

Master thesis : Neuromodulation Controller Enabling Robustness to Transistor Variability in Subthreshold Silicon Neurons

Auteur : Filée, Victoria

Promoteur(s) : Drion, Guillaume

Faculté : Faculté des Sciences appliquées

Diplôme : Master : ingénieur civil électricien, à finalité spécialisée en Neuromorphic Engineering

Année académique : 2025-2026

URI/URL : <http://hdl.handle.net/2268.2/26101>

Avertissement à l'attention des usagers :

Tous les documents placés en accès ouvert sur le site le site MatheO sont protégés par le droit d'auteur. Conformément aux principes énoncés par la "Budapest Open Access Initiative"(BOAI, 2002), l'utilisateur du site peut lire, télécharger, copier, transmettre, imprimer, chercher ou faire un lien vers le texte intégral de ces documents, les disséquer pour les indexer, s'en servir de données pour un logiciel, ou s'en servir à toute autre fin légale (ou prévue par la réglementation relative au droit d'auteur). Toute utilisation du document à des fins commerciales est strictement interdite.

Par ailleurs, l'utilisateur s'engage à respecter les droits moraux de l'auteur, principalement le droit à l'intégrité de l'oeuvre et le droit de paternité et ce dans toute utilisation que l'utilisateur entreprend. Ainsi, à titre d'exemple, lorsqu'il reproduira un document par extrait ou dans son intégralité, l'utilisateur citera de manière complète les sources telles que mentionnées ci-dessus. Toute utilisation non explicitement autorisée ci-avant (telle que par exemple, la modification du document ou son résumé) nécessite l'autorisation préalable et expresse des auteurs ou de leurs ayants droit.



UNIVERSITY OF LIÈGE
Faculty of Applied Sciences

Neuromodulation Controller Enabling Robustness to Transistor Variability in Subthreshold Silicon Neurons

by VICTORIA FILÉE

Thesis submitted in partial fulfillment of the requirements for the degree
of **Master of Science in Electrical Engineering**

Thesis supervisor:

GUILLAUME DRION

Co-supervisor:

ARTHUR FYON

Jury members :

JEAN-MICHEL REDOUTÉ

ALESSIO FRANCI

ARTHUR FYON

LORIS MENDOLIA

– Academic year: 2025-2026 –

ACKNOWLEDGMENTS

First of all, I would like to sincerely thank my supervisor, Guillaume Drion, for giving me the opportunity to work on this very interesting subject and for his invaluable guidance throughout this project. I would also like to express my gratitude to Arthur Fyon and Loris Mendolia for their help and valuable advice during this thesis. They were always willing to promptly share their experience and expertise whenever needed.

I am also deeply grateful to my family, and especially to my siblings, whose unwavering support meant the world to me. They were there whenever I needed them, and even when I did not realize I did, giving me the strength to pursue my goals and never give up.

I would like to thank my friends, who brought laughter, parties, debates, and study sessions to another level, creating memories that I will always look back on with a smile.

Finally, I wish to thank my boyfriend, Hugo, for his constant optimism, advice, humor, and love.

DECLARATION ON THE USE OF AUTOMATIC TOOLS FOR WRITING THE MANUSCRIPT

I hereby certify that I have not used any generative intelligence tool in the writing of text, graphics, images, or data reproduced in this manuscript.

DeepL and ChatGPT were used throughout the project solely for spelling and syntax checking.

In the brain, billions of neurons communicate through neuronal signaling based on both electrical and chemical processes, particularly action potentials. This communication process is implemented in analog neuromorphic neurons using electronic circuits that emulate the behavior of biological neurons. In such implementations, transistor-intrinsic mismatch arises, causing intra- and inter-neuron variability. One consequence of mismatch is the difficulty of implementing neuromodulation, the process by which neuronal activity is regulated, enabling neurons to adapt their behavior to context and to support different functions.

Therefore, this thesis investigates the feasibility of feedback-based regulation of neuronal activity, robust to mismatch, in an analog neuromorphic neuron.

A controller is designed using dynamic input conductances for neuronal activity tracking, allowing control of excitability instead of raw voltage trajectories. This approach is first validated with the computational model of the neuron chip and against artificial mismatch. Then, variability is characterized both globally and individually by incorporating mismatch models, providing insight into the impact of device variability. After controller adaptation and neuron calibration, the control loop is tested on ultra-low-power neurons implemented on the developed chip, while accounting for practical mismatch effects and electronic constraints.

The results show that neuronal activity can be robustly regulated despite variability. However, large-scale implementation appears impractical, as the parameters must be specifically calibrated for each individual neuron. The framework is therefore evaluated at higher power levels to reduce mismatch-induced variability. Under these conditions, neuronal dynamics become more consistent across neurons, enabling a common calibration strategy that is more suitable for large-scale integration.

Overall, this work demonstrates that neuromodulation is achievable in analog subthreshold neurons through activity-based feedback control, despite the pronounced effects of device mismatch, contributing to robust, adaptive behavior in scalable and biologically inspired neurons.

Keywords: neuromorphic chip, neuromodulation, transistor mismatch, controller.

Introduction	1
1 Background	4
1.1 Excitability in neurons	4
1.2 Dynamic input conductances	8
1.3 CMOS technology	9
1.3.1 Basics of MOSFETs	9
1.3.2 Subthreshold region	11
1.3.3 Mismatch in transistors	12
1.4 The neuromorphic chip	13
1.4.1 State-of-the-art	13
1.4.2 A neuromodulable current-mode silicon neuron	14
2 Model-based control for a silicon neuron	24
2.1 Model description	24
2.2 Slow DIC for excitability quantification	26
2.3 Control loop design	28
2.3.1 The PI controller	28

2.3.2	Spike to DIC estimate	30
2.3.3	Controller gains tuning	32
2.4	Robustness to variability caused by mismatch	34
3	Mismatch characterization in Cadence	39
3.1	Mismatch impact on the neuron output	39
3.2	Individual mismatch and sensitive transistors	41
4	Neuromodulation of ultra-low-power neurons	42
4.1	Set-up	42
4.2	Impact of device mismatch on inter-neuron variability	43
4.3	Inter-neuron mismatch consequences on neuromodulation	44
4.4	Compensation and exploitation of mismatch	46
4.5	Control loop integration for robust neuromodulation	46
4.5.1	Discrete slow DIC estimation for activity control	46
4.5.2	Control loop implementation	50
4.6	Results of neuromodulation in ultra-low power neurons	52
4.6.1	Control loop impact on firing pattern variability	52
4.6.2	Evolution of parameters during control	52
4.6.3	Neurons output evolution	57
4.7	Conclusions on neuromodulation for ultra-low power CMOS neurons	57
5	Neuromodulation on low-power neurons	60
5.1	Mismatch effects on low power neurons	60
5.2	Neuromodulation in low-power neurons	63
5.3	Controller adaptation	63
5.4	Results of neuromodulation in low-power neurons	65

Discussion, perspectives and conclusion	67
Appendices	70
A Simulation parameters	70
B Additional results	73
C Additional materials	79
C.1 Control interfaces	79
C.2 Codes	80
Bibliography	84

Motivations

The biological brain is an extraordinarily complex organ that exhibits remarkable flexibility and robustness. It is able to adapt its behavior to changing contexts without requiring any modification of its underlying neuronal wiring. As a result, the same neural circuitry can support multiple functionalities and generate distinct behavioral outputs depending on environmental conditions. A key mechanism enabling this versatility is neuromodulation, which allows neural activity to be dynamically regulated without altering the structural connectivity of the network. Unlike fast synaptic transmission, which conveys specific signals between neurons, neuromodulation operates through chemicals such as serotonin and dopamine, which act diffusely to influence neuronal activity [1]. Among other mechanisms, neuromodulators can indeed bind to metabotropic receptors, initiating intracellular signaling cascades that can adjust a neuron intrinsic excitability and firing patterns. For instance, these cascades can switch neurons between tonic firing and bursting modes, thereby reshaping the dynamics of the entire circuit [2].

A well-studied example of neuromodulation in action is found in central pattern generators (CPGs). These are neural circuits that produce rhythmic motor outputs such as walking, running, or swimming, even in the absence of changing sensory inputs [3]. The activity of CPGs is finely tuned by context-dependent neuromodulatory signals, allowing animals to adapt their movement patterns to different environmental conditions or behavioral goals [4]. Through this mechanism, a single neural circuit can support multiple functional states, highlighting the adaptability conferred by neuromodulation.

In contrast, classical computing systems generally lack this level of adaptability. They are typically based on fixed architectures, most notably the Von Neumann architecture, in which changes in system behavior require explicit reprogramming or, in some cases, structural modifications. Moreover, when compared to the biological brain, such architectures exhibit relatively

high power consumption while offering limited processing efficiency [5].

Hence, neuromorphic engineering is an emerging discipline that seeks to reproduce the computational principles of biological neural networks using analog or digital hardware. By mimicking the event-driven and massively parallel nature of the brain, neuromorphic systems can achieve high energy efficiency, low latency, and scalable adaptability. These systems often integrate memory and processing elements co-located, enabling real-time, context-dependent responses that resemble the flexibility of biological circuits [6, 7]. In the analog domain, current-mode transistor design operating in the subthreshold region has emerged as a powerful approach for neuromorphic implementations [8]. Operating in this regime enables the realization of strong nonlinearities at ultra-low currents, since the current–voltage relationship of field-effect transistors (FETs) is exponential. In addition, current-mode design naturally supports the direct addition and subtraction of signals, which is particularly advantageous for implementing neural computations.

Despite these advantages, most complementary metal oxide semiconductor (CMOS) neuron implementations today primarily emulate integrate-and-fire behavior with fixed dynamics, thereby lacking the adaptive adjustment mechanisms observed in biological neurons, the neuromodulation [9]. A recent and notable development addressing this limitation is the current-mode mixed-feedback CMOS neuron presented by Mendolia et al. [10]. This model relies exclusively on differential pair integrators (DPIs) and sigmoidal nonlinearities, which are interconnected through feedback loops operating at different timescales to mimic key dynamical properties of biological neurons. As a result, this approach enables a compact, modular, and ultra-low-power hardware realization of a neuromodulable neuron.

The goal of this thesis is to design and implement a controller for this neuromodulable neuron in order to achieve robust neuromodulation by managing excitability. Neuromodulation has been extensively studied from a dynamical systems perspective, where it has been shown that a small number of parameters is sufficient to control qualitative neuron behaviors such as spiking and bursting [11]. However, in neuromorphic hardware implementations, these modulatory mechanisms are typically realized in an open-loop manner through bias tuning, making them highly sensitive to unavoidable device mismatch and variability [12]. This creates a gap between the theoretical efficiency of neuromodulation and its practical robustness in analog circuits. In this thesis, we aim to bridge this gap by introducing a closed-loop neuromodulation strategy that enforces desired firing regimes despite hardware non-idealities. The proposed controller is intended to compensate for intra-neuron mismatch, arising from variations among transistors within a single neuron, while a calibration procedure is employed to address inter-neuron mismatch, corresponding to variations between nominally identical transistors across different neurons.

Contributions

The main contributions of this thesis are as follows:

- Design and implementation of a PI controller for activity regulation in Julia.
- Development of a processing pipeline for activity classification on-chip.
- Characterization of inter-neuronal variability on chip across two current regimes.
- Mapping of neuronal activity to estimated slow DIC using both software and hardware implementations.
- Establishing the slow dynamic input conductance is a satisfactory metric for neuronal activity control.

Outline

This thesis is divided into chapters, each addressing a distinct part of the process and allowing for a clear understanding of the work. They are organized as follows.

Chapter 1 This first chapter gathers all the knowledge required to understand this thesis. The topics covered include neuronal excitability, CMOS technology, neuromorphic chips (including a detailed description of the chip used in this study), and dynamic input conductances.

Chapter 2 The model of the mixed-feedback neuron is used to explore its behavior and to design a first controller in this second chapter. The concept of dynamic input conductance is applied to the model and used to design a control loop. Artificial mismatch is then introduced to test the robustness of the approach.

Chapter 3 The goal of this chapter is to characterize variability using a mismatch model in addition to the chip model. The characterization is carried out globally, across all transistors, and individually.

Chapter 4 True mismatch is addressed here, where the electronic chip is used for neuromodulation. This chapter uses calibration techniques and a controller to handle variability in neurons biased in the picoampere regime. Results are presented and discussed.

Chapter 5 This chapter is similar to the previous one; however, neurons are biased in the nanoampere regime, leading to controller adaptations and different results.

Discussion, perspectives and conclusion Results from the two controller implementations are compared and discussed, highlighting the limitations and advantages of both current regimes. Some perspectives and a brief conclusion are also provided.

Appendices This chapter contains parameters used in the process, additional results, and other relevant information.

This chapter gathers all the fundamental knowledge required for a proper understanding of this thesis. It begins with the basics of biological neurons, more specifically their excitability, which helps explain the motivation and inspiration behind neuromorphic engineering. A brief overview of CMOS technology, widely used in neuromorphic systems, is then presented. The chapter continues with a presentation of the chip used in this work, first situating it within the state of the art and then providing a detailed description. Finally, the concept of dynamic input conductances is introduced.

1.1 Excitability in neurons

The brain is a network capable of receiving, processing and transmitting information. It allows light to be converted to perception or pain to lead to withdrawal decision. The flow of information is supported by the basic units, the neurons. Neurons receive, integrate and transmit signals, relying on action potentials, or spikes. A neuron receives information from others through its dendrites, more precisely at its synapses where the chemical signals are transformed into potentials. These postsynaptic potentials reach the axon initial segment, where they are integrated. If the combined inputs are strong enough to reach a threshold, the neuron internal dynamics generate an action potential, which then propagates along the axon to other neurons [13, 14]. The flow of information between neurons is illustrated in Figure 1.1.

The plasma membrane of the neuron mediates this action potential. The membrane is not directly permeable to ions, but it contains transmembrane proteins that allow the crossing of larger molecules and ions. Among these, ion channels regulate the exchange of ions between the intracellular and extracellular spaces, thereby generating currents. At rest, the membrane potential, resulting from the distribution of ions across the membrane and the membrane permeability to these ions, is typically around -70 mV. When inputs are summed, the membrane




FIGURE 1.1 – **Synaptic transmission.** Left: pathway of the electrical signal (action potential) from the axon initial segment of the pre-synaptic neuron to the dendrites of the post-synaptic neuron through the synapse. Right: zoom on the synapse, when the electrical signal of the pre-synaptic neuron arrives at the synapse, the neurotransmitters are released in the synaptic cleft and bind to receptors (chemical synapse). The receptors activation will either trigger a direct flow of ions or second messenger cascades (neuromodulators). Taken from BioRender.

potential changes, altering the opening of ion channels, most of which are voltage-dependent, and thereby further affecting the membrane potential. If this potential crosses a threshold, it rapidly rises to approximately 40 mV; this phase is known as depolarization. Repolarization occurs when the potential returns to its resting value. The ability of the neuron to generate and propagate an action potential in response to stimulus is called neuronal excitability.

There is a wide variety of ion channels in cells; however, to generate an action potential, at least fast transient sodium and delayed rectifier potassium channels are required. Sodium ions are responsible for depolarization: as the membrane potential increases, permeability to sodium ions increases, leading to an influx of positive ions. This corresponds to a fast positive feedback. Potassium ions, on the other hand, are responsible for repolarization. Once the membrane potential reaches its peak, sodium ions can no longer cross the membrane because sodium channels are blocked by the inactivation gate, while potassium channels open, creating an outward flow of potassium ions and driving repolarization. These two mechanisms (inactivation of sodium channels and the activation of potassium channels) operate on a slower timescale, hence resulting in a slow negative feedback. Because potassium channels do not have an inactivation gate, the membrane potential can drop below the resting potential; this phase is called hyperpolarization. During this phase, potassium channels gradually close and sodium inactivation gates reopen. The membrane potential then returns to its resting value (Figure 1.2).

Neuronal signaling exhibits a rich diversity of firing patterns, reflecting distinct types of neuronal excitability that are primarily determined by the distribution of ion channels [15–18].




FIGURE 1.2 – **The action potential.** A diagrammatic representation of the action potential and associated models depicting the sequence of activation and inactivation of the voltage sensitive Na^+ and K^+ channels. Arrows indicate the state of the respective ion channels during different phases of the action potential. The broken line on the action potential trace indicates the resting membrane potential and is used to highlight the rebound after hyperpolarization. Taken from [14].

Neurons can exhibit, for example:

- **Type I neuronal excitability:** this corresponds to a tonic spiking pattern in which the neuron can fire at arbitrarily low frequencies depending on the applied current. The current–frequency (I–f) curve is therefore continuous (Figure 1.3, left). This type of neuron can be interpreted as an encoder of the input current magnitude.
- **Type II neuronal excitability:** this also corresponds to a tonic spiking pattern, but low firing frequencies are not achievable, leading to a discontinuity in the I–f curve (Figure 1.3, right). Once the input current exceeds the threshold, the neuron fires immediately at a relatively high frequency. This behavior can be associated with a binary-like system, in which the neuron is either silent or firing at high frequency.
- **Persistent (tonic) bursting:** this firing pattern is characterized by repetitive groups of close spikes, the bursts, with relatively constant intra-burst and inter-burst frequencies. It can be roughly assimilated to pulse width modulation (PWM), with a duty cycle controlled by the intra-burst frequency and the number of spikes per burst.

Tonic bursting and several other firing behaviors that a neuron can exhibit are shown in Figure 1.4, additional patterns also exist.

The excitability of neurons is not fixed and can adapt to the environment through a process known as neuromodulation. Neuromodulation is mediated by a specific type of neurotransmitter,




FIGURE 1.3 – **Type I and type II excitability.** The top of each panel shows membrane potential (V_m) traces for two different step input currents (I_{app}) (black and orange traces). The bottom of each panel shows neuron firing rate as a function of the input current (I-f curve). Reproduced and adapted from [19].




FIGURE 1.4 – **More firing patterns.** Other types of firing patterns that can be found in biological neurons, reproduced with a neuro-computational model. Taken and adapted from [20].

which alter neuronal dynamics, excitability, and synaptic function [1, 21]. These molecules, such as dopamine and serotonin, are released into the extracellular space and act diffusely over local regions of the brain. Unlike classical neurotransmitters, they do not directly mediate synaptic transmission. Instead, they bind to metabotropic receptors, triggering intracellular second messenger cascades. The resulting effect is a modification of ion channel properties and distribution, thereby modulating neuronal excitability.

By contrast to classical computers inspired by the Von Neumann architecture, which operate sequentially and rely on processing and memory separation, the brain, as a computing tool, offers memory and computation combined in one unit (the neuron) [5]. These units work massively in parallel, while still ensuring low power-consumption and high robustness. In addition, the brain learns and adapts in real-time, compared to the need for a huge training phase for some classical AI models. Therefore, attention towards emulating these biological principles in hardware and algorithms has emerged in a field called neuromorphic engineering, which seeks to exploit these features of biological neuron systems to offer energy-efficient real-time processing [22]. Applications could be low-power artificial intelligence, real-time interactions, sensory processing (e.g. event-based cameras), etc [23].

1.2 Dynamic input conductances

Dynamic input conductances (DICs) were introduced by Drion et al. [24]. They allow, in conductance-based models and biological neurons, to avoid model reduction and, instead, account for all ion channel contributions across voltages and timescales by proposing a low dimensional representation of the dynamics. DICs have been shown to be crucial in neuronal studies, such as understanding neuronal dynamics and exploring neuromodulation [25, 26].

At steady state, the local sensitivity of the membrane voltage of a neuron to a small current deviation ΔI is, using a Taylor expansion,


$$\Delta V = \left(\frac{\partial I}{\partial V} \right)^{-1} \Delta I. \quad (1.1)$$

DICs arise from the separation of neuronal dynamics into three timescales: fast (spike upstroke), slow (spike downstroke and inter-spike period), and ultra-slow (spike adaptation and inter-burst period). Under this decomposition, the change in current ΔI induced by a change in voltage ΔV can be expressed as

$$\Delta I_f + \Delta I_s + \Delta I_u = \Delta I. \quad (1.2)$$

This yields, for the voltage,

$$-g_f(V)\Delta V - g_s(V)\Delta V - g_u(V)\Delta V = -g_{in}(V)\Delta V, \quad (1.3)$$



and thus

$$g_f(V) + g_s(V) + g_u(V) = g_{in}(V). \quad (1.4)$$

The complete DIC, g_{in} , can therefore be decomposed into three components, one for each timescale, under the assumption that for a given timescale, currents associated with faster timescales have already reached steady state, while those associated with slower timescales remain negligible.

These DIC curves shape the feedback gain across the different timescales, thereby determining the dynamical activity. An example of these curves, along with the feedback mechanisms underlying spike generation and bursting, is shown in Figure 1.5. The positive feedback of the fast timescale is associated with the regenerative upstroke of the action potential. In the slow dynamic input conductance, negative feedback is responsible for repolarization, while a small region of positive feedback is essential for regenerative excitability and, consequently, bursting. This positive region in the slow DIC is of particular interest for neuromodulation. Finally, the ultra-slow negative feedback regulates spike adaptation.

This decomposition of a neuron dynamics into three timescales is the foundation of electronic mixed-feedback neurons such as the one presented later in Section 1.4.2.

1.3 CMOS technology

1.3.1 Basics of MOSFETs

The metal–oxide–semiconductor field-effect transistor (MOSFET or MOS transistor) is the fundamental building block of modern electronics [27]. It can be described as a tiny electronic

switch controlled by voltage. A MOSFET has three terminals: the gate, the source, and the drain (Figure 1.6). The voltage applied to the gate controls the current flowing between the drain and the source. In practice, MOSFETs are structurally symmetric devices, the roles of drain and source are determined by the applied bias conditions.




FIGURE 1.6 – **MOSFETs**. Symbols of the NMOS and PMOS, annotated with the gate (G), the drain (D), the source (S) and the drain current (I_D).

There are two types of MOS transistors:

- n-channel MOS transistor (NMOS or nFET): the majority charge carriers are electrons. When the device is turned on, current flows between drain and source through an electron channel (Figure 1.6, left).
- p-channel MOS transistor (PMOS or pFET): the majority charge carriers are holes. When the device is turned on, current flows between source and drain through a hole channel (Figure 1.6, right).




FIGURE 1.7 – **CMOS technology using n-well process**. The green (resp. red) areas are the p-type (resp. n-type) material and the white p^+ and n^+ regions are the bulks. D stands for the drain, S for the source, G for the gate and B for the bulk. The blue insulator under the gate is the gate oxide. Reproduced from [28].

Complementary metal–oxide–semiconductor (CMOS) technology uses both NMOS and PMOS transistors on the same substrate. In CMOS circuits, one transistor type conducts while the other is off, allowing charge carriers (electrons or holes) to flow between the doped terminals in a complementary manner. Two important parameters in CMOS technologies are the width and the length of the channel between the two doped terminals as they determine the amount of current the transistor can carry as well as the switching speed.

MOSFETs are operated either in depletion mode, in which the transistor is on when the gate-to-source voltage is zero, or in enhancement mode, in which the transistor is off at zero gate-to-source voltage. In enhancement mode, the most common in CMOS technology, there are two basic fabrication processes: the n-well process and the p-well process. An n-type material is created by doping silicon with impurities that provide extra electrons, whereas a p-type material is created by doping silicon with impurities that generate holes (missing electrons). Today, the dominant technology is the n-well process implemented on a p-type substrate (Figure 1.7).

1.3.2 Subthreshold region




FIGURE 1.8 – **Subthreshold region.** Drain current in a n-type MOSFET as a function of the gate voltage with delimitation of the subthreshold, superthreshold and transition regions. Reproduced and adapted from [29].

As represented in Figure 1.8, for small gate-to-source voltage V_{gs} the relationship between V_{gs} and the current flowing through the transistor is exponential (linear in log scale). More specifically, in the subthreshold region [30],

$$\text{NMOS: } I_D = I_{n0} e^{\kappa_n \frac{V_g}{U_T}} \left(e^{-\frac{V_s}{U_T}} - e^{-\frac{V_d}{U_T}} \right) \quad (1.5)$$

$$\text{PMOS: } I_D = I_{p0} e^{\kappa_p \frac{V_{dd}-V_g}{U_T}} \left(e^{-\frac{V_{dd}-V_s}{U_T}} - e^{-\frac{V_{dd}-V_d}{U_T}} \right) \quad (1.6)$$

Where I_{n0}/I_{p0} is the technology-dependent current scaling parameter (leakage current), κ_n/κ_p is the subthreshold slope factor, U_t is the thermal voltage, V_g , V_d , and V_s are the gate, drain,

and source voltages respectively and V_{dd} is the supply voltage, connected to the bulk for PMOS, the bulk being the semiconductor material under the channel. For a large enough difference between V_d and V_s ($V_d - V_s > 4U_T$), which is called the saturation regime, these equations can be rewritten as

$$\text{NMOS: } I_D = I_0 e^{\frac{\kappa V_g - V_s}{U_T}} \quad (1.7)$$

$$\text{PMOS: } I_D = I'_0 e^{\frac{V_s - \kappa V_g}{U_T}} \quad (1.8)$$

This region of operation is particularly interesting for neuromorphic analog circuits because of its low power consumption, nonlinear (exponential) behavior, high transconductance efficiency, and the possibility of achieving long time constants [31].

1.3.3 Mismatch in transistors

Mismatch in transistors is inherent due to imperfections in the fabrication process. This implies that two identically designed transistors can exhibit different electrical behaviors [12, 32].

These manufacturing variations lead to differences in device parameters at multiple levels, including lot-to-lot, wafer-to-wafer, die-to-die, and device-to-device [33]. Variability across lots and wafers is generally systematic, and appropriate design techniques exist to make circuit performance insensitive to these variations. At the die level, systematic position-dependent variations arise due to processing gradients. Layout techniques such as symmetry and common-centroid configurations can be used to significantly reduce gradient effects.

Regarding device-to-device variability, Pavasović et al. [34] studied MOS transistor matching in the context of very-large-scale integration (VLSI) systems, with an emphasis on subthreshold operation. They identified the following key sources of mismatch:

- Random local variations: caused by the inherent randomness of the fabrication process, including fluctuations in dopant concentration, oxide thickness, and channel characteristics.
- Edge effects: arising from asymmetric surroundings (e.g. stress, isolation, or missing neighboring devices), causing edge devices to behave differently from interior ones.
- Striation effects: periodic, position-dependent current variations across the wafer or die.
- Gradient effects: long-range trends in parameter variation across the array, similar to die-level gradients.

These effects tend to alter the threshold voltage, which is exponentially related to the current in the subthreshold region, resulting in high sensitivity to mismatch. The study also emphasizes that mismatch can be decomposed into random and spatially correlated components. Moreover,

after removing systematic spatially correlated components, the remaining random mismatch scales approximately with $1/\sqrt{W \cdot L}$, where W and L are the width and length of the transistor channel, respectively. This is consistent with Pelgrom’s law [12].

Therefore, mismatch effects are more pronounced in LSI/VLSI systems that operate small transistors in the subthreshold region, since small variations in voltage translate exponentially into current differences, and because such systems typically rely on small-area transistors. In this regime of operation, mismatch is typically modeled with a normal distribution [35]. Hence, mismatch represents a fundamental limitation in neuromorphic implementations, where VLSI in subthreshold operation is commonly used, amplifying device variability, leading to significant dispersion in circuit behavior.

1.4 The neuromorphic chip

1.4.1 State-of-the-art

Since the emergence of neuromorphic engineering around 1989 with Maher et al. [36], which introduced the idea of emulating biological systems using analog very-large-scale integration, diverse technologies and applications have been explored. Various implementations ranging from analog and mixed-signal to fully digital, as well as hybrid approaches (e.g. analog circuits with digital communication), can be found in Singh Thakur et al. [9].

Regarding digital technologies, a recent example is Loihi from Intel Labs [37], a neuromorphic chip that implements spiking neural networks directly in hardware. The chip contains 128 neuromorphic cores capable of performing on-chip training and inference. With a total of 130 thousand neurons fabricated using Intel 14-nm process, the chip integrates hierarchical connectivity, dendritic compartments, synaptic delays, and programmable synaptic learning rules. The neuron units used are based on leaky integrate-and-fire (LIF) models.

Another example, based on mixed-signal technology, is the DYNAP-SE2 by Richter et al. [38]. It is an asynchronous spiking neural network processor, combining analog processing, for direct emulation of dynamic and realistic neural processing, with digital circuits for routing and mapping events. This provides small networks of spiking neurons that emulate the dynamics of real neurons and synapses.

However, most state-of-the-art silicon neurons designed for large-scale implementations rely on integrate-and-fire models. While these models enable efficient and compact neuron circuits suitable for large-scale integration, they lack the dynamical richness required to reproduce more complex biological behaviors, such as bursting or rest-spike bistability. Among silicon neurons exhibiting a wider range of dynamics, the design presented in Hynna et al. [39] demonstrates both high-frequency bursting (>250 Hz) and lower-frequency spiking (<100 Hz). Similarly, the neuron proposed by Wijekoon et al. [40] is capable of generating both spiking and bursting behaviors. More recently, the neuron developed by Mendolia et al. [10] has demonstrated a

particularly rich set of dynamical regimes together with improved energy efficiency compared to previous implementations. For these reasons, this CMOS neuron will be used in this thesis as the basis for the application of robust neuromodulation.

1.4.2 A neuromodulable current-mode silicon neuron

This section describes the CMOS neuron presented by Mendolia et al. [10] and used in the context of the thesis to implement robust neuromodulation.

The current-mode mixed feedback model

The proposed approach is a mixed-feedback current-mode circuit to ensure more compact, robust and energy efficient designs [41]. The mixed-feedback aspect combines positive and negative feedbacks at different timescales, while current-mode implementation represents state variables, e.g. the membrane potential and the feedback terms, as physical currents. This last choice directly brings trivial addition and subtraction of variables thanks to Kirchhoff's current law.

The structure of the model is rather simple as it relies on two types of building blocks: first-order low-pass filters and static sigmoidal functions. The first order current-mode low-pass filter is described in state-space representation, with input I_{in} and output I_{out} , as

$$\tau \dot{I}_{out} = -I_{out} + GI_{in} \quad (1.9)$$

where τ is the time constant of the filter and G its gain. Moving into the Laplace domain, (1.9) is rewritten as

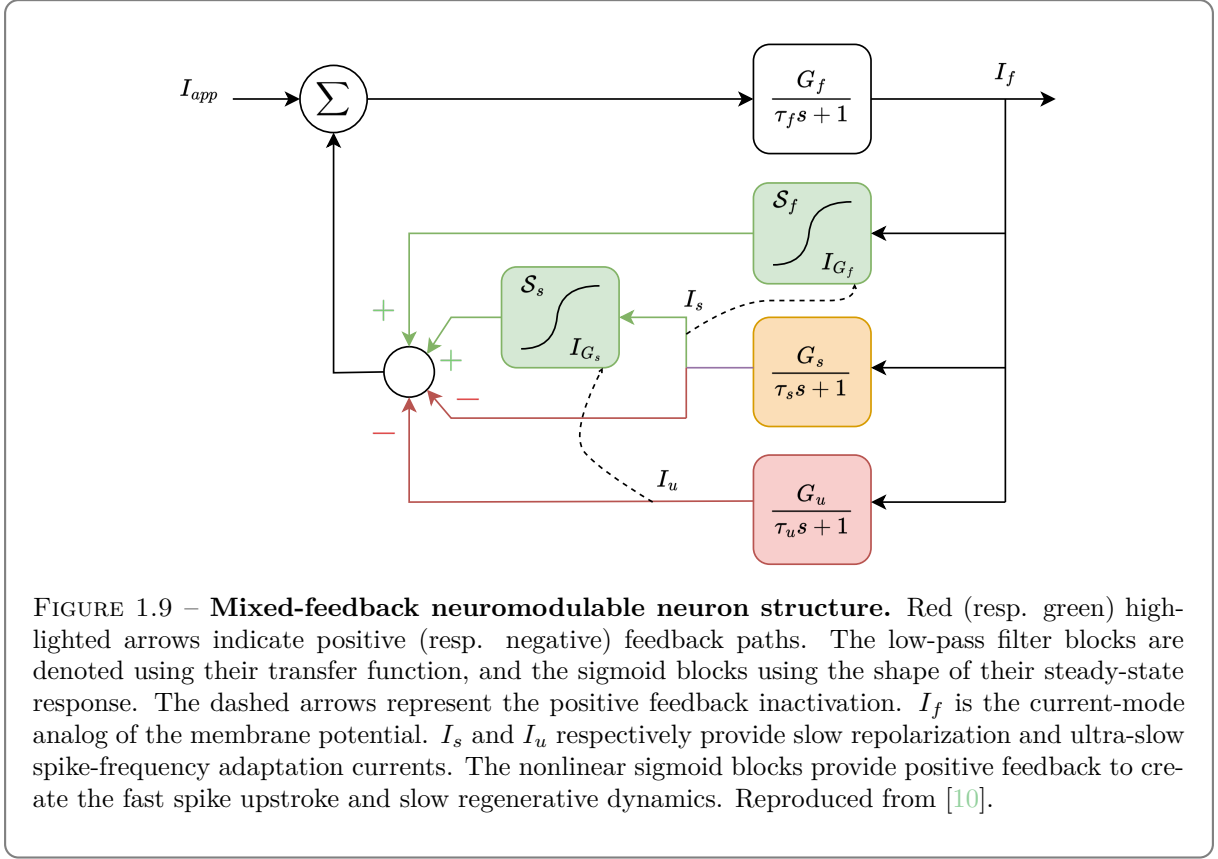
$$\tau I_{out}s = -I_{out} + GI_{in} \quad (1.10)$$

which is represented by the transfer function

$$H(s) = \frac{I_{out}}{I_{in}} = \frac{G}{\tau s + 1} \quad (1.11)$$

The whole structure of the current-mode mixed-feedback model is represented in Figure 1.9. In order to mimic at best biological neurons, three timescales are used (fast, slow and ultra-slow) as well as two linear negative feedback loops and two nonlinear positive ones [24].

The fast low pass-filter models the passive membrane of the neuron. The applied current I_{app} and all feedback loops, including a self-positive feedback through the sigmoidal block $\mathcal{S}_f$, are integrated in this leaky integrator to output I_f , the equivalent of the membrane voltage in current-mode. The slow low-pass filter sets the slow current-mode state variable I_s that provides both a nonlinear slow positive feedback (through $\mathcal{S}_s$) and a linear slow negative feedback. Finally, the ultra-slow low-pass filter identifies the ultra-slow state variable I_u providing a linear ultra-slow negative feedback.



These feedback loops form the basis of the dynamics emulated by ionic currents underlying membrane potential fluctuations in biological neurons. The fast depolarization arising from rapid positive feedback is associated with fast transient sodium channels, whereas the slow negative feedback corresponds to delayed rectifier potassium channels, leading to slow repolarization. As explained in Section 1.1, they are essential for spike generation, while slow depolarization leads to bursting and ultra-slow repolarization to adaptation currents.

In addition, the model incorporates a positive feedback inactivation mechanism to obtain slimmer, shorter and more power-efficient spikes. This mechanism is modeled by dashed arrows in Figure 1.9. In practice, the gain of the fast (resp. slow) nonlinear positive feedback is obtained by subtracting the slow (resp. ultra-slow) current to the fast (resp. slow) sigmoid gain $I_{G_f,0}$ (resp. $I_{G_s,0}$). This gives the modulated sigmoid gains

$$I_{G_f}(t) = I_{G_f,0} - I_s(t) \quad (1.12)$$

$$I_{G_s}(t) = I_{G_s,0} - I_u(t) \quad (1.13)$$

As the fast dynamics are much faster than the other ones, the inactivation mechanism has an impact mainly during the repolarization period where it attenuates the positive feedback and allows the neuron to repolarize faster with the consequences of a lower power consumption per spike and narrower events, thus more localized in time (Figure 1.10).




FIGURE 1.10 – **Inactivation mechanism.** Spike without (A) and with (B) positive feedback inactivation. With the inactivation mechanism the spike is narrower in time but still long enough to signal to physical systems. Reproduced from [10].

The current-mode building blocks

Differential-pair integrator (DPI) circuits are used to implement the current-mode low-pass filters necessary in the mixed-feedback architecture. The circuit used is depicted in Figure 1.11. The DPI is common in analog mixed-signal neuromorphic systems as it offers tunable gain and timescale along with low-power operation. It behaves closely to a first-order low-pass filter under the hypotheses that, in addition to subthreshold operation, the input current exceeds the leakage current ($I_{in} > I_\tau$) so that the output current largely overcomes the threshold current ($I_{out} \gg I_{th}$). Then, the filter parameters in (1.9) are given by

$$G = \frac{I_{th}}{I_\tau} \quad (1.14)$$

$$\tau = \frac{CU_T}{\kappa I_\tau} \quad (1.15)$$

with $I_\tau = I_0 e^{\frac{\kappa V_T}{U_T}}$, where I_0 is the transistor leakage current, κ is the subthreshold slope factor, and U_T is the thermal voltage. In the DPI, the different timescales needed in the mixed-feedback model are expressed as different time constants spanning several orders of magnitude, implemented through the scaling of the capacitors. They use capacitors ranging from hundreds of fF to tens of pF, resulting in timescales from milliseconds (fast timescale) to hundreds of milliseconds (ultra-slow timescale). These capacitors are required for timescales similar to biology, even though they are costly in terms of space on the neuron chip area.

The circuit implementing the two non-linear sigmoids $\mathcal{S}_f(I_f; I_{G_f})$ and $\mathcal{S}_s(I_s; I_{G_s})$ is a custom current-mode circuit shown in Figure 1.12. It brings together a comparator with diode-connected transistors and a gain branch for a fully controllable sigmoidal relationship. Three bias voltages independently tune the input-output characteristic:




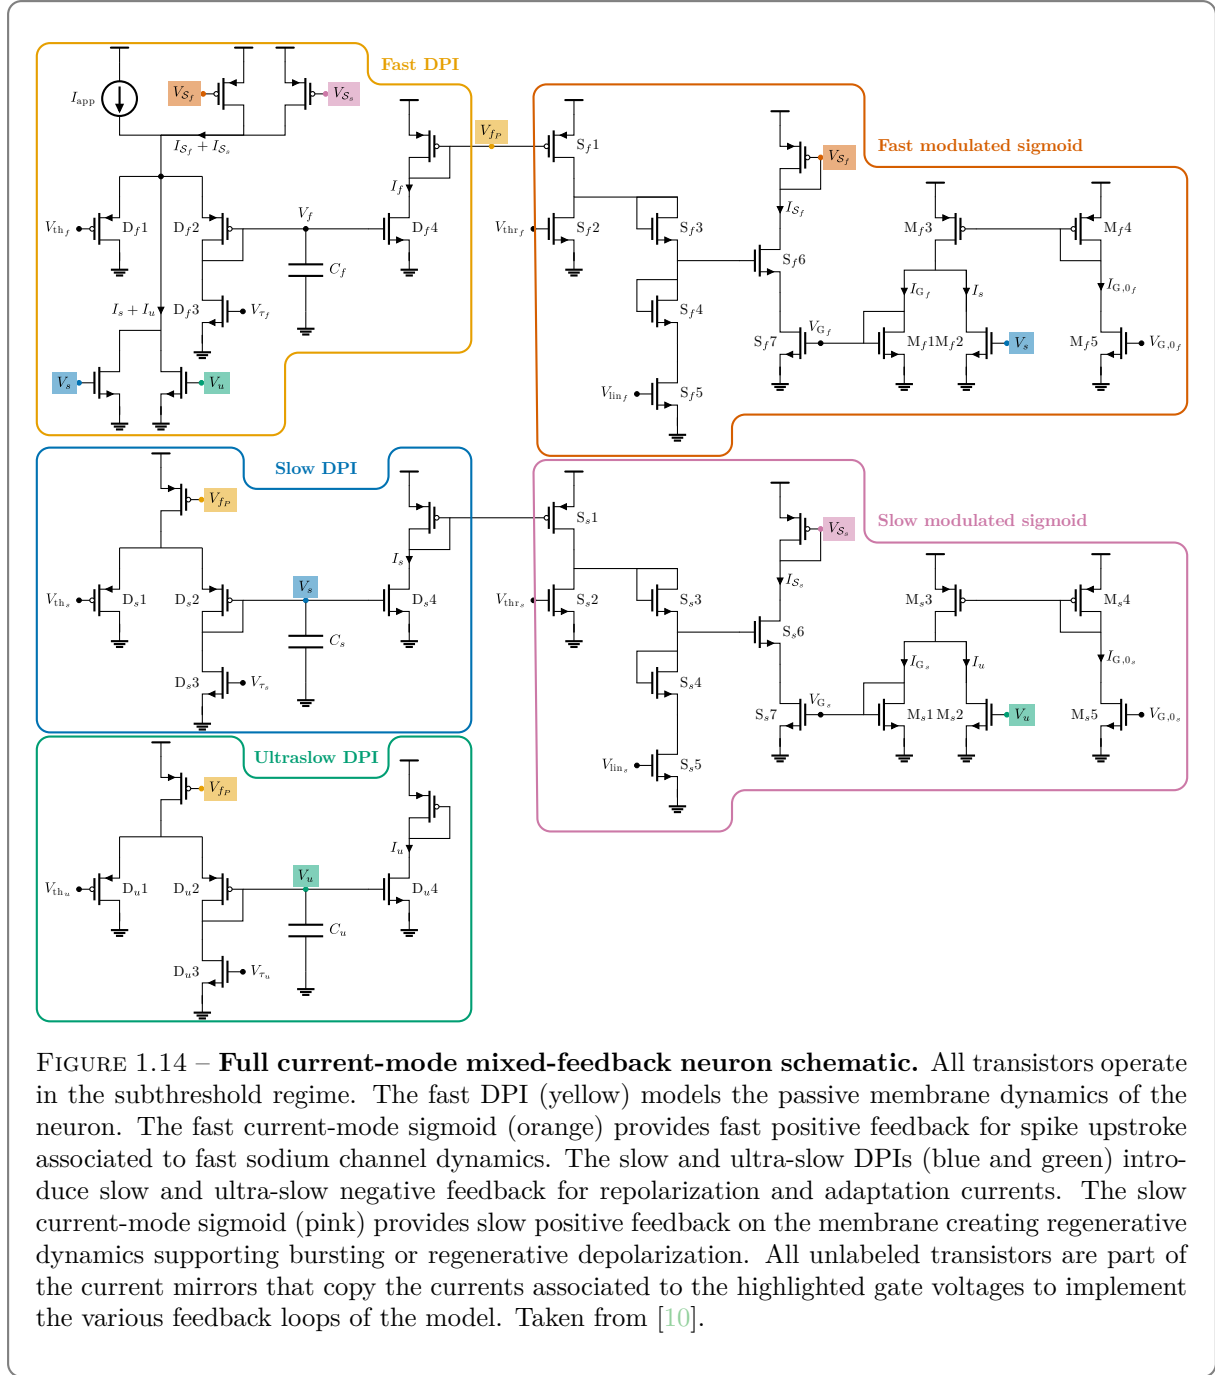
FIGURE 1.13 – **Current-mode sigmoid circuit response.** For a set of bias parameter values: I_{thr} sets the threshold, I_{lin} sets the width of the increasing range and I_G sets the maximum achievable output current. Reproduced and adapted from [10].

- V_{thr} determines the input threshold current I_{thr} at which the circuit starts outputting current,
- V_{lin} establishes the range in input current where the output is roughly linearly and monotonically increasing,
- V_G sets the maximum output current.

The impact of these three bias currents on the input-output curve is exemplified in Figure 1.13. The circuit also includes a sub-circuit to implement the positive feedback inactivation mechanism using gain modulation. The inactivation current I_{mod} (either I_s or I_u) is subtracted from the baseline gain current $I_{G,0}$ using current mirrors. This results in a bias voltage V_G for the gain branch producing a gain current $I_G(t) = I_{G,0} - I_{mod}(t)$, giving (1.12) and (1.13) for the fast and slow sigmoids respectively.

Using these two building blocks allows the implementation of the entire current-mode mixed-feedback model, with the full schematic presented in Figure 1.14. All state variables are currents and interact through linear DPI filters and current-mode summation, leading to a model invariant to scaling of the bias currents, as long as the system operates in the subthreshold regime. Therefore, biasing the neuron in picoamperes, nanoamperes, or even higher ranges will not affect its qualitative behavior. For the same reasons, subthreshold currents scale with temperature, ensuring proportional changes of current and thus preserved relative ratios, leading to qualitative invariance with temperature (rescaling in temporal evolution).

The choice of the bias range is the result of a balance between energy efficiency and robustness to mismatch. Higher currents reduce mismatch effects as the relative error due to mismatch becomes smaller. This comes at the cost of higher power consumption, while decreasing the currents favors ultra-low-power applications but increases sensitivity to mismatch.



Model analysis

The continuous-time dynamics of the mixed-feedback model are expressed using the structure of the model. Writing the input-output relationships of the low-pass filters gives:

$$\begin{aligned}\tau_f \dot{I}_f &= -I_f + G_f(\mathcal{S}_f(I_f; I_{G_f}) + \mathcal{S}_s(I_s; I_{G_s}) - I_s - I_u + I_{app}) \\ \tau_s \dot{I}_s &= -I_s + G_s I_f \\ \tau_u \dot{I}_u &= -I_u + G_u I_f\end{aligned}\tag{1.16}$$

Using these equations, the technique introduced by Ribar et al. [41] for voltage-mode mixed-feedback circuits can be applied. It provides robust and flexible control of excitability by shaping the circuit input-output I-V curves across multiple timescales, thereby revealing the interaction between the different feedback loops. This results in a method for analyzing and controlling the emergence of spiking and bursting behaviors through fast parameter tuning. This further supports the neuromodulation capabilities offered by this neuron implementation.

Experimental results

To test the model, an array of 16 mixed-feedback neurons was fabricated on a test chip in a 180 nm CMOS X-FAB process, with an active area of $750 \times 90 \mu\text{m}^2$. The dimensions of all transistors are $W \times L = 3 \mu\text{m} \times 300 \text{nm}$, not yet optimized for mismatch. All neurons share the same set of 12 bias voltages, and the output is delivered through a transmission gate multiplexer controlled by on-chip shift registers. In this way, one of the internal currents (I_f , I_s , or I_u) of a selected




FIGURE 1.15 – **Experimental result to an increasing input.** Measured neuron responses to increasing input current (bottom), in spiking (top) and bursting (middle) configurations. The neuron exhibits tonic excitability, showing firing frequencies increasing with input current in both modes, starting as low as a few Hz. Taken from [10].

neuron is transmitted to an output buffer stage, which amplifies an analog copy of the desired current for the output pad.

The chip is mounted on a test printed circuit board (PCB) (Figure 4.1), where a digital-to-analog converter (DAC) (Analog Devices AD5674, 12-bit, 16-channel) is used to generate the bias voltages. The DAC and the shift registers are configured by an Arduino Nano Every, which can also optionally stream output data. Otherwise, the current is converted into a voltage using two amplification stages, after which it can be measured with an oscilloscope, the Analog Discovery 3. The nominal supply voltage is 1.8 V.

The neuron response to an increasing input current, in both spiking and bursting regimes, is shown in Figure 1.15. This behavior is consistent with biological neurons, as both the spiking and bursting frequencies increase with the applied current. Similarly to biological systems, in the bursting regime and for sufficiently large input current, the neuron transitions from bursting to fast spiking.

Using the slow positive feedback gain I_{G_s} as a neuromodulatory input highlights the neuromodulation capabilities of the neuron (Figure 1.16). As this current increases, the neuron transitions from tonic spiking to tonic bursting, where both the burst duration and the number of spikes per burst increase with I_{G_s} . This demonstrates the strong neuromodulatory capabilities of this model using a single bias current. This behavior is fully consistent with the slow positive feedback theory: increasing the slow positive gain enhances the contribution of the slow current to spike and burst generation without altering the fast positive feedback responsible for spike initiation. The importance of slow positive feedback is also consistent with biophysical conductance-based models [42].

The ultra-low-power advantage of this silicon neuron was also verified through SPICE simulations of its power consumption. Overall, the neuron operates between 40 pJ/spike (high-frequency tonic spiking) and 200 pJ/spike (low-frequency tonic spiking).




FIGURE 1.16 – **Experimental result to modulation.** Measured neuromodulation from spiking to bursting (top) through an increase in slow positive feedback $I_{G_{s,0}}$ (bottom). The neuron undergoes a smooth transition from spiking to bursting, with a gradual increase in the number of spikes per burst and a consistent inter-burst interval. Taken from [10].

Finally, the robustness of the neuron to temperature variations, discussed in Section 1.4.2, was experimentally validated using a climatic chamber. As expected, in both the spiking and bursting regimes, the frequency increases with temperature, while the qualitative behavior remains robust and largely unaffected.

Optimized layout

The prototype used for the experimental phase, which integrates transistors of identical size, allowed the validation of the architecture and the expected behavior of the current-mode mixed-feedback neuron. However, it is highly sensitive to mismatch, as neither the transistor sizing nor the overall circuit design accounted for mismatch effects. Therefore, a second design was developed, incorporating transistor sizing and layout techniques aimed at reducing mismatch.

For this new design, 200-point Monte Carlo mismatch simulations of the neuron were performed for both spiking and bursting bias conditions in the picoampere range, under constant input. The simulations were then repeated after increasing the bias current to the nanoampere range and scaling the input to enable frequency comparison. The resulting distributions of spiking and bursting (inter-burst) frequencies are shown in Figure 1.17. As expected, higher currents lead to fewer transitions between firing regimes and reduced frequency dispersion within a given regime. These results also confirm that mismatch effects at low currents (picoampere range) result in transitions between spiking and bursting, as well as a large dispersion in frequency.

Although these results were obtained using the optimized layout model, it is worth noting that the first design would exhibit the same qualitative differences between the two current




FIGURE 1.17 – **Current regime impact on mismatch.** Monte Carlo mismatch analysis ($N = 200$) results of nominal spiking and bursting configurations biased in the picoampere range and in the nanoampere range. Violin plots depict the kernel density of the frequency distribution for spiking and bursting neurons. The green bar marks the median, the diamond indicates the mean, the box spans the interquartile range (IQR), and the whiskers extend to $1.5 \times \text{IQR}$, setting the bounds for outliers shown as hollow circles. At ultra-low currents, mismatch leads to regime mixing and larger frequency variability. Increasing bias currents restores full yield and reduces frequency dispersion. Taken from [10].

regimes, with more frequent transitions and larger distributions due to stronger mismatch effects. As the optimized chip is still under development, the initial design was used in this thesis.

CHAPTER 2

MODEL-BASED CONTROL FOR A SILICON NEURON

In this chapter, the Julia model of the neuron presented in Section 1.4.2 will be used to build a controller for robust neuromodulation of the CMOS neuron. After a brief description, since the model does not include mismatch, it will be exploited to extract a relationship between the neuron activity and the dynamic input conductances. This link will then be the foundation for the design of the controller, allowing to directly control excitability, instead of raw electrical variables. Finally, some artificial mismatch will be added in the model to test its theoretical robustness.

2.1 Model description

The model implements the structure presented in Figure 1.9 using Julia language, a language combining speed of compiled languages like C with the simplicity of high-level languages like Python [43].

As the model is invariant to current scaling without mismatch, it is biased in the nanoampere range, although this choice of scale is arbitrary. The bias current used for neuromodulation is I_{gain_s} , corresponding to the baseline gain of the slow sigmoid, as explained by Mendolia et al. [10]. Unless otherwise stated, the default values of the other biases, allowing neuromodulation via the neuromodulatory bias, are summarized in Table A.1, where the sigmoids shaping currents are I_{thr_f} , I_{lin_f} , I_{gain_f} for the fast sigmoid and I_{thr_s} , I_{lin_s} , I_{gain_s} for the slow one.

A first test was conducted to explore the different firing patterns obtained by varying the neuromodulatory bias. The results are shown in Figure 2.1. As expected, starting from a state of no activity, increasing the bias leads to type I tonic spiking: the frequency is initially low and increases with the bias. The activity then transitions to bursting, with the number of spikes per burst increasing until tonic spiking reappears, now of type II. The spiking frequency then continues to increase. Unlike electronic circuits, the model does not exhibit upper saturation.




FIGURE 2.1 – **Model output examples.** Examples of the model simulation output for different values of I_{gain_s} . With $I_{gain_s} = 225$ nA resulting in spiking type I (left), $I_{gain_s} = 250$ nA resulting in bursting (middle) and $I_{gain_s} = 320$ nA resulting in spiking type II (right).




FIGURE 2.2 – **Output firing pattern.** Classification of the neuron output activity depending on the applied input current and the slow sigmoid gain. The green area corresponds to no activity, while other areas are annotated with the either spiking type I (pink), bursting (blue) or spiking type II (orange).

However, the current must remain low to ensure operation within the subthreshold regime of the circuit analogy.

The applied current also impacts the output activity, as shown in Figure 2.2. In this section, the applied current is kept constant at 400 nA to ensure that all firing patterns are reachable. A potential next step toward robust neuromodulation would be to incorporate the effect of the applied current into the controller, thereby expanding the range of achievable behaviors.

2.2 Slow DIC for excitability quantification

The goal of this thesis is to control the firing pattern output by the neuron. To this end, instead of feeding the controller with the whole neuron activity recording, a value should be extracted that quantifies the firing pattern as precisely as possible. In this context, the dynamic input conductances, introduced in Section 1.2, will be exploited.

The DICs as described by Drion et al. [24] can be computed for the current-mode mixed-feedback neuron, using I_f as the analog of the membrane voltage, with

$$g_f = -\frac{\partial \dot{I}_f}{\partial I_f} \cdot \frac{\partial \dot{I}_{f\infty}}{\partial I_f} \quad (2.1)$$

$$g_s = -\frac{\partial \dot{I}_f}{\partial I_s} \cdot \frac{\partial \dot{I}_{s\infty}}{\partial I_f} \quad (2.2)$$

$$g_u = -\frac{\partial \dot{I}_f}{\partial I_u} \cdot \frac{\partial \dot{I}_{u\infty}}{\partial I_f} \quad (2.3)$$

Using (1.16),

$$g_f = 1 - G_f \left. \frac{\partial \mathcal{S}_f}{\partial I_f} \right|_{I_s=I_f} \quad (2.4)$$

$$g_s = G_f \left(1 - \left. \frac{\partial \mathcal{S}_f}{\partial I_s} \right|_{I_s=I_f} - \left. \frac{\partial \mathcal{S}_s}{\partial I_s} \right|_{I_s=I_u=I_f} \right) \quad (2.5)$$

$$g_u = G_f \left(1 - \left. \frac{\partial \mathcal{S}_s}{\partial I_u} \right|_{I_s=I_u=I_f} \right). \quad (2.6)$$

The DIC curves obtained with the computational model of the neuron, biased for bursting, are shown in Figure 2.3. The positive feedback responsible for spike upstroke can be clearly identified in the orange curve. The slow and ultra-slow negative feedbacks are less clear because of computational limits. Nonetheless, it can be seen that both the slow and ultra-slow DICs increase to a positive value (negative feedback gain) right after the spiking threshold, which is the current at which the fast DIC starts decreasing. Importantly, the slow DIC decreases sharply before the threshold current, corresponding to the slow positive feedback essential for bursting. This part of the slow DIC curve near the spike threshold controls the transition from slow restorative (spiking) to slow regenerative (bursting) excitability [24, 44].

To implement the control of the neuron activity easily, a single value should be selected for each DIC. Inspired by Fyon et al. [25], the threshold value should be such that $g_{in}(I_{th}) = g_f(I_{th}) + g_s(I_{th}) + g_u(I_{th}) < 0$, while ensuring that $g_{in}(I_{th} - \delta I) \geq 0$ for any arbitrary $\delta I > 0$. This reduces to finding the current where the total DIC crosses the zero line. After simulations with the model, this threshold was found to be very close to the spike threshold; therefore, for the sake of computation, the fast sigmoid threshold was used, $I_{thrf} = 150$ nA. The DICs at the threshold current $g_f(I_{th})$, $g_s(I_{th})$ and $g_u(I_{th})$ will now be referred to as g_f , g_s , and g_u respectively.




FIGURE 2.3 – **DICs computed with the neuron model.** Dynamic input conductances versus I_f , computed with the mixed-feedback neuron model in bursting ($I_{gain_s} = 250$ nA). The fast DIC (orange) indicates a fast positive feedback responsible for spike upstroke, the slow DIC (blue) shows a narrow slow positive feedback leading to bursting and the ultra-slow DIC (green) exhibit an ultra-slow negative feedback for adaptation. The dotted line is the vertical line at the threshold current $I_{thr_f} = 150$ nA.




FIGURE 2.4 – **Evolution of the DICs with modulation.** Values of the three DICs at the threshold current plotted against the neuromodulatory bias I_{gain_s} .

To explore the evolution of the DICs through neuromodulation, these values are plotted as a function of the gain of the slow sigmoid. The result is shown in Figure 2.4, where it can be seen that while both g_f and g_u remain approximately constant, the slow DIC decreases monotonically. This demonstrates that the value of the slow DIC, g_s , is sufficient to uniquely quantify the neuron excitability, and that only this DIC is needed to implement the controller. By combining Figures 2.2 and 2.3, for a fixed applied current, we obtain $|g_{s_{spiking\ I}}| < |g_{s_{bursting}}| < |g_{s_{spiking\ II}}|$.

2.3 Control loop design

2.3.1 The PI controller

To control the neuron and achieve the desired activity, a feedback control system is used. This loop consists of measuring the output to estimate the slow DIC, comparing it to the desired excitability type, and adjusting the system based on the computed error. This error is generally defined as the difference between the desired output and the current output. Hence, the goal of a controller is to reduce this error over time.


Among controllers, the proportional-integral (PI) controller is one of the most commonly used. The general formula for a PI controller is the following:

$$u(t) = K_p \left[e(t) + \frac{1}{T_i} \int_0^t e(\tau) d\tau \right] \quad (2.7)$$

$$= K_p \cdot e(t) + K_i \int_0^t e(\tau) d\tau \quad (2.8)$$

where $u(t)$ is the control input, $e(t)$ is the error, K_p is the proportional gain, and $K_i = K_p/T_i$ is the integral gain [45]. The proportional term allows a fast response to the current error, but alone it does not ensure zero steady-state error. The integral term responds to the accumulated error over time, which eliminates the steady-state error. Increasing K_p improves the speed of the response but can create oscillations, while increasing K_i improves accuracy but can also lead to overshoot. Hence, the controller gains require tuning to find a good balance.

The complete structure of the CMOS neuron controlled by the PI controller is represented






FIGURE 2.6 – **PI controller.** Structure of the PI controller including saturation and anti-windup with back calculation with K_p the proportional gain, K_i the integral gain, $K_w = K_i/K_p$ the windup gain and I'_{gain_s} the control input, corresponding to I_{gain_s} after saturation.

in Figure 2.5. The input to the controller is the error computed as

$$e(t) = |g_{sref}| - |\hat{g}_s| \quad (2.9)$$

where g_{sref} is the desired slow dynamic input conductance associated with the neuron desired firing pattern, and $\hat{g}_s$ is the slow DIC estimated from the neuron output activity through a mapping from spikes to DIC estimate. This mapping is developed in Section 2.3.2 and will be referred as spike2DIC for simplicity. The control signal $u(t)$ of the PI controller is the gain of the slow sigmoid I_{gain_s} . The absolute values applied to the reference and the estimated slow DIC are required because the modulated gain should increase if $g_{sref} < \hat{g}_s$. The absolute value notation will be omitted in the remaining of this thesis for notation simplicity.

Due to the limitations of the model, the value of I_{gain_s} has a minimum value of $I_{gain_{s_{min}}} = 210 \text{ nA}$, corresponding to lower saturation. However, when saturation is present in a PI controller, the system may attempt to reach a value outside the allowable range, causing the integral term to continue accumulating the error. This can lead to an excessive integral term and, consequently, to overshoot, sluggish recovery, or sustained oscillations [46], this is called wind-up. Several anti-windup strategies exist; the one chosen here is the tracking back-calculation method, based on the difference between the controller output and the actual control signal after saturation. According to Li et al. [46], the anti-windup gain K_w is chosen as K_i/K_p . The final structure of the controller is presented in Figure 2.6.

The control loop can be summarized in a few steps, starting from $I_{gain_s} = 210 \text{ nA}$:

- a) the neuron is simulated, and its activity is recorded over 3 s;
- b) from the recorded activity, the estimated DIC is extracted using the spike2DIC mapping;
- c) this estimate is compared to the desired slow DIC, and the resulting error is fed to the controller;




FIGURE 2.7 – **Activity quantification for spiking type I.** Computed slow DIC value plotted against the inter-spike frequency for a range of I_{gain_s} leading to spiking type I.




FIGURE 2.8 – **Activity quantification for bursting.** Computed slow DIC value plotted against the intra-burst frequency for a range of I_{gain_s} leading to bursting.

- d) the output of the controller is the updated slow sigmoid gain, which is then used to simulate the neuron again, and the process repeats.

2.3.2 Spike to DIC estimate

The feedback control loop utilizes the slow DIC, however, this metric is not directly output by the neuron and must be extracted from the firing pattern. To this end, for each type of spiking activity, a curve linking g_s to a characteristic of the activity has been obtained. The methodology is the same in all cases: for a range of I_{gain_s} values leading to the desired firing pattern, both the slow DIC and a metric quantifying this activity are computed and plotted against each other.

- **Spiking type I:** g_s is monotonically increasing with the inter-spike frequency (ISF) (Figure




FIGURE 2.9 – **Activity quantification for bursting.** Computed slow DIC value plotted against the number of spikes per burst for a range of I_{gain_s} leading to bursting.




FIGURE 2.10 – **Activity quantification for spiking type II.** Computed slow DIC value plotted against the inter-spike frequency for a range of I_{gain_s} leading to spiking type II.

2.7). Therefore, if the activity is classified as spiking type I, the spiking frequency is computed, and the corresponding estimated g_s is obtained by looking it up on this curve.

- **Bursting:** the behavior observed when increasing I_{gain_s} is more complex. As shown in Figure 2.8, the intra-burst frequency exhibits bumps corresponding to the addition of a new spike per burst, which is confirmed in Figure 2.9. For a fixed number of spikes per burst, g_s increases as the intra-burst frequency increases. Therefore, in the bursting regime, g_s is evaluated using only the intra-burst frequency.
- **Spiking type II:** similarly to spiking type I, g_s is monotonically increasing with the inter-spike frequency (Figure 2.10). Once again, the slow DIC value is evaluated from the activity using this curve and the computed spiking frequency.

The spike2DIC block first identifies the excitability type among spiking type I, spiking type II, or bursting, and then refers to the corresponding curve to estimate g_s . The classification of




FIGURE 2.11 – **Controller step responses, good balance gains.** Step response of the controller for different step sizes using $K_p = 10^{-9}$ and $K_i = 5 \cdot 10^{-9}$. The initial value of I_{gain_s} is 210 nA resulting in no activity and thus $g_s = 0$. Multiple steps are tested in each activity target, the final activity is illustrated on the right. A fixed number of 100 iterations is used.

the firing pattern is based on a simple rule: if the ratio between the maximum and minimum inter-spike intervals is greater than 2, the activity is classified as bursting. Otherwise, if the mean inter-spike frequency is greater than 20 Hz, it is classified as spiking type II; in all other cases, it is classified as spiking type I.

2.3.3 Controller gains tuning

This section briefly explains how the controller gains are selected for this model-based study. The controller structure, as well as its gains, may need to be adjusted when using the physical chip, as electronic limitations will be encountered.

Since the system is nonlinear, Bode and Nyquist plots are not applicable, and the gain selection is based on a qualitative analysis of the controller response. After several tests, the best results were obtained with $K_p = 10^{-8}$ and $K_i = 5 \cdot 10^{-9}$, yielding $K_w = 1/2$. These gains produce the step responses shown in Figure 2.11. Using smaller gains results in Figure 2.12, where the response is too slow and does not converge in 100 steps. On the other side, higher gains lead to oscillations for some step responses, as depicted in Figure 2.13.

The order of magnitude of K_p and K_i is justified by the difference in magnitude between the slow DIC values and the gain current used for modulation. It can be observed that the discontinuities in Figure 2.8 lead to small oscillations in the step response. Although hysteresis could be introduced to address this issue, it is not considered in this chapter, which focuses on understanding the system.




FIGURE 2.12 – **Controller step responses, too low gains.** Step response of the controller for different step sizes using $K_p = 5 \cdot 10^{-9}$ and $K_i = 10^{-9}$. The initial value of I_{gain_s} is 210 nA resulting in no activity and thus $g_s = 0$. Multiple steps are tested in each activity target, the final activity is illustrated on the right. A fixed number of 100 iterations is used.




FIGURE 2.13 – **Controller step responses, too high gains.** Step response of the controller for different step sizes using $K_p = 2 \cdot 10^{-8}$ and $K_i = 10^{-8}$. The initial value of I_{gain_s} is 210 nA resulting in no activity and thus $g_s = 0$. Multiple steps are tested in each activity target, the final activity is illustrated on the right. A fixed number of 100 iterations is used.

Before introducing variability in the model, the controller was tested for different slow DIC targets. Some output activity are displayed in Figure 2.14, to validate the controller on the neuron model.




FIGURE 2.14 – **Examples of activity obtained with control.** Firing pattern obtained after 100 steps of control with the neuron model for different values of the slow DIC reference (g_s). Depending on the reference, spiking I, bursting or spiking II is obtained. Both the spiking frequency and the number of spikes per burst depend on g_s .

2.4 Robustness to variability caused by mismatch

As explained in Section 1.3, two equally designed neurons do not behave exactly the same in practice due to mismatch in transistors. This means that no silicon current-mode mixed-feedback neuron operates exactly like the Julia model. In this part, an artificial mismatch on some transistors is imposed to test the controller robustness. To this end, few bias currents are randomly modified with

$$I' = I \cdot (1 + m) \quad \text{with} \quad m \sim \mathcal{N}(0, 0.075^2). \quad (2.10)$$

The currents concerned are I_{thr_f} , I_{thr_s} , I_{lin_f} , I_{lin_s} , I_{gain_f} . In addition, the slow sigmoid gain I_{gain_s} is also altered when the controller is not used, to allow a better visualization of the variability caused by mismatch. These biases have been chosen because they are fed to the model directly as currents compared to the time constants and gain of the DPIs. The mismatch is modeled using a Gaussian distribution, which best represents variations in the model parameters, with the mismatch magnitude expressed in terms of variance [32]. A value of 7 % was chosen arbitrarily, but is sufficiently large to represent significant mismatch.

The robustness of the controller is therefore evaluated by running and analyzing control results across several realizations of mismatch. The results for the three excitability-type targets are displayed in Figure 2.15. Due to mismatch, for the same target value g_s , the steady-state

value of the bias current varies across realizations. The same firing pattern is achieved, but the bias current must compensate for mismatch. This demonstrates the robustness of the proposed controller to substantial mismatch in a few transistors, while robustness to real mismatch will be assessed later using the electronic chip itself.

Another way to demonstrate the controller robustness to mismatch is to compare the distribution of activity metrics between two cases: neurons simulated with mismatch using fixed parameters, and neurons simulated with mismatch and in a feedback control loop. The values of the biases in the first case are chosen such that, without variability, they would lead to the activity that the controller will try to reach. These values and the targets used in parallel in the control feedback in this section are listed in Table A.2.

For the inter-spike frequency in spiking types I (Figure 2.16) and II (Figure 2.17), the fixed-parameter case shows a spread of more than 1 Hz between the lowest and highest frequencies for type I, and a similar spread for type II. Although this dispersion may appear small, it should be noted that mismatch is applied to a subset of the transistors and other electronic constraints are disregarded. In contrast, the controller yields an almost uniform frequency across neurons in type I, and a very narrow range (about 0.3 Hz) for type II.

In bursting mode, the analyzed metrics are the intra-burst frequency and the number of spikes per burst, as shown in Figure 2.18. Without control, both distributions are quite broad: several hertz for the intra-burst frequency and 10 to 16 spikes per burst. When control is applied, the frequency distribution becomes narrower, around one hertz, although it is not completely flat. This is mainly due to the conversion of activity into slow DIC being less precise in bursting mode than in spiking mode. Moreover, due to the artificial mismatch, the reachable ranges of frequency and spike count per burst may differ from those obtained with the perfect model. Regarding the number of spikes, applying control reduces the range to 10 to 13 spikes. This distribution remains relatively broad for the same reason as the intra-burst frequency, with the additional factor that the estimated g_s is derived from the frequency rather than from the number of spikes.




FIGURE 2.15 – **Step response over mismatch realizations.** Evolution of I_{gain_s} and the estimated g_s along the step response of the controller for different realization of mismatch, with a fixed number of 100 steps. The reference is chosen to achieve either spiking type I (A), bursting (B) or spiking type II (C).




FIGURE 2.16 – **ISF distributions for spiking type I.** Violin plot of the distribution of the inter-spike frequency with (green) and without (pink) control for spiking type I, with an illustration of the activity after control.




FIGURE 2.17 – **ISF distributions for spiking type II.** Violin plot of the distribution of the inter-spike frequency with (green) and without (pink) control for spiking type II, with an illustration of the activity after control.




FIGURE 2.18 – **Bursting parameters distributions.** Violin plot of the distribution of bursting parameters with (green) and without (pink) control. **A:** intra-burst frequency. **B:** number of spikes per burst. Two illustrations of the activity after control are given, as the controller does not always reach the same activity characteristics.

Before moving to hardware, it is interesting to estimate actual mismatch using transistor level simulations. To that end, the software Cadence Virtuoso is used, enabling Monte Carlo analysis of mismatch on the desired transistors, using the optimized transistor sizes. In the first part of this chapter, the total impact of mismatch is evaluated, with mismatch applied to all transistors. Then, mismatch on individual transistors is investigated to identify the most sensitive devices.

3.1 Mismatch impact on the neuron output

When mismatch is applied to all transistors of the neuron, the output varies accordingly. Cadence allows the simulation of multiple mismatch realizations, where each transistor characteristic follows a normal distribution under manufacturer characterization. This simulation approach, using repeatedly random sampling within the distribution to model uncertainty is the Monte Carlo method. The number of Monte Carlo samples chosen is 200.

The circuit is biased here in picoamperes to maximize mismatch impact, which is higher at lower currents, and the values used for the different modules shown in Figure 1.14 are summarized in Table A.3. The currents are used to generate the corresponding bias voltages using diode connected transistors, while the gains G_{fDPI} , G_{sDPI} , and G_{uDPI} determine V_{th_f} , V_{th_s} , and V_{th_u} , respectively.

The input signal is composed of two parts: a DC component fixed at 300 pA, and a time-varying component I_{pwl} , which changes according to the desired stimulation. The slow sigmoid gain $I_{G,0_s}$ is also adjusted during the simulations.

Three simulations were carried out, each focusing on a different metric:

- $I_{G,0_s} = 210$ pA and $I_{pwl} = 70$ pA after 10 ms (zero before): the output exhibits a type I

spiking firing pattern, and the inter-spike frequency is recorded using Monte Carlo analysis.

- $I_{G,0s} = 300$ pA and $I_{pwl} = 70$ pA after 10 ms (zero before): the output exhibits a type II spiking firing pattern, and the inter-spike frequency is recorded using Monte Carlo analysis.
- $I_{G,0s} = 300$ pA and I_{pwl} is a ramp from 0 to 200 pA over 500 ms: the neuron starts firing when the total input (DC + I_{pwl}) crosses a threshold called the rheobase. This metric is recorded using Monte Carlo analysis.

If the neuron never spikes, in either simulations, the sample is discarded. This gives rise to another metric, the yield, corresponding to the percentage of active neuron, i.e. the number of sample kept over the total number of samples.

The results of the three tests are shown in Figure 3.1. As can be seen, the inter-spike frequency (ISF) distribution in spiking type I is broader than 41 Hz and broader than 100 Hz in type II. Once again, since there is no distinction in activity, some outliers might be due to a change in firing regime. Regarding the rheobase, excluding the outliers around 200 pA, the computed values span a 100 pA range, which is of the same order as the bias voltages. These results illustrate how total mismatch can significantly alter the characteristics of the neuronal output, both in terms of timescales, leading to wide inter-spike frequency distributions, and in terms of threshold current, which impacts the rheobase range. Moreover, the inter-spike frequency distributions found in Section 2.4 are confirmed and are even broader when all transistors are affected.




FIGURE 3.1 – Mismatch impact on activity metrics. **A:** Inter-spike frequency histogram for spiking type I, with a yield of 42 %. The mean (red line) is 35.69 Hz and the standard deviation 10.70 Hz. **B:** Inter-spike frequency histogram for spiking type II, with a yield of 71.5 %. The mean (red line) is 138.77 Hz and the standard deviation 28.13 Hz. **C:** Inter-spike frequency histogram for the rheobase, the minimum input to produce a spike, with a yield of 94.5 %. The mean (red line) is 353.22 pA and the standard deviation 28.69 pA.

3.2 Individual mismatch and sensitive transistors

To identify the critical transistors in the circuit, mismatch is evaluated individually for each transistor. The setup uses a slow sigmoid gain of 200 pA, $I_{pwm} = 70$ pA, and other bias conditions as described in Table A.3. This configuration results in type I spiking, with the inter-spike frequency used as the reference metric.

The results for all transistors are presented in Table 3.1. From this analysis, the most critical transistors, those exhibiting the highest standard deviations, appear to be the bias transistors of all the DPIs, as well as the threshold bias transistors of both sigmoids. Regarding the DPI transistors, they determine the behavior of the analog low-pass filter. Mismatch in this module can lead to variations in the time constants and, in particular, in the output frequency. Moreover, mismatch effects are relatively larger at very low currents, such as those used in the DPIs (ranging from 10 to 100 pA).

On the other hand, the sigmoid threshold transistors play a crucial role in spike generation and, therefore, in the overall output behavior. Focusing on these two transistors, an additional test was performed. For each of them, still using Monte Carlo simulations, the mean current flowing through the transistor, as well as its standard deviation, were measured. A relative mismatch metric was then computed as the ratio of the standard deviation to the mean, yielding relative mismatch values of 3% and 2.8% for the threshold currents of the fast and slow sigmoids, respectively. This confirms the sufficiently large mismatch value of 7% used for the analysis in Section 2.4.

As the revised layout was used in this analysis, mismatch effects are expected to be worse with the initial chip design.

Transistor (see Figure 1.14)	biased	ISF mean [Hz]	ISF standard deviation [Hz]	Yield [%]
V_{τ_u}		33.87	16.29	78
V_{th_u}		28.35	6.87	91.5
V_{τ_f}		30.50	4.68	74
V_{thr_s}		30.15	4.41	74
V_{th_f}		28.73	3.73	86
V_{τ_s}		28.20	3.09	71
V_{thr_f}		28.19	2.78	98
V_{th_s}		27.94	2.72	91
V_{lin_s}		27.84	1.22	99
V_{G_s}		27.94	0.40	99.5
V_{G_f}		28.06	0.18	100
V_{lin_f}		28.08	0.10	100

TABLE 3.1 – Individual mismatch impact. Mean and standard deviation of the inter-spike spike frequency induced by mismatch on one transistor, along with the yield of the simulation. Lines are sorted in decreasing order of the mismatch-induced ISF standard deviation.

CHAPTER 4

NEUROMODULATION OF ULTRA-LOW-POWER NEURONS

Now moving to hardware, true mismatch effects and electronic constraints, such as saturation, will be encountered, where the controller designed in Chapter 2 will be adapted for real-life implementation. In this chapter, after a brief description of the setup, the mismatch effects on the neurons dynamics will be demonstrated and discussed. Then, an exploration of the different behaviors achievable by each neuron will be conducted. The feedback control loop will thus be integrated and results analyzed, to finish with a conclusion on neuromodulation of CMOS neurons operating at ultra-low power, i.e. in the picoampere range.

4.1 Set-up

The PCB used to interface the chip is represented in Figure 4.1, with a computer used for exploration first and then neuromodulation. Although the chip integrates 16 neurons, only a subset of six was used in this thesis. This choice was made to facilitate more efficient modulation and in-depth exploration of their behavior. Since all neurons share the same underlying architecture, the analysis of six representative units is sufficient to capture the relevant dynamics without introducing unnecessary repetition.

The neurons share the same set of biases, which are voltages directly related to the ones of the circuit depicted in Figure 1.14. Only one neuron and timescale can be recorded at a time. The fast timescale was chosen as it is the analog of the membrane voltage, while the choice of the neuron will vary. From the measured voltage, the corresponding current can be obtained using

$$I_f = V_f \cdot 10^{-8}. \quad (4.1)$$




FIGURE 4.1 – **PCB set-up.** Custom PCB used for characterization of the on-chip mixed-feedback neurons. The board hosts the fabricated CMOS die (center, yellow), an external 12-bit DAC (AD5674, red) for bias generation, and an operational amplifier configured as a feedback ammeter (blue) for on-board current-to-voltage conversion. A microcontroller (Arduino Nano Every, green) controls the on-chip multiplexer for output selection and programming of the DAC for bias generation. Output signals were acquired with an oscilloscope (Analog Discovery 3, purple) for data acquisition. A computer, (dark blue) acquiring output signal by the oscilloscope and setting biases and input through the microcontroller, is used for control. Taken and adapted from [10].

4.2 Impact of device mismatch on inter-neuron variability

In Section 3, the effects of mismatch on one neuron output and across neurons were investigated using Monte Carlo simulations on the neuron model. With the chip, integrating an array of these neurons, other mismatch sources can appear. Along with random local variations which lead one neuron behavior to differ from the perfect model and each neuron to differ from one another, edge as well as striation and gradient effects appear in an array of neuron (Section 1.3.3). These new effects introduce spatially structured variability in VLSI systems.

To explore the effects of mismatch-induced variability across integrated CMOS neurons, the chip was used to measure the output of multiple neurons biased with the same set of voltages. Recordings of 5 seconds for the first six neurons are shown in Fig. 4.2 and in Fig. 4.3 for two sets of bias values, shared across neurons, reported in Table A.4.

Both bias configurations result in currents in the picoampere range, which is favorable for ultra-low power operation but increases sensitivity to mismatch. The effects of mismatch are clearly visible: for the same set of biases, some neurons exhibit spiking activity, others display bursting behavior, while some remain silent. Consequently, the output firing patterns vary significantly from one neuron to another. Using all the neurons on the chip, the first set of biases results in 81.25 % of the neurons silent, 12.5 % spiking and 6.25 % bursting. For the second set of biases, 12.5 % are spiking, 68.75 % are silent and 18.75 % are saturated, similarly to neuron 1 and 5 in Figure 4.3. Together, these results confirm the output pattern variability

between neurons induced by the random modification of the dynamics of each neuron.




FIGURE 4.2 – **Inter-neuron mismatch effect.** Output measurement from six neurons biased with the same set of voltages reported in Table A.4 (set 1). Among the six neurons, three are silent (3,4,6), one is low-amplitude high-frequency spiking (1), one is low-amplitude low-frequency spiking (5) and one is bursting (2).

4.3 Inter-neuron mismatch consequences on neuromodulation

To simply and reliably control the firing pattern of a given neuron, it must be possible to span different excitability regimes by adjusting one or multiple parameters. In our framework, it should be possible using only the neuromodulatory bias $V_{G,0_s}$. However, exploration of neuronal




FIGURE 4.3 – **Inter-neuron mismatch effect.** Output measurement from six neurons biased with the same set of voltages reported in Table A.4 (set 2). Among the six neurons, two are silent (4,2), two seem saturated (1, 5), one is spiking at low frequency with spike extended in time (3) and one is spiking at low frequency (6).

behavior under multiple bias conditions shows that each neuron requires individual calibration with a specific set of bias values prior to control, as no universal bias configuration could be identified.

Among the first six neurons, only four were successfully calibrated to allow the full spanning of all excitability regimes, with the calibration parameters given in Table A.5. This highlights another limitation induced by mismatch: the possible loss of excitability. For the retained neurons (1, 2, 5, and 6), a neuromodulation analysis was performed. This enables the extraction

of possible firing patterns, including reachable firing rates and number of spikes per burst, for each neuron. The results are presented in Fig. 4.4 for neurons 1 and 2, and in Fig. 4.5 for neurons 5 and 6.

It can be observed that, even when calibrated, different neurons exhibit differences in their modulation capabilities. For example, neuron 1 does not appear to reach more than three spikes per burst, whereas neuron 2 displays stable bursting of five to six spikes per burst. This is an important consideration for output control, as certain target firing patterns may be unreachable for some neurons due to intrinsic mismatch constraints.

4.4 Compensation and exploitation of mismatch

Mismatch in VLSI systems is an intrinsic constraint that should be dealt with. To get rid of its effects, techniques such as optimized layout, higher current biases, and thorough calibration can be used. On the other hand, when mismatch effects are solely reduced, the remaining variability it offers can be exploited. As seen in Figures 4.4 and 4.5, after small compensation using calibration, mismatch offers a diversity in the firing pattern ranges the neurons can reach. While some neurons offers long range of spiking frequencies, other offer stable and constant bursting.

This intrinsic heterogeneity introduced by mismatch mirrors biological systems, where neuronal processes exhibit variability, as neurons of the same type can possess different properties. The functional role of this variability is not yet fully understood and remains an open area of research. However, several studies suggest that, rather than hindering performance, predictable and reliable behavior in neural systems may depend crucially on such variability, which can improve performance in terms of speed and responsiveness [47, 48].

In this thesis, the remaining differences between the neuromodulatory ranges of neurons after calibration are exploited. The controller is designed such that the target activity is selected from a set specific to each observed neuron. This approach avoids restricting the system to the subset of activities that are achievable by all neurons. Instead, the individual capabilities of each neuron are leveraged to provide greater flexibility in modulation.

4.5 Control loop integration for robust neuromodulation

4.5.1 Discrete slow DIC estimation for activity control

To reliably control the neuron output, the chip is placed in a control loop, where the output of one neuron at a time is shaped by the controller through the neuromodulatory gain $V_{G,0_s}$. The idea is to implement a feedback mechanism, using the true recorded activity of the selected neuron, similar to the model-based control described in Section 2.3. To achieve this, the neuron





FIGURE 4.4 – **Calibrated neuromodulation.** Transition from silent or low-frequency spiking to high-frequency spiking through bursting activity. Each neuron is individually calibrated to enable neuromodulation, and the range of the neuromodulatory input, labeled v_{gain_s} , is adapted accordingly. **Neuron 1:** The neuron exhibits type I spiking with approximately constant frequency, before transitioning to bursting activity with 2 then 3 spikes per burst, and finally to high-frequency spiking (type II) with increasing frequency. **Neuron 2:** The neuron initially exhibits low-frequency spiking with increasing rate, followed by bursting with an approximately constant number of spikes per burst (from 2 to 5), and finally transitions to high-frequency spiking.



	Spiking I				Bursting				Spiking II			
	2.5 Hz	4 Hz	6 Hz		2-3 spb	4-5 spb	6-7 spb	> 7 spb	25 Hz	35 Hz	60 Hz	
neuron 1												
neuron 2												
neuron 3												
neuron 4												
Reference	0.5	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0	5.5	6.0

FIGURE 4.6 – **Slow DIC discrete estimation.** Diagram dividing each activity type into its four subset depending either in the rate (spiking I and II) or the number of spikes per burst, spb (bursting). Each neuron output possibilities are marked by a colored box. The last line corresponds to the estimated slow DIC for each subset.

output must be quantified so that it can be compared to a reference and used to update the slow sigmoid gain accordingly. In Section 2.3, continuous values of the slow DIC at the threshold current were used to characterize the activity.

With the chip, however, neurons exhibit different dynamics, both between each other and compared to the model. As a result, the ranges of spiking frequency, number of spikes per burst, and intra-burst frequency are altered, and the curves from Figures 2.7, 2.9, 2.8, and 2.10 cannot be used directly. To address this issue, a discrete control approach is considered. Instead of quantifying the firing pattern using continuous values of g_s , subsets of activity are assigned to specific values of the slow DIC. Thus, the DIC estimation outputs a value from a discrete set, and the target activity is quantified using the same set. This categorization also helps prevent small oscillations in the control caused by perturbations in neuron dynamics due to environmental changes or time.

For this thesis, twelve subsets of activity have been selected, with four per activity type, providing a balance between control precision and stability. As previously discussed, variability between neurons leads to differences in their neuromodulation capabilities. Therefore, for each activity category, each neuron was tested to determine whether the desired pattern could be achieved. The results are presented in Figure 4.6, including activity categorization, neuron capabilities, and the discretely estimated slow DIC.

The last point that needs to be discussed is the conversion from membrane voltage recordings to activity characterization, namely, determining whether the neuron is spiking or bursting, as well as identifying the frequency and number of spikes per burst. Compared to the model-based approach, a new difficulty arises due to measurement noise added to the signal. This constraint is illustrated in Figure 4.2 for neuron 3, where, even when silent, the neuronal activity recording exhibits high-frequency oscillations corresponding to this noise. Using the same spike detection technique as in Chapter 2 leads to the detection of peaks in the noise, and thus falsely

indicates consistently high-frequency activity. To mitigate the effects of noise, several methods are employed.

The first step is to determine whether the neuron is silent. One way to do this is by comparing the variability of the signal to that of the noise: if they are sufficiently close, the neuron is considered inactive. The variability of the signal is computed using the classical standard deviation of the recorded membrane voltage. However, since noise is included in the recording, it cannot be directly estimated in the same way. Instead, the signal is filtered by retaining only values below the 50th percentile, thereby excluding part of the spikes. The variability is then computed using the Median Absolute Deviation (MAD), a measure that is more robust to outliers than the standard deviation [49]. The MAD is defined as:

$$\text{MAD} = 1.4826 \cdot M_i(|x_i - M_j(x_j)|) \quad (4.2)$$

where x_j represents the n original observations and M_i denotes the median of the series. Finally, if the variability of the signal is smaller than twice that of the noise, the neuron is considered silent.

If the neuron is active, spikes must then be extracted for further classification. All peaks in the signal are detected based on three criteria:

- Peaks must exceed a threshold defined as the baseline plus 2.5 times the signal standard deviation. The baseline is computed as the median of the portion of the signal below the 50th percentile, making it more robust to outliers than the mean.
- Peaks must be sufficiently prominent, with a prominence greater than 10 times the noise deviation (estimated using the MAD as described above).
- Peaks must be separated by at least 10 ms, enforcing a refractory period consistent with the chip dynamics and preventing multiple detections of the same spike.

Together, these criteria allow solid detection of spikes despite noise and electronic artifacts.

Finally, spiking and bursting behaviors must be distinguished in order to apply the DIC estimation summarized in Figure 4.6. Using the spike times, all inter-spike intervals are computed, along with their mean and standard deviation. The coefficient of variation, defined as the ratio of the standard deviation to the mean of the intervals, is then used to differentiate tonic spiking (coefficient < 0.4) from bursting (coefficient > 0.4). Depending on the type of activity, specific metrics are estimated to obtain the slow DIC estimation (Figure 4.6).

4.5.2 Control loop implementation

The feedback control loop is similar to that shown in Figure 2.5, with several adaptations for chip implementation. The output of the slow DIC estimation block is discrete, as explained


in the previous section. Moreover, the chip is voltage-biased, and the output corresponding to membrane activity is the voltage V_f . Finally, only a proportional controller is used. Since the tracking variable is discrete, this choice enables faster tuning and reduces stability issues, however a steady-state error will remain. The controller output, $V_{G,0_s}$, is applied to the neurons through the DAC. The updated controller scheme is presented in Figure 4.7.

In addition, multiple neurons are available on a single chip. All neurons share the same biases, even though only one output can be observed at a time. The controller is therefore connected to an interface that allows the user to select the recorded neuron, as well as the desired activity. Each time the neuron of interest is changed, the calibration parameters and inputs are updated according to Table A.5. The set of achievable activities is also adjusted to prevent selecting a reference that cannot be reached by the chosen neuron. When the desired activity is modified, the controller reference is updated accordingly. To allow the neuron behavior to stabilize, the controller is paused for 1 s after a neuron change.

In parallel with these operations, given the current target, the neuron output is recorded and converted into a discrete slow DIC estimate. This estimate is then used to compute the error fed to the proportional controller, which determines the new slow sigmoid gain to be applied to the neuron. The gain is chosen to ensure precise and fast control, with $k_p = 0.005$. The update rule for the modulatory gain is:

$$V_{G,0_s}^t = V_{G,0_s}^{t-1} + k_p \cdot (g_{s_{ref}} - \hat{g}_s) \quad (4.3)$$

The initial value is chosen arbitrarily below saturation, which is set to 0.245 V for all neurons. The controller is again paused for 1 s after a bias change to allow the system to stabilize. The controller runs continuously; however, since the update is proportional to the error, once the target is reached, the error converges to zero and the neuromodulatory gain remains stable.



4.6 Results of neuromodulation in ultra-low power neurons

4.6.1 Control loop impact on firing pattern variability


The goal of this thesis was to achieve robust neuromodulation for CMOS neurons that are sensitive to electronic mismatch. Such mismatch introduces variability both between a neuron and its ideal model, and among neurons on the same chip as well as across different chips. In ultra-low power operation, particularly in the subthreshold regime, the impact of mismatch has been shown to be significant, leading to substantial differences in neuronal dynamics under identical bias conditions. To mitigate part of this variability, neuron-specific calibration sets were employed, enabling a more consistent range of neuromodulation across neurons. The remaining variability manifests in differences in neuromodulation capabilities: some neurons exhibit a wide frequency range in spiking, others are capable of stable bursting with a high number of spikes per burst, while some can reach very low spiking rates. The approach adopted in this thesis is to exploit this heterogeneity. The chip was therefore integrated into a control loop to robustly neuromodulate each neuron, thereby handling intra-neuron mismatch.


As an initial validation of the feedback control loop, the activity of the four neurons was recorded before control, when the slow sigmoid gain was identical for all neurons, and after control, once the reference slow DIC had been reached. These comparisons are shown for two different initial gains and targets in Figures 4.8 and 4.9. It can be observed that, even with the calibration sets, the initial activity differs across neurons for both conditions. However, all neurons are active—some exhibit spiking behavior, others bursting—but none are silent or saturated. This demonstrates the reduction of mismatch effects achieved through calibration.

Once the target is reached, all neurons exhibit similar firing patterns, with a consistent number of spikes per burst and comparable frequencies in the spiking regime (Figures 4.8 and 4.9, right panels). Due to the discrete estimation of the DIC, which results in activity being categorized into ranges, the output activity metrics are not exactly identical. In particular, the intra-burst frequency and the spiking rate still vary across neurons. This limitation is a consequence of the activity classification approach and represents a trade-off for the robustness and stability it provides. Additionally, this frequency dispersion is partly due to residual mismatch effects, which lead to differences in the underlying neuronal dynamics. However, such differences in activity can be related to biology. Indeed, even if functional output is ensured in the absence of pathology, neuronal activity is never identical.

4.6.2 Evolution of parameters during control

To further analyze the effects of the control loop, the evolution of several parameters was recorded during the control of the four neurons. These metrics include the estimated slow DIC, the frequency (either the firing rate or the intra-burst frequency), and the number of spikes per burst, which is equal to one in tonic spiking. The experiment was repeated for three different





sets of initial slow sigmoid gain and target values. The results are shown for two of them in Figures 4.10 and 4.11 and will be analyzed. The last one is presented in Figure B.1 in Appendix B.

For both simulations, the slow DIC and the number of spikes converge to the same value for all neurons, except for neuron 2 in the second simulation (Figure 4.11), where a different target was imposed because the common one was not reachable. This demonstrates that using a control loop with DIC categories as targets ensures convergence of neuronal activity to the desired state. The number of spikes is indirectly imposed by the reference, since tonic spiking corresponds to one spike per burst, while in bursting regimes, the classification is based on the number of spikes.

Regarding frequency, in the first simulation (Figure 4.10), the target activity is bursting with 2 or 3 spikes per burst; therefore, the steady-state frequency corresponds to the intra-burst frequency. The final frequencies of the first three neurons are close to each other, around 27 Hz, while neuron 6 exhibits a lower rate of approximately 13 Hz. This difference can be explained by both mismatch and control design choices, as discussed for Figure 4.2.




FIGURE 4.10 – Parameters evolution during control. Evolution of the estimated slow DIC, the frequency (inter-spike for spiking and intra-burst for bursting), the number of spikes per burst (equal to one for spiking) and the slow sigmoid gain for the four neurons, during control. A recording window of 2s between control update is used, but the time axis is normalized for clarity. The initial neuromodulatory gain is $V_{G,0s} = 0.15$ V and the slow DIC reference is $g_s = 2.5$. All neurons reach the desired activity, with two spikes per burst, while final frequencies and neuromodulatory gains differ.




FIGURE 4.11 – **Parameters evolution during control.** Evolution of the estimated slow DIC, the frequency (inter-spike for spiking and intra-burst for bursting), the number of spikes per burst (equal to one for spiking) and the slow sigmoid gain for the four neurons, during control. A recording window of 2 s between control update is used, but the time axis is normalized for clarity. The initial neuromodulatory gain is $V_{G,0_s} = 0.15$ V and the slow DIC reference is $g_s = 5$, except for neuron 2, where $g_s = 5.5$ (5 being unreachable for this neuron). All neurons reach the desired activity of spiking type II. The final frequencies are close except for neuron 2 as the target is slightly different. The final neuromodulatory gains differs.

In the second simulation (Figure 4.11), neuron 2 fires at a higher rate than the others, as its target is spiking II, with a firing frequency between 35 and 60 Hz. In contrast, the remaining neurons reach frequencies between 25 and 35 Hz, as desired.

While the final activity is constant across neurons, even if the frequency may differ, the steady-state slow sigmoid gain required to reach the target does not converge to a unique value. Each neuron requires a different neuromodulatory gain, determined by the controller, in order to impose its activity. This demonstrates the need for a control loop to achieve neuromodulation that is robust to mismatch-induced variability in neuronal dynamics. The distribution of the final gain confirms results from the model-based simulations in Figure 2.15.

It should also be noted that both the initial parameter values and the evolution of these metrics differ between neurons. This confirms the results shown in Figures 4.2 and 4.3, highlighting the initial variability in activity prior to the effect of the control loop. The different trajectories followed by the neurons result from mismatch and further corroborate the neuron-specific neuromodulation characteristics observed in Figures 4.4 and 4.5. For the same reasons, the time

needed to reach the target activity differs between neurons.

4.6.3 Neurons output evolution

Lastly, the transition of the neuronal output from the initial activity to the target activity can be observed. The results corresponding to the simulations presented in Figures 4.10 and 4.11 are shown in Figures 4.12 and 4.13, respectively. These results confirm the previous observations: despite different initial behaviors, trajectories, and convergence times, all neurons exhibit the same firing pattern at steady state.

4.7 Conclusions on neuromodulation for ultra-low power CMOS neurons

Mismatch in neuromorphic circuits, particularly in subthreshold implementations in the picoampere regime, leads to variability between transistors on a chip. Considering the current-mode mixed-feedback neuron proposed by Mendolia et al. [10], this variability drives neuronal dynam-




FIGURE 4.12 – **Activity evolution during control.** Evolution of the activity of the four neurons during control. A recording window of 2s between control update is used. The initial neuromodulatory gain is $V_{G,0_s} = 0.15$ V and the slow DIC reference is $g_s = 2.5$. All neurons reach the desired activity of bursting with two spikes per burst. The time to reach this activity differ between neurons.




FIGURE 4.13 – **Activity evolution during control.** Evolution of the activity of the four neurons during control. A recording window of 2 s between control update is used. The initial neuromodulatory gain is $V_{G,0_s} = 0.15$ V and the slow DIC reference is $g_s = 5$, except for neuron 2, where $g_s = 5.5$ (5 being unreachable for this neuron). All neurons reach the desired activity of spiking type II. The time to reach this activity differ between neurons.

ics away from the nominal model and introduces differences between neurons, making open-loop neuromodulation unfeasible.

Following the demonstration of mismatch effects, a calibration procedure and a feedback control loop were introduced to enable robust neuromodulation. In this thesis, the chosen approach exploits the remaining variability in neuronal behavior by adapting the set of achievable dynamics to each individual neuron. The proposed controller shows promising results for the four calibrated neurons among the first six, demonstrating that closed-loop robust neuromodulation can compensate for mismatch in CMOS neurons, enabling activity control.

However, the scope of this study is limited by the small number of neurons considered (four out of sixteen), although the results remain encouraging. Another limitation lies in the calibration process, which is specific to each neuron and not directly transferable across chips. This highlights the need for an automated calibration procedure.

Finally, the chip used in this work is not optimized for mismatch reduction. The question of achieving robust neuromodulation on a less mismatch-sensitive chip remains open and warrants further investigation. Therefore, the next section focuses on the analysis and modulation of the

same neurons under higher bias currents, which reduce mismatch effects.

CHAPTER 5

NEUROMODULATION ON LOW-POWER NEURONS

In this chapter, robust neuromodulation is implemented in low-power neurons, corresponding to the nanoampere range, following the same pipeline as in Chapter 5.

The new set of biases was obtained by adding 200 mV to each voltage, corresponding to scaling the current by a factor of 1000 (from picoamperes to nanoamperes), except for those controlling the timescales. Each bias was subsequently fine-tuned to achieve a complete range of activity regimes (type I spiking, bursting, and type II spiking) across the neurons. This current scaling is possible thanks to the current scale invariance of the mixed-feedback model.

The resulting set is reported in Table A.6, and the neuromodulatory gain $V_{G,0s}$ is varied throughout the simulations.

5.1 Mismatch effects on low power neurons

Similarly to the ultra-low power case, mismatch-induced variability is first assessed using a fixed set of biases across neurons and by recording the output activity. Only the slow sigmoid gain is changed between the two sets, from 0.43 V to 0.48 V.

The first set, whose comparison is presented in Figure 5.1, shows that all neurons are firing, with spikes reaching up to 1.7 V, corresponding to 17 nA (Equation 4.1). Only one neuron exhibits bursting behavior, making the activity much more consistent than ultra-low power simulations.

In Figure 5.2, it can again be observed that no neurons are silent: two exhibit bursting behavior, while the others are spiking.

Considering the sixteen neurons on the chip, Set 1 results in 6.25% bursting neurons and 93.75% low-frequency spiking neurons. For Set 2, 25% exhibit type I spiking, 31.25% are burst-

ing, and 43.75% show type II spiking. This transition in regime is expected, since the first set uses a lower value of the slow sigmoid gain, resulting mainly in type I spiking. In contrast, the second set, with an increased gain, switches behavior toward bursting or high-frequency spiking. Compared to the ultra-low-power variability, the neurons still exhibit different firing patterns, but these activities are less dispersed. Whereas, in the ultra-low-power regime, neurons could vary from silent to saturated states through tonic spiking, all neurons remain active in the low-power regime.




FIGURE 5.1 – **Low-power inter-neuron mismatch effect.** Output measurement from six neurons biased with the same set of voltages reported in Table A.6 and a slow sigmoid gain $V_{G,0_s} = 0.43$ V. Among the six neurons, one is bursting (1) and the others are spiking at low frequency.




FIGURE 5.2 – **Low-power inter-neuron mismatch effect.** Output measurement from six neurons biased with the same set of voltages reported in Table A.6 and a slow sigmoid gain $V_{G,0_s} = 0.48$ V. Among the six neurons, two are bursting (2, 3) and the others are spiking at high frequency.

This demonstrates a reduction in mismatch effects due to the increase in current in the circuit. The neurons behaviors are more similar, both in firing patterns and frequencies, and no neurons are silent or saturated while others spike.

5.2 Neuromodulation in low-power neurons

Since variability is reduced, calibration is no longer required, and the single set of biases presented in Table A.6 enables each neuron to transition from type I spiking to bursting, and finally to type II spiking, using only the slow neuromodulatory bias.

These transitions, along with the corresponding slow sigmoid gain, are illustrated in Figure 5.3 for neurons 1 and 2, and in Figures B.3 and B.4 from Appendix B, for the others. A more detailed analysis shows that all neurons exhibit type I spiking at a constant rate, followed by bursting with an increasing number of spikes, and finally high-rate spiking with increasing frequency. However, the spiking frequencies remain distributed across neurons, which complicates the precise modulation of tonic spiking. In addition, the spiking rates are increased compared to ultra-low power simulations, with frequencies up to 25 Hz in type I spiking for some neurons.

5.3 Controller adaptation

Since the neuronal dynamics were modified by the scaling of the current, the control loops also required slight adjustments.

First, because higher firing frequencies are encountered, the refractory period used for spike detection was reduced to 1 ms.

As all neurons share the same range of possible activities, a common classification is used. However, since frequencies in both spiking regimes vary significantly across neurons, tonic spiking is not divided into rate-dependent subcategories. Only the bursting regime is subdivided, according to the number of spikes per burst.

The estimation of the DIC based on neuronal activity is summarized in Table 5.1. Smaller intervals are used within the bursting regime to allow precise control, whereas larger separations between regimes are employed to ensure faster convergence.

	Spiking I		Bursting			Spiking II
			2-3 spb	4-5 spb	6-7 spb	> 7 spb
g_s	0.5	1.5	2.0	2.5	3.0	4.0

TABLE 5.1 – Low-power slow DIC estimation. Classification of the firing pattern into a discrete slow DIC value for activity control in low-power neurons. All categories are reachable by all neurons.


Then, the choice of the desired activity becomes independent of the neuron of interest, since all neurons share the same set of possible activities. The only parameter that changes from one neuron to another is the upper saturation limit of the neuromodulatory gain. Indeed, during the experiments, it was observed that if this gain becomes too high, the neurons saturate at undesirable membrane voltages. Moreover, once a neuron is saturated, a large decrease in the

slow sigmoid gain is required to recover spiking activity. To prevent this issue, the following rule is applied whenever the controller output $V_{G,0_s}$ exceeds the saturation limit:

$$V'_{G,0_s} = V_{sat} - 0.7 \cdot (V_{G,0_s} - V_{sat}) \quad (5.1)$$

where V_{sat} is the upper saturation value.

The final adaptation concerns the gain of the proportional controller. This gain, K_p , is set to 0.002 to ensure sufficiently precise control, allowing each neuron to be directed toward all activity types. Such a small value is mainly required because neuron 1, in the bursting regime, needs only a very small gain increase to switch behavior. For clarity, a larger gain of 0.005 is used in some simulations where precise control is not required.



5.4 Results of neuromodulation in low-power neurons

Identically to ultra-low power simulations, the controller results will be analyzed using the evolution of some metrics and the neuronal activities along the control till target reaching.

Changes in parameters are presented in Figure 5.4 for a target activity of 4-5 spikes per burst. All neurons converge to the desired firing pattern, while the final frequency and $V_{G,0_s}$ differ across neurons. Similarly to the ultra-low-power results, this demonstrates the need for a feedback loop to achieve robust neuromodulation. In contrast to the ultra-low power regime, the initial activity of the neurons is closer across neurons, emphasizing the effect of using higher currents.

The transient behavior of the neurons before convergence can be seen in Figure 5.5. It can be observed that, apart from the initial firing patterns, the trajectories toward the desired activity are also more similar compared to those in the ultra-low-power regime.




FIGURE 5.4 – Parameters evolution during control in low-power. Evolution of the estimated slow DIC, the frequency (inter-spike for spiking and intra-burst for bursting), the number of spikes per burst (equal to one for spiking) and the slow sigmoid gain for the six neurons, during control. A recording window of 1 s between control update is used, but the time axis is normalized for clarity. The proportional gain used is $K_p = 0.005$, the initial neuromodulatory gain is $V_{G,0_s} = 0.41$ V and the slow DIC reference is $g_s = 2$. All neurons reach the desired activity of spiking type bursting with 4 or 5 spikes per burst. The final frequencies and gains are spread.

Additional results are presented in Figures B.5 and B.6 in Appendix B for a type II spiking target, from which the same conclusions can be drawn.




FIGURE 5.5 – Activity evolution during control. Evolution of the activity of the six neurons during control. A recording window of 1 s between control update is used but the time axis is normalized for clarity. The proportional gain used is $K_p = 0.005$, the initial neuromodulatory gain is $V_{G,0_s} = 0.41$ V and the slow DIC reference is $g_s = 2$. All neurons reach the desired activity of bursting with four or five spikes per burst.

Discussion

This thesis addressed the problem of robust neuromodulation in the presence of mismatch-induced variability in neuromorphic chips. In a subthreshold neuron implementation, a control loop was integrated in two current regimes: the picoampere (or ultra-low-power) and the nanoampere (or low-power) ranges. In both regimes, activity regulation was successfully achieved, although with differences in implementation and resulting behavior. These variations are therefore compared and discussed, highlighting the limitations and advantages of each regime. Consequently, the associated trade-offs, possible applications, and future perspectives are also addressed.

Low-power and ultra-low-power neuromodulation. In both regimes, a major contribution of this work is the demonstration of robust neuromodulation in neuromorphic neurons despite variability in neuronal dynamics, which is also demonstrated here. Output activity feedback, using the slow DIC as a reference metric, enables the neurons to converge toward the desired firing pattern. The target can be selected, at least, among the three main firing-pattern categories, i.e., type I tonic spiking, type II tonic spiking and bursting. This minimal categorization ensures functional control of the neurons, even though the exact activity may differ across neurons, similarly to biological systems.

Ultra-low-power neuromodulation: limitations and advantages. In the picoampere range, mismatch induces strong intrinsic variability in neuronal dynamics, ranging from silent to spiking or saturated behaviors, as well as increased activity instability. This heterogeneity leads to the need for individual calibration of each neuron in order to obtain a minimal set of firing patterns achievable by the neuron. This calibration was performed manually and was time-consuming, as the parameter sets are widely spread due to the strong mismatch effects. While this remains feasible for small-scale integration, it becomes impractical for LSI/VLSI systems.

The same limitation applies to the reachable neuronal activities of each neuron, which also had to be extracted manually. However, the strong mismatch effects at very low currents introduce diversity in neuronal behaviors that can be exploited and related to biological systems. Each neuron exhibits its own neuromodulation capabilities, mimicking the heterogeneity observed in biological brains. Nonetheless, the ultra-low-power consumption remains the main advantage of this operating regime.

Low-power neuromodulation: limitations and advantages. Moving to the nanoampere range increases power consumption but provides several benefits for the neuromodulation of neuromorphic neurons. This regime reduces neuron-to-neuron variability in neuronal dynamics, thereby removing the need for individual calibration. A calibration set of biases is still required, but it is shared across all neurons. As a result, neurons exhibit more similar neuromodulation capabilities and share a common range of achievable firing patterns. However, the range of accessible activities is reduced, with, for example, less control over tonic spiking rates. In addition, the common set of biases still needs to be determined to ensure robust neuromodulation. On the other hand, increasing the current leads to more stable neuronal activity over time, thereby improving the stability of the control loop.

Trade-offs and application-specific uses. The ultra-low-power implementation, despite requiring individual calibration and exhibiting possible instabilities, is better suited for small-scale applications where power consumption is highly critical, such as implantable systems constrained by battery capacity. In contrast, the low-power regime should be preferred in applications where stability and reproducibility are prioritized over extremely low power consumption, such as large-scale computing architectures. However, this duality is not strict, and other current ranges could be explored to achieve a balance between power consumption and stability.

Perspectives

This work opens several future directions for the robust neuromodulation of silicon neurons. One possible extension is the integration of the ultra-low-power implementation into very-large-scale integration systems. This would require automated calibration and neuron-specific analyses to exploit large heterogeneous neuronal populations effectively. The results could also be generalized to a larger number of neurons, to chips that are less sensitive to mismatch, or even to multi-chip systems. On the controller side, one perspective would be to extend the controller to include the input current as a control variable, since it also influences neuronal excitability. More generally, this work validates the slow DIC metric as a relevant measure of neuronal excitability within a control loop, which could be extended to a wide range of neuronal models.

Conclusion

Overall, this work demonstrates that robust neuromodulation of silicon neurons can be achieved despite high mismatch-induced variability, encountered in subthreshold and ultra-low current implementations. The proposed control approach regulates excitability structure through slow dynamic input conductances, across multiple operating regimes, highlighting a trade-off between ultra-low-power efficiency and low-power stability. These results contribute toward the development of scalable and biologically inspired neuromorphic systems capable of robust adaptive behavior.

Appendix A Simulation parameters

This appendix groups all the parameters necessary to reproduce the different graphs or results from this thesis. Each table is referenced in the core text in its corresponding part of the process.

Bias	Default value
τ_f	0.0001 s
τ_s	0.02 s
τ_u	0.4 s
I_{thr_f}	150 nA
I_{gain_f}	700 nA
I_{lin_f}	250 nA
I_{thr_s}	110 nA
I_{lin_s}	50 nA
G_{fDPI}	1
G_{sDPI}	1
G_{uDPI}	1
I_{app}	400 nA

TABLE A.1 – Default values of the biases for the Julia model of the current-mode mixed-feedback model used for controller design.

	Spiking type I	Bursting	Spiking type II
$g_{s_{ref}} [-]$	0.7683	2.5216	4.9111
$I_{gain_s} [\text{nA}]$	223	275	340

TABLE A.2 – Parameters used to analyze the robustness of the controller to mismatch. $g_{s_{ref}}$ is the slow DIC reference used in the control loop. I_{gain_s} is the neuromodulatory gain used in the fixed parameter configuration, before applying mismatch.

Bias	Default value
$I_{\tau f}$	100 pA
$I_{\tau s}$	20 pA
$I_{\tau u}$	10 pA
$I_{thr f}$	110 pA
$I_{G,0f}$	600 pA
$I_{lin f}$	250 pA
$I_{thr s}$	90 pA
$I_{lin s}$	50 pA
G_{fDPI}	1
G_{sDPI}	1
G_{uDPI}	1

TABLE A.3 – Default values of the biases for the mixed-feedback current-mode neuron simulated in Cadence.

Bias	Value in set 1 [V]	Value in set 2 [V]
$V_{th f}$	0.13	0.13
$V_{th s}$	0.11	0.11
$V_{th u}$	0.04	0.04
$V_{\tau f}$	0.13	0.13
$V_{\tau s}$	0.11	0.11
$V_{\tau u}$	0.07	0.07
$V_{G,0f}$	0.27	0.333
$V_{G,0s}$	0.16	0.19
$V_{thr f}$	0.13	0.15
$V_{thr s}$	0.11	0.1
$V_{lin f}$	0.29	0.29
$V_{lin s}$	0.15	0.15
I_{app}	1.57	1.56

TABLE A.4 – Sets of bias voltages used for inter-neuron variability demonstration. Each voltage refers to a voltage in Figure 1.14.

Bias	Neuron 1	Neuron 2	Neuron 5	Neuron 6
V_{th_f}	0.13	0.13	0.13	0.13
V_{th_s}	0.11	0.11	0.11	0.11
V_{th_u}	0.04	0.04	0.04	0.04
V_{τ_f}	0.13	0.13	0.13	0.13
V_{τ_s}	0.11	0.11	0.11	0.11
V_{τ_u}	0.07	0.07	0.07	0.07
$V_{G,0_f}$	0.3	0.27	0.32	0.333
V_{thr_f}	0.15	0.13	0.16	0.15
V_{thr_s}	0.1	0.11	0.1	0.1
V_{lin_f}	0.29	0.29	0.34	0.29
V_{lin_s}	0.15	0.15	0.09	0.15
I_{app}	1.57	1.57	1.566	1.56

TABLE A.5 – Calibration sets of bias voltages for the four neurons used, allowing neuromodulation using only the slow sigmoid gain for each neuron.

Bias	Low-power value
V_{th_f}	0.33
V_{th_s}	0.1279
V_{th_u}	0.08
V_{τ_f}	0.16
V_{τ_s}	0.11
V_{τ_u}	0.07
$V_{G,0_f}$	0.533
V_{thr_f}	0.35
V_{thr_s}	0.3
V_{lin_f}	0.3
V_{lin_s}	0.4
I_{app}	1.28

TABLE A.6 – Sets of bias voltages used for lower-power simulations.

Appendix B Additional results

This appendix presents additional results on experiments from Chapters 4 and 5.




FIGURE B.1 – Parameters evolution during control. Evolution of the estimated slow DIC, the frequency (inter-spike for spiking and intra-burst for bursting), the number of spikes per burst (equal to one for spiking) and the slow sigmoid gain for the four neurons, during control. A recording window of 2s between control update is used, but the time axis is normalized for clarity. The initial neuromodulatory gain is $V_{G,0_s} = 0.17$ V and the slow DIC reference is $g_s = 1.5$, except for neuron 5, where $g_s = 0.5$ (1.5 being unreachable for this neuron). All neurons reach the desired activity of spiking type I. The final frequencies are close except for neuron 5 as the target is slightly different. The final neuromodulatory gains differs.






FIGURE B.2 – Activity evolution during control. Evolution of the activity of the four neurons during control. A recording window of 2s between control update is used, but the time axis is normalized. The initial neuromodulatory gain is $V_{G,0_s} = 0.17\text{V}$ and the slow DIC reference is $g_s = 1.5$, except for neuron 5, where $g_s = 0.5$ (1.5 being unreachable for this neuron).



Neuron 5



Neuron 6




FIGURE B.4 – **Low-power neuromodulation.** Transition from low-frequency spiking to high-frequency spiking through bursting activity for neuron 5 and 6. The range of the neuromodulatory input, labeled V_{gain_s} , is adapted according to the neuron.




FIGURE B.5 – Parameters evolution during control in low-power neurons. Evolution of the estimated slow DIC, the frequency (inter-spike for spiking and intra-burst for bursting), the number of spikes per burst (equal to one for spiking) and the slow sigmoid gain for the six neurons, during control. A recording window of 1 s between control update is used, but the time axis is normalized for clarity. The proportional gain used is $K_p = 0.005$, the initial neuromodulatory gain is $V_{G,0_s} = 0.41$ V and the slow DIC reference is $g_s = 4$. All neurons reach the desired activity of spiking type bursting with 4 or 5 spikes per burst. The final frequencies and gains are spread.




FIGURE B.6 – Activity evolution during control in low-power neurons. Evolution of the activity of the six neurons during control. A recording window of 1 s between control update is used but the time axis is normalized for clarity. The proportional gain used is $K_p = 0.005$, the initial neuromodulatory gain is $V_{G,0_s} = 0.41$ V and the slow DIC reference is $g_s = 4$. All neurons reach the desired activity of high frequency spiking.

Appendix C Additional materials

C.1 Control interfaces

This appendix shows the control interface used in Chapter 4 and 5, enhancing the difference between ultra-low power regime, where the choice of desired activity depends on the selected neurons while it does not in the low-power regime.




FIGURE C.1 – **Ultra-low-power control interface example.** Image of the interface used to modify the target activity in the picoampere regime, for the neuron 2. The possible activities are modified according to the desired neuron.




FIGURE C.2 – **Ultra-low-power control interface example.** Image of the interface used to modify the target activity in the picoampere regime, for the neuron 6. The possible activities are modified according to the desired neuron.

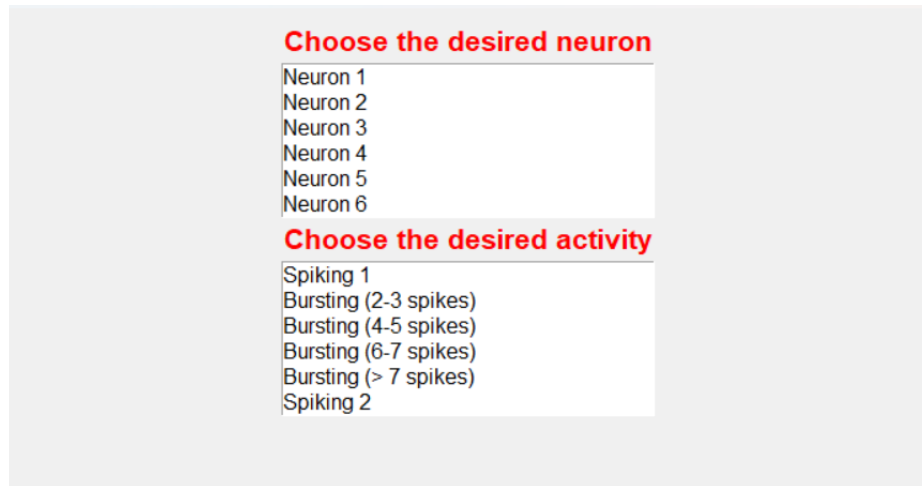


FIGURE C.3 – **Low-power control interface example.** Image of the interface used to modify the target activity in the nanoampere regime. The possible activities are common to all neurons.

C.2 Codes

All codes used to obtain the results presented in this thesis can be found in the following repository: <https://gitlab.uliege.be/Victoria.Filee/pcb-chip-1-interface>. The source repository containing the codes for the Julia model simulations can be found here: <https://github.com/LorisMendolia/IMF-Neuron-Circuit-Sim>. The codes for the chip simulations can be found in the following repository: <https://gitlab.uliege.be/LMendolia/pcb-chip-1-interface>.

BIBLIOGRAPHY

- [1] Cornelia I Bargmann. Beyond the connectome: how neuromodulators shape neural circuits. *Bioessays*, 34(6):458–465, 2012.
- [2] Farzan Nadim and Dirk Bucher. Neuromodulation of neurons and synapses. *Current opinion in neurobiology*, 29:48–56, 2014.
- [3] Eve Marder and Dirk Bucher. Central pattern generators and the control of rhythmic movements. *Current biology*, 11(23):R986–R996, 2001.
- [4] Eve Marder, Timothy O’Leary, and Sonal Shruti. Neuromodulation of circuits with variable parameters: single neurons and small circuits reveal principles of state-dependent and robust neuromodulation. *Annual review of neuroscience*, 37(1):329–346, 2014.
- [5] Hao Wang, Bin Sun, Shuzhi Sam Ge, Jie Su, and Ming Liang Jin. On non-von neumann flexible neuromorphic vision sensors. *npj Flexible Electronics*, 8(1):28, 2024.
- [6] Giacomo Indiveri and Shih-Chii Liu. Memory and information processing in neuromorphic systems. *Proceedings of the IEEE*, 103(8):1379–1397, 2015.
- [7] Daniele Ielmini and Stefano Ambrogio. Emerging neuromorphic devices. *Nanotechnology*, 31(9):092001, 2020.
- [8] Elisabetta Chicca, Fabio Stefanini, Chiara Bartolozzi, and Giacomo Indiveri. Neuromorphic electronic circuits for building autonomous cognitive systems. *Proceedings of the IEEE*, 102(9):1367–1388, 2014.
- [9] Chetan Singh Thakur, Jamal Lottier Molin, Gert Cauwenberghs, Giacomo Indiveri, Kundan Kumar, Ning Qiao, Johannes Schemmel, Runchun Wang, Elisabetta Chicca, Jennifer Olson Hasler, et al. Large-scale neuromorphic spiking array processors: A quest to mimic the brain. *Frontiers in neuroscience*, 12:891, 2018.
- [10] Loris Mendolia, Chenxi Wen, Elisabetta Chicca, Giacomo Indiveri, Rodolphe Sepulchre, Jean-Michel Redouté, and Alessio Franci. A neuromodulable current-mode silicon neuron for robust and adaptive neuromorphic systems. *arXiv preprint arXiv:2512.01133*, 2025.

- [11] Fernando Castanos and Alessio Franci. Implementing robust neuromodulation in neuromorphic circuits. *Neurocomputing*, 233:3–13, 2017.
- [12] Marcel JM Pelgrom, Aad CJ Duinmaijer, and Anton PG Welbers. Matching properties of mos transistors. *IEEE Journal of solid-state circuits*, 24(5):1433–1439, 2003.
- [13] Greg Stuart, Nelson Spruston, Bert Sakmann, and Michael Häusser. Action potential initiation and backpropagation in neurons of the mammalian cns. *Trends in neurosciences*, 20(3):125–131, 1997.
- [14] Mark W Barnett and Philip M Larkman. The action potential. *Practical neurology*, 7(3):192–197, 2007.
- [15] Barry W Connors and Michael J Gutnick. Intrinsic firing patterns of diverse neocortical neurons. *Trends in neurosciences*, 13(3):99–104, 1990.
- [16] Nataša Savić, Paola Pedarzani, and Marina Sciancalepore. Medium afterhyperpolarization and firing pattern modulation in interneurons of stratum radiatum in the ca3 hippocampal region. *Journal of neurophysiology*, 85(5):1986–1997, 2001.
- [17] Arthur Fyon, Alessio Franci, Pierre Sacré, and Guillaume Drion. Dimensionality reduction of neuronal degeneracy reveals two interfering physiological mechanisms. *PNAS nexus*, 3(10):pgae415, 2024.
- [18] Arthur Fyon and Guillaume Drion. Activity-dependent neuromodulation and calcium homeostasis cooperate to produce robust and modifiable neuronal function. *PLOS Computational Biology*, 22(4):e1014177, 2026.
- [19] Guillaume Drion, Timothy O’Leary, and Eve Marder. Ion channel degeneracy enables robust and tunable neuronal firing rates. *Proceedings of the National Academy of Sciences*, 112(38):E5361–E5370, 2015.
- [20] Eugene M Izhikevich. Which model to use for cortical spiking neurons? *IEEE transactions on neural networks*, 15(5):1063–1070, 2004.
- [21] Eve Marder. Neuromodulation of neuronal circuits: back to the future. *Neuron*, 76(1):1–11, 2012.
- [22] Hyunho Seok, Dongho Lee, Sihoon Son, Hyunbin Choi, Gunhyoung Kim, and Taesung Kim. Beyond von neumann architecture: Brain-inspired artificial neuromorphic devices and integrated computing. *Advanced Electronic Materials*, 10(8):2300839, 2024.
- [23] Dylan Richard Muir and Sadique Sheik. The road to commercial success for neuromorphic technologies. *Nature communications*, 16(1):3586, 2025.
- [24] Guillaume Drion, Alessio Franci, Julie Dethier, and Rodolphe Sepulchre. Dynamic input conductances shape neuronal spiking. *eneuro*, 2(1), 2015.
- [25] Arthur Fyon, Pierre Sacré, Alessio Franci, and Guillaume Drion. Reliable neuromodulation from adaptive control of ion channel expression. *IFAC-PapersOnLine*, 56(2):458–463, 2023.

- [26] Julijana Gjorgjieva, Guillaume Drion, and Eve Marder. Computational implications of biophysical diversity and multiple timescales in neurons and synapses for circuit performance. *Current opinion in neurobiology*, 37:44–52, 2016.
- [27] Namrata Mendiratta and Suman Lata Tripathi. A review on performance comparison of advanced mosfet structures below 45 nm technology node. *Journal of Semiconductors*, 41(6):061401, 2020.
- [28] ElProCus. Cmos technology : Working principle, characteristics & its applications. <https://www.elprocus.com/cmos-working-principle-and-applications/>. Accessed: 2026-02-25.
- [29] Tim Grotjohn and Bernd Hoefflinger. A parametric short-channel mos transistor model for subthreshold and strong inversion current. *IEEE journal of solid-state circuits*, 19(1):100–112, 2003.
- [30] Shih-Chii Liu. *Analog VLSI: circuits and principles*. MIT press, 2002.
- [31] Eric Vittoz and Jean Fellrath. Cmos analog integrated circuits based on weak inversion operations. *IEEE journal of solid-state circuits*, 12(3):224–231, 2003.
- [32] Gregor Cijan, Tadej Tuma, and A Burmen. Modeling and simulation of mos transistor mismatch. In *Proceeding of the 6th Eurosim Congress on Modeling and Simulation (EUROSIM 07)*, pages 1–8, 2007.
- [33] Peter R Kinget. Device mismatch and tradeoffs in the design of analog circuits. *IEEE Journal of Solid-State Circuits*, 40(6):1212–1224, 2005.
- [34] Aleksandra Pavasović, Andreas G. Andreou, and Charles R. Westgate. Characterization of subthreshold mos mismatch in transistors for vlsi systems. *Journal of VLSI signal processing systems for signal, image and video technology*, 8(1):75–85, Feb 1994.
- [35] Ben Varkey Benjamin, Richelle L Smith, and Kwabena A Boahen. An analytical mos device model with mismatch and temperature variation for subthreshold circuits. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 70(6):1826–1830, 2023.
- [36] Mary Ann C Maher, Stephen P Deweerth, Misha A Mahowald, and Carver A Mead. Implementing neural architectures using analog vlsi circuits. *IEEE transactions on circuits and systems*, 36(5):643–652, 1989.
- [37] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham China, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, et al. Loihi: A neuromorphic manycore processor with on-chip learning. *Ieee micro*, 38(1):82–99, 2018.
- [38] Ole Richter, Chenxi Wu, Adrian M Whatley, German Köstinger, Carsten Nielsen, Ning Qiao, and Giacomo Indiveri. Dynap-se2: a scalable multi-core dynamic neuromorphic asynchronous spiking neural network processor. *Neuromorphic computing and engineering*, 4(1):014003, 2024.

- [39] Kai M Hynna and Kwabena Boahen. Silicon neurons that burst when primed. In *2007 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 3363–3366. IEEE, 2007.
- [40] Jayawan HB Wijekoon and Piotr Dudek. Compact silicon neuron circuit with spiking and bursting behaviour. *Neural Networks*, 21(2-3):524–534, 2008.
- [41] Luka Ribar and Rodolphe Sepulchre. Neuromodulation of neuromorphic circuits. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 66(8):3028–3040, 2019.
- [42] Alessio Franci, Guillaume Drion, and Rodolphe Sepulchre. Robust and tunable bursting requires slow positive feedback. *Journal of neurophysiology*, 119(3):1222–1234, 2018.
- [43] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1):65–98, 2017.
- [44] Alessio Franci, Guillaume Drion, Vincent Seutin, and Rodolphe Sepulchre. A balance equation determines a switch in neuronal excitability. *PLoS computational biology*, 9(5):e1003040, 2013.
- [45] Rohit G Kanojiya and PM Meshram. Optimal tuning of pi controller for speed control of dc motor drive using particle swarm optimization. In *2012 international conference on advances in power conversion and energy technologies (APCET)*, pages 1–6. IEEE, 2012.
- [46] Xin-lan Li, Jong-Gyu Park, and Hwi-Beom Shin. Comparison and evaluation of anti-windup pi controllers. *Journal of power Electronics*, 11(1):45–50, 2011.
- [47] Johannes Lengler, Florian Jug, and Angelika Steger. Reliable neuronal systems: the importance of heterogeneity. *PloS one*, 8(12):e80694, 2013.
- [48] Nicolas Perez-Nieves, Vincent CH Leung, Pier Luigi Dragotti, and Dan FM Goodman. Neural heterogeneity promotes robust learning. *Nature communications*, 12(1):5791, 2021.
- [49] Christophe Leys, Christophe Ley, Olivier Klein, Philippe Bernard, and Laurent Licata. Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of experimental social psychology*, 49(4):764–766, 2013.