53$ Hz and a standard deviation $\sigma = 39.32$ Hz. The minor down-shift of the mean frequency with respect to the nominal target (1160 Hz) is a known phenomenon caused by systematic parasitic capacitances at the internal nodes of the flipped voltage follower cell, which slightly augment the active capacitance value. The controlled standard deviation demonstrates that the filter channels maintain a distinct spectral segregation, preventing adjacent band overlap even in the presence of severe mismatch variations.

Figure 4.9 reports the probability density function for the closed-loop peak gain G at resonance over $N = 3617$ simulation runs.

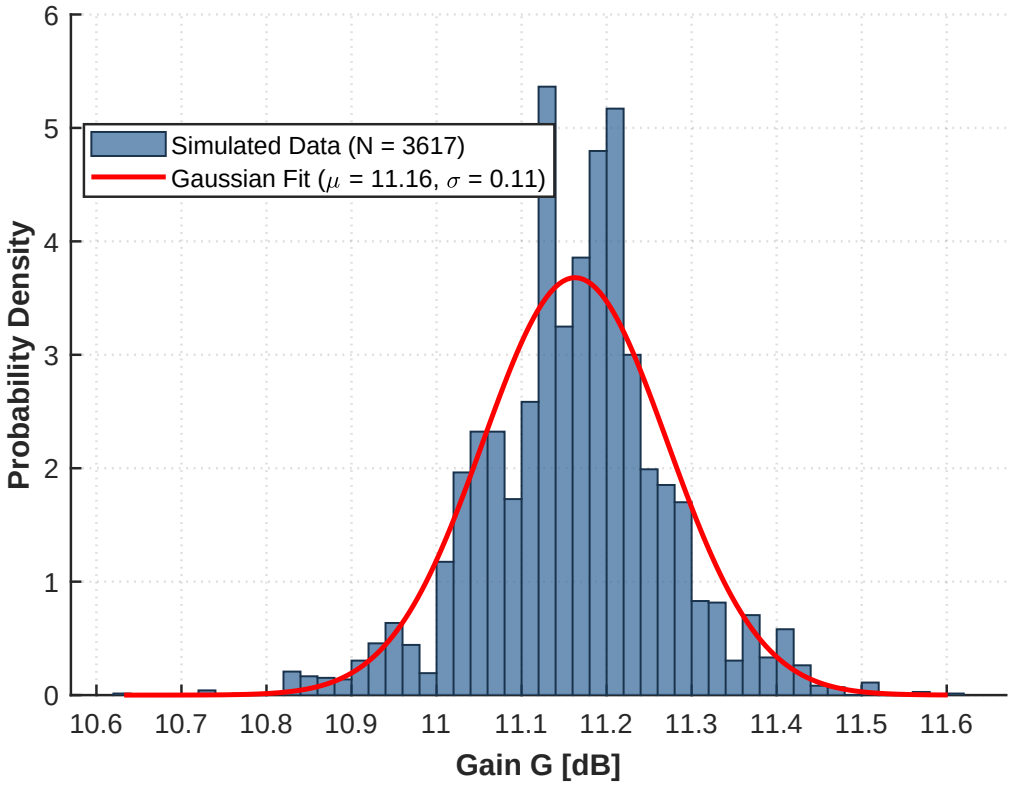

Figure 4.9: Monte Carlo statistical distribution ($N = 3617$ runs) and Gaussian fit of the resonant peak gain G for the channel centered at 1160 Hz.

The distribution exhibits a mean peak gain $\mu = 11.16\,\text{dB}$ and a very tight standard deviation $\sigma = 0.11\,\text{dB}$. This exceptionally narrow spread indicates that the voltage gain is highly decoupled from subthreshold transistor parameters, as it is strictly dictated by the geometric matching of the integrated capacitive partition ratio ($G = -C_2/C_1 \approx 4 \equiv 12\,\text{dB}$). The minor attenuation of the mean relative to the ideal 12 dB target is attributed to the finite open-loop transconductance and body-effect loading of M_1.
Finally, the robustness of the filter selectivity is confirmed by the quality factor Q distribution shown in Figure 4.10, simulated over $N = 3491$ runs.

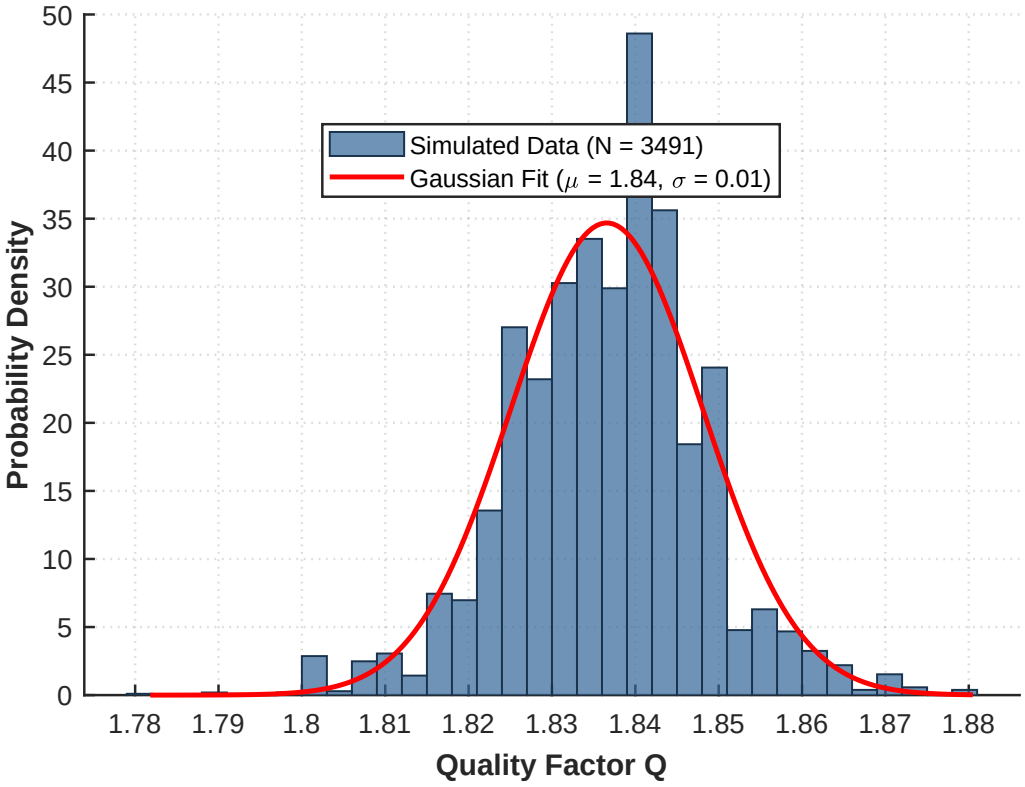

Figure 4.10: Monte Carlo statistical distribution ($N = 3491$ runs) and Gaussian fit of the filter quality factor Q for the channel centered at 1160 Hz.

The calculated Gaussian fit yields a mean selectivity $\mu = 1.84$ with a standard deviation $\sigma = 0.01$. The minimal deviation confirms that the filter sharpness remains stable across all fabrication corners, validating the reliability of this subthreshold topology for stable, robust continuous-time neuromorphic feature extraction.

5 | Voltage-to-Current Converter

5.1. Topology Introduction and Requirements

To rectify and convert from voltage to current the filtered audio signal different approaches were explored, before reaching a circuit prototype that roughly satisfied the necessary requirements.
Note that the half-wave rectifier proposed in [12] was discarded as it represents a simplistic approach based on a single, voltage-controlled MOSFET. Since the device is biased in the sub-threshold region, this configuration exposes the circuit to unacceptable exponential variations with temperature.
The implemented conversion circuit is required to work in the whole frequency range covered by the filterbank, from 100 Hz to 6 kHz, and to translate the input voltage into a positive current exiting the output, which will be integrated into the membrane capacitance of the neuron.
To prevent the neuron from firing when no input signal is present, the DC output current needs to be as close as possible to zero, even considering non-idealities and process/mismatch variations.

5.1.1. First Design

One of the topologies investigated was the one in Figure 5.1.

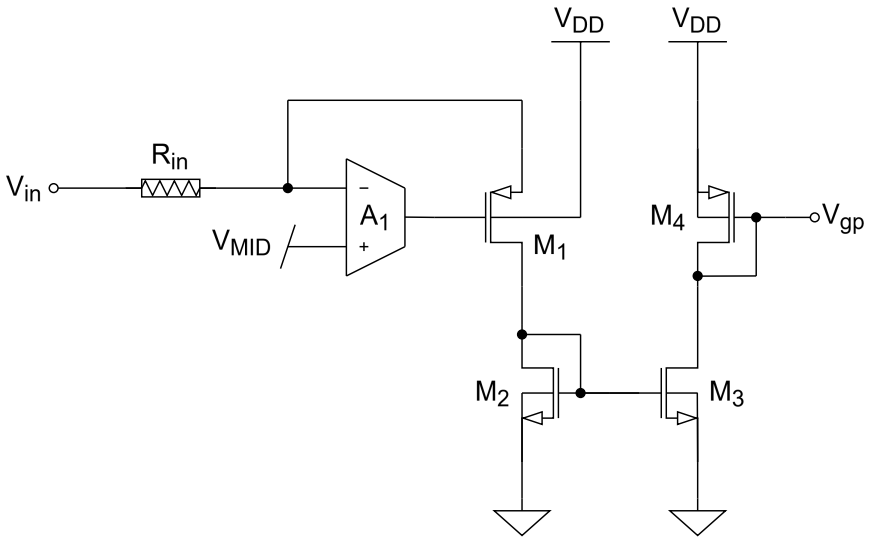

Figure 5.1: Schematic of the first voltage-to-current converter topology investigated.

The amplifier A_1 is connected in a negative feedback loop by means of the pMOS M_1, keeping its negative input at the reference voltage V_{MID}, which is also the DC output voltage of the BPF cell. The input voltage is applied to a pseudoresistor R_{in}, whose other terminal is connected to the virtual ground of A_1. Thanks to the feedback loop no DC current flows into M_1, which is off in absence of input signal: a negative input signal cannot therefore produce a current flowing into M_1, while any positive input voltage leads to an output current $i_{rec} = v_{in}/R_{in}$ flowing into M_1, which can be read and replicated by means of current mirrors.

This initial approach presented two significant drawbacks rendering it unsuited for the application.

First, the loop gain magnitude strongly depends on the input signal and its related current, that flowing into M_1 increases its polarization.

Second, because the minimum input voltage is comparable to the offset of A_1 (a few mV), a DC current of the same order as the minimum signal current is constantly present at the output, thereby degrading the performance of the circuit.

5.1.2. Second Design

An alternative architecture that was evaluated next is the one proposed in [13], reported in Figure 5.2.

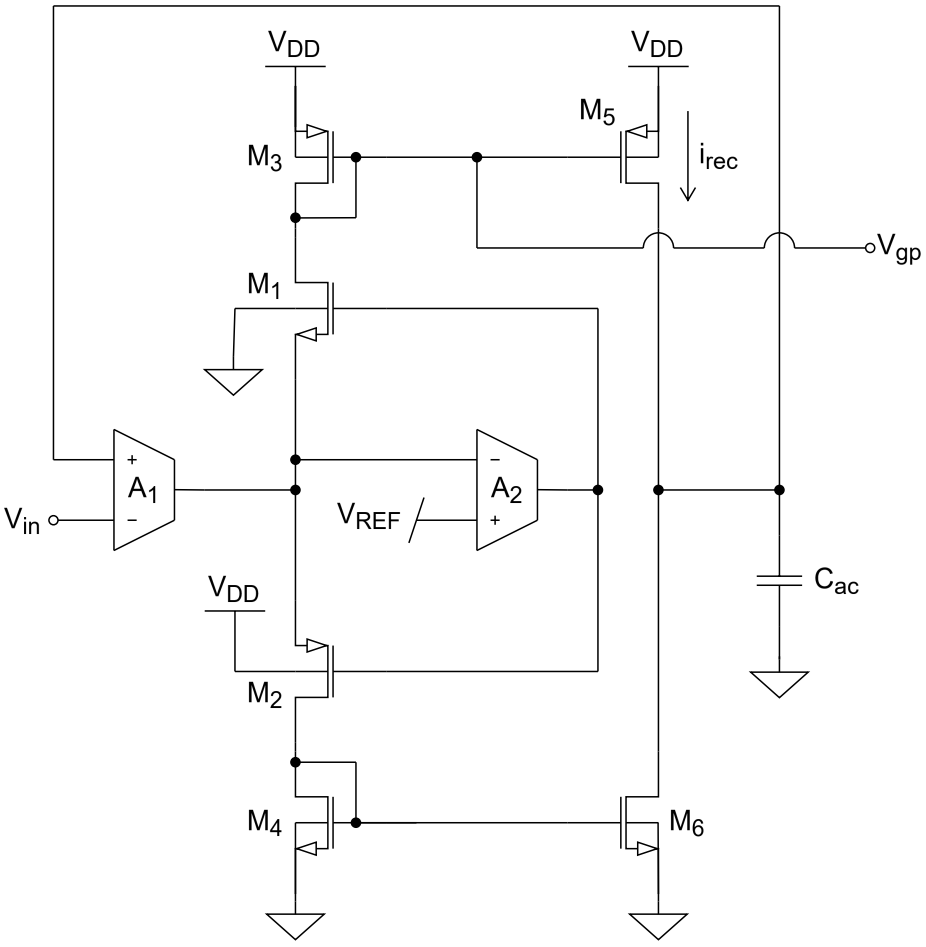

Figure 5.2: Schematic of the second voltage-to-current converter topology investigated.

In this topology the transconductor A_1 is closed in a negative feedback loop by the transistor network, so that the positive input of A_1 is kept equal to the negative one as long as the capacitor C_{ac} can be considered an open circuit.
The amplifier A_2 closes an internal loop that keeps the output voltage of A_1 to a reference voltage V_{REF}, by setting the gate voltage of M_1 and M_2 so that their DC drain current is zero.
For an input signal at higher frequency than the pole introduced by C_{ac}, given by

$$f_{ac} = \frac{1}{2\pi\, g_{m,A_1}\, C_{ac}}\,, \tag{5.1}$$

the capacitor C_{ac} can be considered a short circuit and the positive input of A_1 at ground, so that an output current $i_{A_1} = g_{m,A_1}\, v_{in}$ is generated and perturbates the potential of the virtual ground of A_2 that, as a consequence, switches on either M_1 or M_2 so that i_{A_1} flows through the upper or lower branch of the transistor network, reaching ground through C_{ac}.

Being that no DC current flows into the transistors, the end result is that the drain current of M_5, i_{rec}, is the half-wave rectification of i_{A_1}.
Full-wave rectification could be achieved by copying and flipping, by means of current mirrors, the drain current of M_6 and then adding the two currents together.
The main drawback of this topology is the direct link between the high-pass cut-off frequency f_{ac} and the generated signal current i_{A_1}, because both depend on the transconductance g_{m,A_1}. To keep the value of C_{ac} reasonable, for example 1 pF, the value of g_{m,A_1} required to get a cut-off frequency of 10 Hz, much lower than the first BPF band to avoid attenuation, is:

$$g_{m,A_1} = \frac{1}{2\pi\, C_{ac}\, f_{ac}} \approx 160\,\text{pS} \tag{5.2}$$

which would give, for a minimum input signal of a few mV, a current in the order of hundreds of fA, that is too low to provide reliable results.

5.2. Final Design Description

The final topology adopted for the voltage-to-current converter is the one depicted in Figure 5.3.
It was designed to address the limitations of the topology described in Section 5.1.2 and it proved to be an effective compromise in fulfilling the design specifications.
The fundamental concept of this architecture, discussed in greater detail below, relies on generating a signal current flowing through the pseudoresistor R_{in} and the differential amplifier A_2. The transistor network between nodes v_x and v_y performs a full-wave rectification of this current, which is subsequently injected into the neuron membrane capacitance via two transistors driven by the gate voltages v_{gp1} and v_{gp2}.
Furthermore, this topology decouples the high-pass cut-off frequency of the signal from the value of R_{in}.

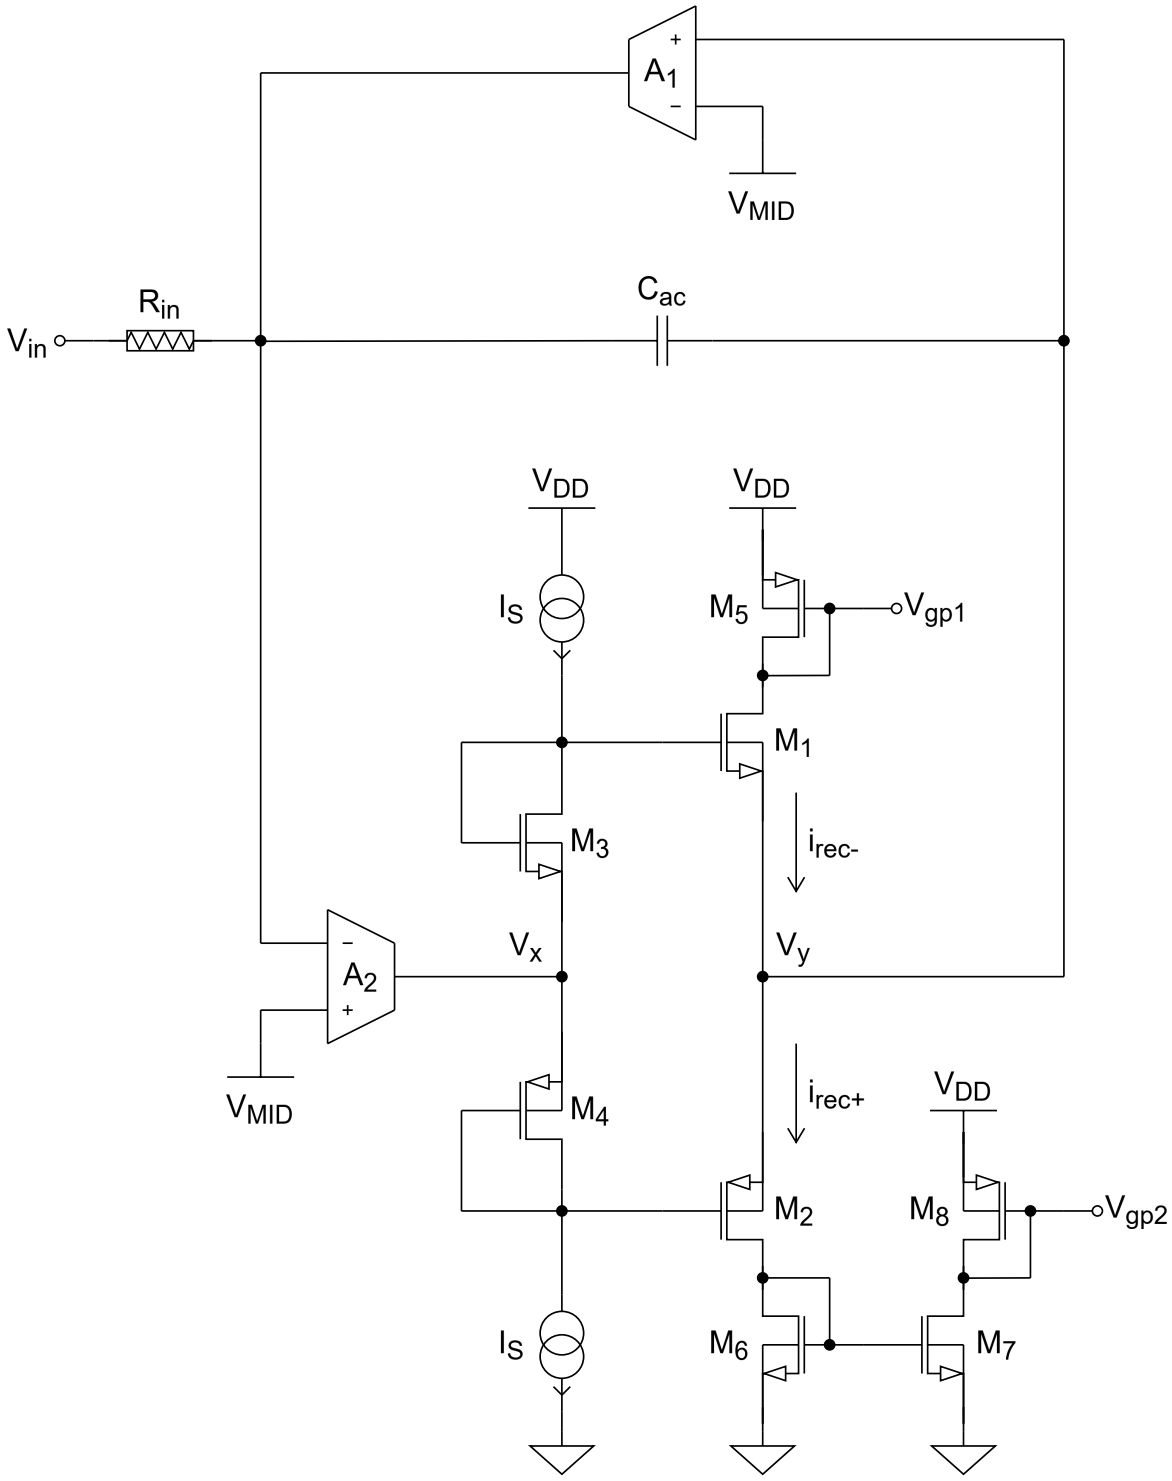

Figure 5.3: Schematic of the adopted voltage-to-current converter topology.

5.2.1. DC Analysis

For the DC analysis the network between nodes v_x and v_y can neglected, as if the two nodes were directly connected together, as shown in Figure 5.4. The reason will be later

explained in Section 5.2.3. The resulting unified node has been called v_{out}.

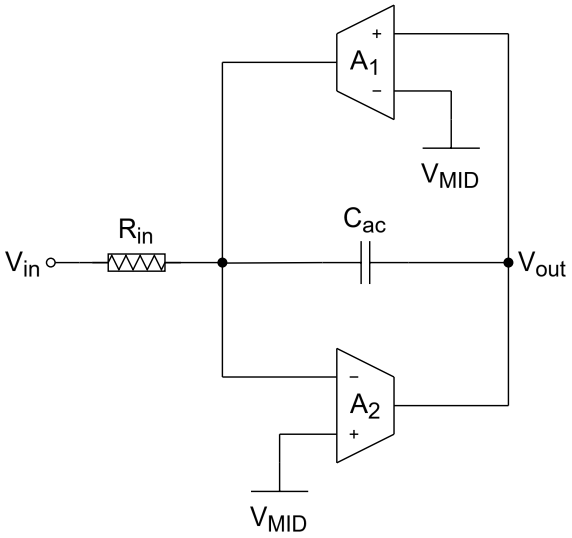

Figure 5.4: Simplified schematic of the proposed voltage-to-current converter, for DC analysis.

At DC, the capacitor C_{ac} can be seen as an open circuit and the amplifiers A_1 and A_2 are connected in a negative feedback loop, which keeps the output of A_1 at V_{MID}.
Being V_{MID} also the DC voltage of v_{in}, ideally A_1 doesn't have to provide any current and also the output node will be at V_{MID}. If for whatever reason the previous assumption isn't true, then A_1 will have to provide the current:

$$i_{A_1} = \frac{\Delta V}{R_{in}} \tag{5.3}$$

leading to an output voltage:

$$v_{out} = \frac{\Delta V}{R_{in}} \frac{1}{g_{m,A_1}} \,. \tag{5.4}$$

Being that R_{in} will be chosen lower than ${g_{m,A_1}}^{-1}$ for reasons that will become clear in Section 5.2.2, the result is that any mismatch in the DC voltage of v_{in} and V_{MID}, including the offset voltage of A_2, is amplified at the output node by a factor:

$$\alpha = \frac{1}{g_{m,A_1} R_{in}} > 1 \,. \tag{5.5}$$

Another detrimental effect of this mismatch is that it sets a lower bound on the bias current of A_1, which needs to be able to provide the required DC current flowing into R_{in}.

Because of both these limitations, it is of paramount importance to limit as much as possible the voltage offset of A_2, which leads to an inevitable increase in the size of its transistors and therefore of its input and output capacitance.

5.2.2. AC Coupling and Ideal Transfer Function

The AC analysis of the circuit can be performed under the same topological conditions of Section 5.2.1, reported again for convenience in Figure 5.5.

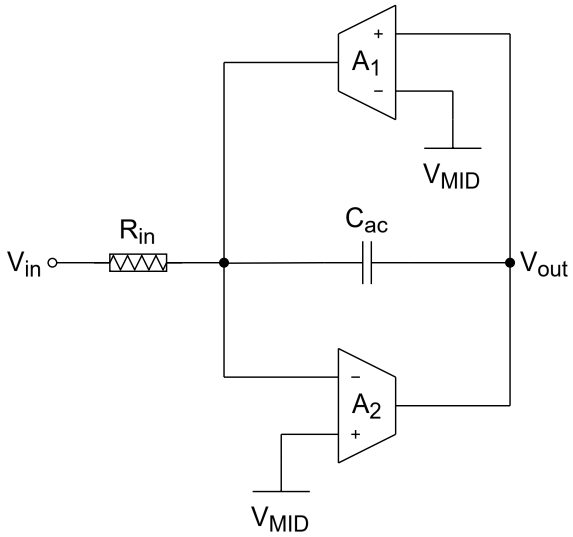

Figure 5.5: Simplified schematic of the proposed voltage-to-current converter, for AC analysis.

The input signal is applied to a pseudoresistor R_{in} whose other terminal is connected to the virtual ground of A_2.
Note that the following analysis is carried out under the assumption of a stable circuit and of a loop gain much greater than 1. The loop gain will be better investigated in Section 5.2.4.
The Kirchhoff current law at that node can be written as:

$$\frac{v_{in}(s)}{R_{in}} = g_{m,A_1}\, v_{out}(s) + sC_{ac}\, v_{out}(s) \tag{5.6}$$

yielding:

$$v_{out}(s) = \frac{v_{in}(s)}{g_{m,A_1}\, R_{in}} \frac{1}{1 + sC_{ac} \,/\, g_{m,A_1}} \tag{5.7}$$

and:

$$i_{out,A_2}(s) = \frac{v_{in}(s)}{g_{m,A_1}\, R_{in}} \frac{sC_{ac}}{1 + sC_{ac} \,/\, g_{m,A_1}} \,. \tag{5.8}$$

As per the formula above, the transfer function of the output current of A_2 has a high-pass filter shape, with a zero in the origin and a pole at:

$$f_{ac} = \frac{g_{m,A_1}}{2\pi\, C_{ac}} \tag{5.9}$$

that has to be placed, ideally, at a much lower frequency than the first band of the filterbank.
The in-band transconductance of the circuit is:

$$G_m = \frac{i_{out,A_2}\,(s \to \infty)}{v_{in}\,(s \to \infty)} = \frac{1}{R_{in}} \tag{5.10}$$

which means that the generated current is inversely proportional to R_{in}, as expected.
A first compromise in the design of the circuit can be noticed: being the output voltage of the BPF cell quite small, a few mV in the worst case, the value of R_{in} must be chosen in the order of hundreds of $M\Omega$, so that a large enough current is produced for the neuron (at least a few tens of pA); on the other hand, the value of R_{in} also determines the DC current due to the offset voltage of A_2, as analyzed in Section 5.2.1, that has to be provided by A_1.
A solution to this compromise could be seen in biasing A_1 with a tail current large enough for the chosen value of R_{in}, but the high-pass cut-off frequency of generated current depends on g_{m,A_1} and therefore on the bias current of A_1.
As a consequence, after reasonably choosing the size of C_{ac} and minimizing the offset of A_2, a trade-off has to be made between the high-pass cut-off frequency and the rectified current produced for the minimum signal.
In our implementation, after reasonable transistor sizing, the offset of A_2 was found to be at most $\sim \pm 6$mV, accounting for a $\pm 3\sigma$ process and mismatch variation.
Note that, as opposed to the topology described in Section 5.1.1, in this case the current due to the DC offset of A_2 doesn't mix with the signal current as it follows a different path.
The capacitor C_{ac} was chosen to be 1 pF to limit area occupation.
The best trade-off for such values of offset and capacitance was achieved by placing the high-pass cut-off frequency at $\sim$ 100 Hz, accepting some slight attenuation in the first bands of the filterbank.

Knowing that:

$$f_{ac} = \frac{g_{m,A_1}}{2\pi\, C_{ac}} \tag{5.11}$$

the required (rounded) value for the tail bias current of A_1 was then calculated as:

$$I_{tail,A_1} = 2 \cdot g_{m,A_1}\, n\, V_T = 2 \cdot n\, V_T\, 2\pi\, C_{ac}\, f_{ac} \approx 40\,\text{pA} \tag{5.12}$$

leading also to:

$$\begin{cases} g_{m,A_1} = \dfrac{I_{tail,A_1}}{2\, n\, V_T} \approx 530\,\text{pS} \\ f_{ac} \approx 85\,\text{Hz}\,. \end{cases} \tag{5.13}$$

Leaving a margin of 10% to account for some mirroring error on the available tail current, the value of R_{in} can be then chosen as:

$$R_{in} = \frac{|V_{os,max}|}{I_{tail,A_1} \cdot 90\%} = 165\,\text{M}\Omega \tag{5.14}$$

so that the requirement on the minimum signal current can be fullfilled.

5.2.3. Rectifying Behavior

The transistor network between nodes v_x and v_x works as a class AB rectifying stage. Taking into account that for the in-band signal C_{ac} can be considered a short circuit, when the input signal v_{in} drives a current through R_{in} into the virtual ground of A_2, the loop reacts by modulating the node v_y.
During the positive half-cycle of the input current, M_1 is driven into conduction, steering the rectified small-signal component into the upper branch transdiode M_5, which defines the control voltage v_{gp1} needed to replicate the negative-wave rectified current i_{rec-}.
Conversely, during the negative half-cycle, M_1 switches off while M_2 turns on, redirecting the current through the lower branch NMOS current mirror made of M_6 and M_7. This current is subsequently mirrored and flipped via the transdiode M_8, generating the control voltage v_{gp2} needed to replicate the positive-wave rectified current i_{rec+}.
By combining these two complementary contributions as

$$i_{rec} = i_{rec+} + i_{rec-} \tag{5.15}$$

a continuous full-wave rectified current can be delivered to the neuron membrane capacitance.
A DC bias current I_S is imposed into M_3 and M_4 and as a consequence into M_1 and M_2, so that the loop can work properly and the block can be considered a simple voltage buffer between nodes v_x and v_y .
The minimum input signal can be considered the one that generates a rectified current equal to 50% of its non-rectified one, which means:

$$i_{in,min} = \frac{v_{in,min}}{R_{in}} = 2 \cdot I_S \,. \tag{5.16}$$

Aiming for a minimum input voltage of $\sim 2\,\mathrm{mV}$ and considering the value of R_{in} found in Section 5.2.2 it was chosen:

$$I_S = \frac{v_{in,min}}{2\,R_{in}} \approx 6\,\mathrm{pA} \tag{5.17}$$

leading also to:

$$g_{m1} = \frac{I_S}{n\,V_T} \approx 160\,\mathrm{pA} \tag{5.18}$$

as the transconductance of M_1 and M_2, assuming both equal for the sake of simplicity.

5.2.4. Loop Gain Analysis

To verify the actual stability of the circuit, the loop gain has to be analyzed. To this purpose it is convenient to cut the loop at the negative input of A_2: the resulting schematic, including the parasitics required for the analysis, is depicted in Figure 5.6.

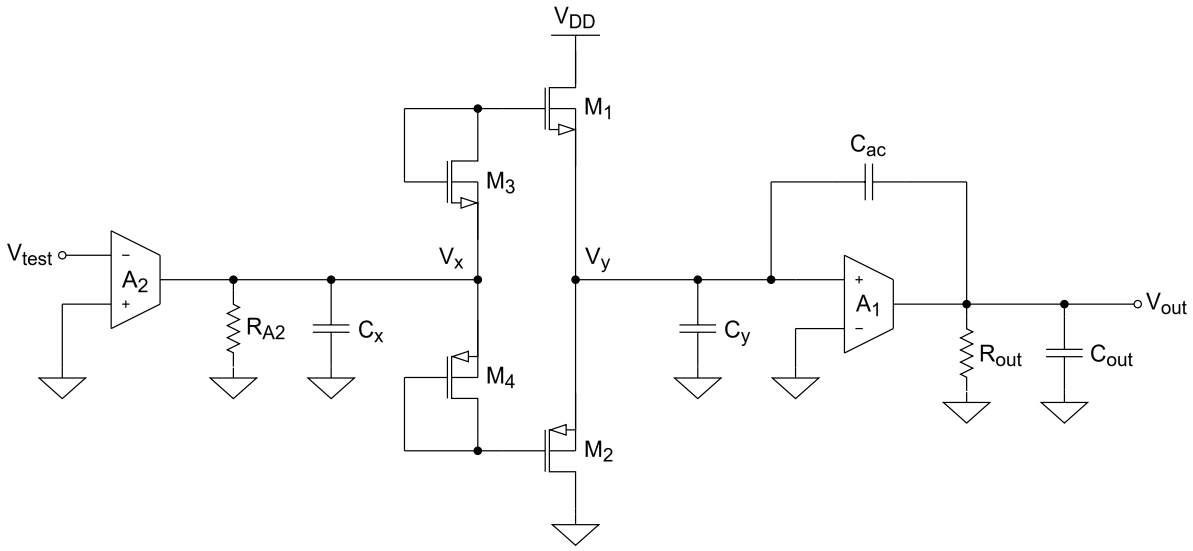

Figure 5.6: Simplified schematic for the computation of the loop gain of the proposed voltage-to-current converter.

The first step is to compute the DC value of the loop gain, from the test voltage v_{test} to the output voltage v_{out}: taking into account that transistors M_1 and M_2 act as complementary voltage followers, it is easy to compute:

$$G_{loop}(0) = g_{m,A_2}\, R_{A_2}\, g_{m,A_1}\, R_{out} \tag{5.19}$$

where R_{A_2} is the output resistance of A_2 and R_{out} is the resistance seen from the output node, given by the parallel of R_{A_1} and R_{in}. As already discussed in Section 5.2.2, $R_{in} = 165\,\mathrm{M\Omega}$ and $I_{tail,A_1} = 40\,\mathrm{pA}$, so clearly $R_{A_1} \gg R_{in}$ and $R_{out} \simeq R_{in} = 165\,\mathrm{M\Omega}$ can be safely assumed.
A_2 was biased with a tail current of 700 pA , for reasons that will be clear in the following, and its transistors were sized quite large to reduce its offset, leading to:

$$\begin{cases} g_{m,A_2} \approx 10\,\mathrm{nS} \\ R_{A_2} \approx 24\,\mathrm{G\Omega}\,. \end{cases} \tag{5.20}$$

The resulting value of the DC loop gain is then:

$$G_{loop}(0) = g_{m,A_2}\, R_{A_2}\, g_{m,A_1}\, R_{in} \approx 21 \approx 26\,\mathrm{dB}\,. \tag{5.21}$$

The next step is to compute the poles and zeros of the loop gain transfer function, which are due to the capacitors depicted in Figure 5.6: C_x is approximately the output capac-

itance of A_2, C_y is approximately the input capacitance of A_1, C_out is the sum of the output capacitance of A_1 and the input capacitance of A_2 and, finally, C_{ac} is the already known AC coupling capacitance.
The values of these capacitances in our implementation is:

$$\begin{cases} C_x \approx 20\,\text{fF} \\ C_y \approx 10\,\text{fF} \\ C_{out} \approx 20\,\text{fF} \\ C_{ac} = 1\,\text{pF}\,. \end{cases} \tag{5.22}$$

Note that the large size of C_{ac} in comparison to C_x and C_y allows to consider it independent in the analysis of poles and zeros, as its singularities will be at much lower frequencies.
The impact of the parasitic capacitances of M_1 and M_2 will be analyzed later, so they were neglected in Figure 5.6.
To compute the pole given by C_{ac} its equivalent resistance $R_{eq,p,ac}$ has to be computed: to this purpose a test current generator can be applied in place of the capacitor, that injects a positive current i_{test} towards the node v_y.
The voltage v_y at the left node of the test generator is then:

$$v_y = i_{test}\,\frac{1}{g_{m1} + g_{m2}} \simeq i_{test}\,\frac{1}{2\,g_{m1}} \tag{5.23}$$

where the transconductances of M_1 and M_2 were assumed to be equal as a simplification.
The current flowing into R_{out} can be found writing a Kirchhoff current law at the output node:

$$i_{out} = i_{test}\left(\frac{g_{m,A_1}}{2\,g_{m1}} - 1\right) \tag{5.24}$$

so that the output voltage is then:

$$v_{out} = i_{test}\left(\frac{g_{m,A_1}}{2\,g_{m1}} - 1\right) R_{out}\,. \tag{5.25}$$

The voltage across the current generator ends up being:

$$v_{test} = v_y - v_{out} = i_{test}\left[\frac{1}{2\,g_{m1}} + \left(1 - \frac{g_{m,A_1}}{2\,g_{m1}}\right) R_{out}\right] \tag{5.26}$$

which yields:

$$R_{eq,p,ac} = \frac{1}{2\,g_{m1}} + R_{out}\left(1 - \frac{g_{m,A_1}}{2\,g_{m1}}\right) . \tag{5.27}$$

Considering that $R_{out} \simeq R_{in}$ and that, with the values computed in Sections 5.2.2 and 5.2.3:

$$\begin{cases} g_{m,A_1} \approx 530\,\text{pS} \\ g_{m1} \approx 160\text{pS} \\ 1 - \dfrac{g_{m,A_1}}{2\,g_{m1}} \approx -0.66 \end{cases} \tag{5.28}$$

the value of $R_{eq,p,ac}$ can be computed as:

$$R_{eq,p,ac} = \frac{1}{2\,g_{m1}} \; - \; 0.66\,R_{in} \approx 3\text{G}\Omega\,. \tag{5.29}$$

The pole given by C_{ac} is then at:

$$f_{p,loop,ac} = \frac{1}{2\pi\,R_{eq,p,ac}\,C_{ac}} \approx 50\,\text{Hz}\,. \tag{5.30}$$

Note that this pole cannot easily be moved, as both C_{ac} and g_{m,A_1} were fixed in Section 5.2.2 and g_{m1} was fixed in Section 5.2.3.
To find the zero introduced by C_{ac}, we can consider $v_{out} = 0$ and a signal current i_y flowing into the capacitor, which leads to A_1 producing the current:

$$i_{out,A_1} = g_{m,A_1}\,i_y\,sC_{ac} \tag{5.31}$$

that can only flow back into C_{ac}, because the voltage drop across R_{out} is zero, so that $i_{out,A_1} = -i_y$ and therefore:

$$-i_y = i_{out,A_1} = g_{m,A_1}\,i_y\,sC_{ac} \;\Rightarrow\; s = -\frac{g_{m,A_1}}{C_{ac}} \;\Rightarrow\; f_{z,loop,ac} = \frac{g_{m,A_1}}{2\pi\,C_{ac}} \approx 85\,\text{Hz} \tag{5.32}$$

as expected from Section 5.2.2, as this zero corresponds to the high-pass pole in the ideal gain transfer function $v_{out}\,/\,v_{in}$.

As a next step, C_{ac} can be assumed to be a short circuit so that the pole given by the equivalent output capacitance $C_{out,eq} = C_{out} + C_y$ can be computed.
Note that $C_{out,eq}$ and C_x are independent capacitors as, for $f > f_{ac}$, M_1 and M_2 do not act as followers, being its load impedance much lower than ${g_{m1}}^{-1}$.
A test voltage generator can be applied in place of $C_{out,eq}$. The following Kirchhoff current law applies at the output node:

$$2\, g_{m1}\, v_{test} + \frac{v_{test}}{R_{out}} - g_{m,A_1}\, v_{test} = i_{test} \tag{5.33}$$

where i_{test} is the current provided by the test voltage generator.
The equivalent resistance seen by $C_{out,eq}$ is then:

$$\begin{aligned} R_{eq,p,out} &= \left(2\, g_{m1} + \frac{1}{R_{out}} - g_{m,A_1}\right)^{-1} = \frac{R_{out}}{1 - (g_{m,A_1} - 2\, g_{m1})\, R_{out}} \simeq \\ &\simeq \frac{R_{in}}{1 - (g_{m,A_1} - 2\, g_{m1})\, R_{in}} \approx 170\,\mathrm{M\Omega} \end{aligned} \tag{5.34}$$

and the pole given by $C_{out,eq}$ ends up being:

$$f_{p,loop,out} = \frac{1}{2\pi\, C_{out,eq}\, R_{eq,p,out}} \approx 32\,\mathrm{kHz}\,. \tag{5.35}$$

Note that the pole above cannot be easily moved at higher frequency: $C_{out,eq}$ is mainly set by the constraint on the offset of A_2 ,that limits the minimum transistor size and, as a consequence, its minimum capacitance; $R_{eq,p,out}$ is alredy fixed as discussed in Section 5.2.2.
The pole given by C_x is trivial, as the only path for the current coming from a test voltage generator is represented by R_{A_2}:

$$f_{p,loop,x} = \frac{1}{2\pi\, R_{A_2}\, C_x}\,. \tag{5.36}$$

This pole can be set, by properly tuning the current of A_2, as the dominant pole that determines the bandwidth of the circuit, being that the other singularities have each its own constraint already.
To guarantee $f_{G_{loop}=0dB} \approx 6\,\mathrm{kHz}$ it was chosen:

$$\begin{cases} I_{tail,A_2} = 3\,\text{nA} \\ R_{A_2} \approx 5\,\text{G}\Omega \\ g_{A_2} \approx 45\,\text{nS} \\ f_{p,loop,x} \approx 1.6\,\text{kHz}\,. \end{cases} \tag{5.37}$$

To better complete the analysis, we may compute the feedforward zero introduced by the gate-to-drain capacitance C_{gd,A_2} of the differential pair of A_2, which is a right-half-plane (positive) zero, and the singularities introduced by the parasitic capacitances of M_{1-4}, that were neglected in the analysis above.
The zero introduced by A_2 can be computed, knowing that $C_{gd,A_2} \approx 4.5\,\text{fF}$, as:

$$f_{z,loop,A_2} = \frac{g_{m,A_2}}{2\pi\, C_{gd,A_2}} \approx 1.5\,\text{MHz}\,. \tag{5.38}$$

Being it at such a high frequency that it doesn't have any influence on the phase margin.
To compute the singularities introduced by the parasitic capacitances of M_{1-4}, it must be noted that C_x can be considered a short circuit, being its pole at a much lower frequency. To achieve a more favorable condition in regard to the obtainable phase margin and stability, the bulk of transistors M_{1-4} was connected to their source terminal, so that the two negative zeros they introduce are pushed at lower frequency, being then $C_{gb} \gg C_{gs}$ in parallel to C_{gs}.
The two zeros can be calculated as:

$$f_{z,loop,M_3} = f_{z,loop,M_1} = \frac{g_{m1}}{2\pi\ (C_{gs1} + C_{gb1})} = \frac{160\,\text{pS}}{2\pi\ (165\,\text{aF} + 735\,\text{aF})} \approx 30\,\text{kHz}\,. \tag{5.39}$$

The capacitance between the gate node of M_1 / M_2 and ground introduces a pole in the loop gain. The following equations hold:

$$\begin{cases} R_{eq,p,loop,rect} = \dfrac{1}{g_{m1}} \\ C_{eq,p,loop,rect} = C_{gd1} + C_{gs3} + C_{gb3} \\ f_{p,loop,rect} = \dfrac{1}{2\pi\, R_{eq,p,loop,rect}\, C_{eq,p,loop,rect}} \approx 25\,\text{kHz}\,. \end{cases} \tag{5.40}$$

Note how the introduction of these singularities, in the same frequency range of $f_{p,loop,out}$,

lead to two pole/zero cancellations that greatly improve the final phase margin, leading to a value close to 90 deg .

The stability of the loop gain may be questioned because $f_{p,loop,ac}$ moves at higher frequency as the input signal increases, being that to rectify the current the circuit doesn't work in small-signal conditions and, for a larger input signal than $2I_S$ (the minimum detectable signal), the transconsuctance seen at the node v_y increases proportionally with the signal.

Examining the loop transfer function more carefully, though, it is evident that, actually, also all of the high-frequency singularities move when g_{m1} increases, except for $f_{z,loop,A_2}$ and $f_{p,loop,out}$: the former is fixed while the latter increases less than linearly.

The simulations performed agree with the above statement, showing stability up to $\sim$ 30dB of dynamic range for the input signal.

Figure 5.7 shows the theoretical Bode plot of the loop gain transfer function, while Figure 5.8 shows the result of the real simulation.

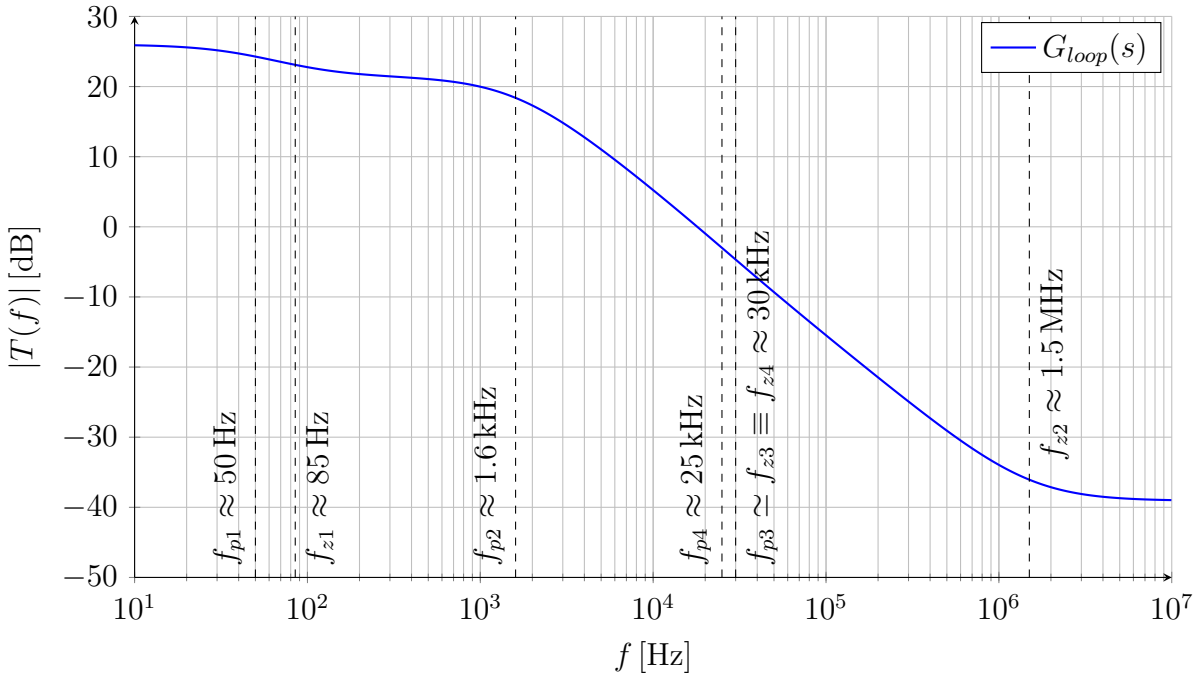

Figure 5.7: Bode plot of the loop gain of the voltage-to-current converter.

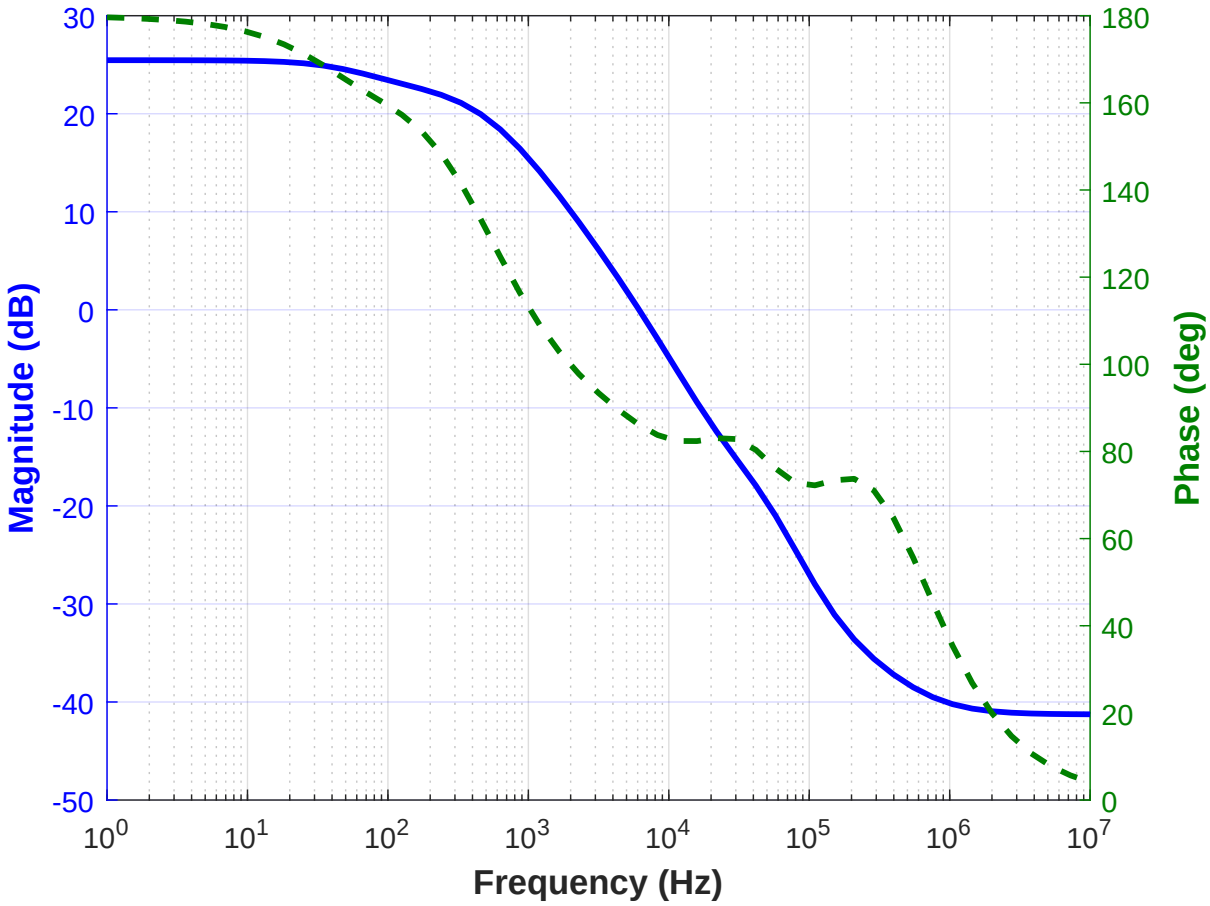

Figure 5.8: Simulated loop gain transfer function of the voltage-to-current converter.

5.3. Performance Analysis

The electrical and statistical performance of the implemented voltage-to-current converter has been thoroughly evaluated to verify its compliance with the strict energy and stability constraints dictated by the always-on neuromorphic processing chain.

5.3.1. Power Consumption

The power budget of the VTC block is tightly optimized since 16 independent converter channels operate in parallel, each processing the output of a specific BPF band. Under nominal conditions, each individual VTC core circuit consumes a static bias current of 6 nA. To evaluate the true power footprint, the current drawn by the subthreshold input pseudoresistor R_{in} must be added; each pseudoresistor accounts for an additional 2.5 nA of current overhead.

The total current consumption for a single channel is therefore 8.5 nA, leading to a com-

bined current for the entire 16-channel VTC block of:

$$I_{tot,VTC} = 16 \times (6\,\text{nA} + 2.5\,\text{nA}) = 136\,\text{nA} \tag{5.41}$$

Operating from a standard 1.8 V supply voltage rail, the total power dissipation of the full VTC stage is calculated as exactly 244.8 nW. This sub-microwatt performance confirms the efficacy of the class-AB subthreshold configuration, which keeps the quiescent current at the nanoampere scale while delivering sufficient dynamic range during signal rectification.

5.3.2. Process and Mismatch Variability

A major design concern in high-gain, subthreshold negative feedback loops is ensuring closed-loop stability across manufacturing corners and intra-die mismatch. To verify the robustness of the VTC topology, a Monte Carlo simulation was performed over $N = 15876$ runs, tracking the statistical distribution of the loop gain phase margin (PM). The resulting probability density function and its corresponding Gaussian fit are illustrated in Figure 5.9.

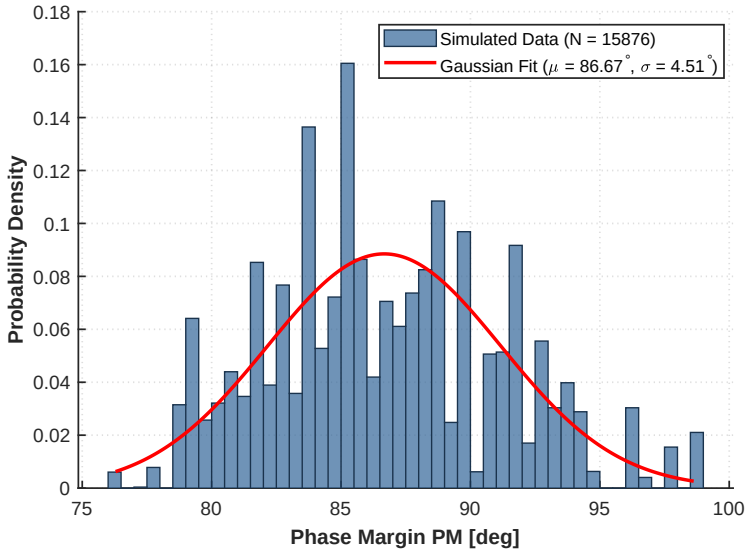

Figure 5.9: Monte Carlo statistical distribution ($N = 15876$ runs) and Gaussian fit of the loop gain phase margin (PM) for the proposed voltage-to-current converter.

The extracted Gaussian fit features a mean phase margin $\mu = 86.67°$ with a standard deviation $\sigma = 4.51°$[cite: 13]. The extremely high mean value confirms that the internal pole/zero cancellations introduced by the bulk-driven active devices (M_{1-4}), dis-

cussed during the analytical loop evaluation, operate robustly even in the worst-case mismatch scenarios. The narrow standard deviation guarantees that the phase margin never drops below 70° across the entire statistical ensemble, ensuring a completely stable, non-oscillatory step response and uniform settling behavior for all 16 channels of the filterbank interface.

6 | Conclusions and Future Developments

This thesis presented the design and analysis of an ultra-low-power analog front-end intended for acoustic feature extraction in neuromorphic audio classification systems. The proposed architecture was conceived to operate under the stringent power constraints imposed by always-on edge devices, where minimizing energy consumption is often more critical than achieving highly accurate waveform reconstruction.

The work began with the analysis of the principles underlying analog feature extraction and its advantages over conventional digital approaches based on high-resolution ADCs and MFCC computation. This design philosophy motivated the development of a fully analog processing chain composed of a low-noise amplifier, a tunable band-pass filter bank and a voltage-to-current conversion stage capable of directly interfacing with analog Leaky Integrate-and-Fire neurons.

Each circuit block was studied through analytical derivations and validated by extensive circuit simulations. Particular emphasis was placed on understanding the operating principles of the proposed topologies, deriving their transfer functions, evaluating stability conditions and assessing the effects of process and mismatch variability. The use of tunable pseudo-resistors enabled the realization of the large time constants required for low-frequency operation while maintaining a compact silicon area and preserving circuit tunability.

Overall, the obtained results confirm that the proposed architecture represents a viable solution for ultra-low-power analog preprocessing in acoustic inference applications. The designed building blocks satisfy the intended specifications while maintaining the simplicity and energy efficiency required by always-on neuromorphic systems.

Future developments may include the complete integration of the processing chain with the target analog Spiking Neural Network, the fabrication and experimental characterization of the proposed circuits, and the optimization of the architecture for more advanced acoustic inference tasks. Additional investigations could also address automatic calibration techniques to compensate for process variations and the extension of the proposed

methodology to other ultra-low-power sensing applications.

Bibliography

[1] P. J. Andres-Barquin. Ramón y Cajal: A century after the publication of his masterpiece. *Endeavour*, 25(1):13–17, Mar. 2001. ISSN 0160-9327. doi: 10.1016/S0160-9327(00)01334-X.

[2] D. Auge, J. Hille, E. Mueller, and A. Knoll. A Survey of Encoding Techniques for Signal Processing in Spiking Neural Networks. *Neural Processing Letters*, 53(6):4693–4710, Dec. 2021. ISSN 1573-773X. doi: 10.1007/s11063-021-10562-2.

[3] R. Carvajal, J. Ramirez-Angulo, A. Lopez-Martin, A. Torralba, J. Galan, A. Carlosena, and F. Chavero. The flipped voltage follower: A useful cell for low-voltage low-power circuit design. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 52(7):1276–1291, July 2005. ISSN 1558-0806. doi: 10.1109/TCSI.2005.851387.

[4] F. Currao. Ultra-Low-Power Analog CMOS Spiking Neural Network with Programmable Floating Gate Memory. Master's thesis, Politecnico di Milano, Milan, Italy, July 2024.

[5] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski. *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge University Press, Cambridge, UK, 2014.

[6] E. Guglielmi, F. Toso, F. Zanetto, G. Sciortino, A. Mesri, M. Sampietro, and G. Ferrari. High-Value Tunable Pseudo-Resistors Design. *IEEE Journal of Solid-State Circuits*, 55(8):2094–2105, Aug. 2020. ISSN 1558-173X. doi: 10.1109/JSSC.2020.2973639.

[7] I. P. Marciano. Custom Neuron and Synapse Models of an Analog Neuromorphic Chip for Hardware-Aware Training. Master's thesis, Politecnico di Milano, Milan, Italy, Apr. 2025.

[8] M. Mastella. Analog spiking neural network with floating gate synapses in standard CMOS technology. Master's thesis, Politecnico di Milano, Milan, Italy, Oct. 2019.

[9] M. Piccolino. Luigi Galvani and animal electricity: Two centuries after the foundation

of electrophysiology. *Trends in Neurosciences*, 20(10):443–448, Oct. 1997. ISSN 0166-2236. doi: 10.1016/s0166-2236(97)01101-6.

[10] M. Steyaert and W. Sansen. A micropower low-noise monolithic instrumentation amplifier for medical purposes. *IEEE Journal of Solid-State Circuits*, 22(6):1163–1168, Dec. 1987. ISSN 1558-173X. doi: 10.1109/JSSC.1987.1052869.

[11] M. Yang, C.-H. Yeh, Y. Zhou, J. P. Cerqueira, A. A. Lazar, and M. Seok. Design of an Always-On Deep Neural Network-Based 1- μ W Voice Activity Detector Aided With a Customized Software Model for Analog Feature Extraction. *IEEE Journal of Solid-State Circuits*, 54(6):1764–1777, June 2019. ISSN 1558-173X. doi: 10.1109/JSSC.2019.2894360.

[12] M. Yang, H. Liu, W. Shan, J. Zhang, I. Kiselev, S. J. Kim, C. Enz, and M. Seok. Nanowatt Acoustic Inference Sensing Exploiting Nonlinear Analog Feature Extraction. *IEEE Journal of Solid-State Circuits*, 56(10):3123–3133, Oct. 2021. ISSN 1558-173X. doi: 10.1109/JSSC.2021.3076344.

[13] S. Zhak, M. Baker, and R. Sarpeshkar. A low-power wide dynamic range envelope detector. *IEEE Journal of Solid-State Circuits*, 38(10):1750–1753, Oct. 2003. ISSN 1558-173X. doi: 10.1109/JSSC.2003.817599.

[14] Z. Zhu and L. Feng. A Review of Sub- μ W CMOS Analog Computing Circuits for Instant 1-Dimensional Audio Signal Processing in Always-On Edge Devices. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 71(9):4009–4018, Sept. 2024. ISSN 1558-0806. doi: 10.1109/TCSI.2024.3408819.

List of Figures

List of Tables