

Achieving ultra-slow timescales in neuromorphic circuits - Application to neural bursting

Auteur : Graindorge, Pierre

Promoteur(s) : Franci, Alessio

Faculté : Faculté des Sciences appliquées

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

Année académique : 2023-2024

URI/URL : <http://hdl.handle.net/2268.2/20956>

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

Master's thesis

*completed in order to obtain the degree of Master of Science in
Electrical Engineering: Neuromorphic Engineering*

Achieving ultra-slow timescales in neuromorphic circuits - Application to neural bursting

Author: Pierre Graindorge

Supervisor: Pr. Alessio Franci

Academic year 2023-2024

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor, Professor Alessio Franci, for proposing this very interesting thesis subject. His guidance and insights were invaluable to the success of this work.

I am also extremely thankful to Professor Jean-Michel Redouté for his exceptional availability and the crucial help he provided during the development of this thesis. His encouragement and expertise played a significant role in shaping this project.

A special thank you goes to Loris Mendolia, to whom I owe immense gratitude. His generous assistance, willingness to answer my countless questions, and the invaluable tips he provided were indispensable. His high availability and meticulous review of this document significantly improved the quality of this work.

I would like to extend my sincere thanks to Marie Garcia-Bardon, Bowen Wang, and Fernando Garcia-Redondo for the warm welcome they gave me during my internship at imec. Their guidance and support made it the best possible internship I could have wished for. I am particularly grateful for the opportunity they provided me to write my first scientific paper and to present it at a conference, an experience that has been invaluable to my growth as a researcher.

I would also like to thank my friends, who shared this journey with me. The joy they brought into my life and the countless memories we created together during these years will always be cherished.

Finally, I want to express my deepest appreciation to my family—my parents, my sister and her husband, and my grandparents. Their infinite support, wisdom, and encouragement helped me persevere through the challenging times. They have continuously inspired me to become a better version of myself, and for that, I am forever grateful.

Abstract

Neuromorphic engineering seeks to replicate the brain's computational power and energy efficiency in hardware. Current neuromorphic designs, however, face challenges in achieving ultra-slow timescales critical for replicating biological neural behaviors such as realistic bursting patterns. This thesis focuses on addressing these limitations through the design and simulation of neuromorphic circuits capable of ultra-slow dynamics while optimizing area efficiency. Using the Cadence Virtuoso software and a general purpose development kit (GSDK), the work reproduces a reference circuit which mimics biological homeostasis, and incorporates this system to an existing neuron circuit, leading to a new modifiable neuron design. Key advancements include the combined use of a differential pair integrator (DPI) and an automatic gain control (AGC) loop to achieve ultra-slow temporal filtering and new neuromodulation capabilities while avoiding the need for excessively large capacitors. Simulation results demonstrate significant improvements in achieving the desired dynamics with enhanced area efficiency. This work represents a step towards more practical large-scale neuromorphic hardware capable of mimicking complex neural behaviors.

Table of Contents

1	Introduction	1
1.1	Background Information	1
1.2	Problem Statement	1
1.3	Motivation	2
1.4	Methodology	3
2	A Modulable Neuromorphic Neuron	4
2.1	The basic building blocks	4
2.1.1	Differential Pair Integrator	4
2.1.2	Current-mode sigmoid	8
2.2	Feedback structure	10
2.3	Neuromorphic modulation	11
2.4	Limitations of the current design	12
3	The Automatic Gain Control Loop	14
3.1	Homeostasis	14
3.2	System Overview and Functionality	15
3.3	The Low-Leakage Cell	16
3.3.1	The Low-Leakage PMOS	17
3.3.2	Operational Transconductance Amplifiers	19
3.4	Comparator circuit	21
3.4.1	Ideal comparator	22
3.4.2	1-PMOS 1-NMOS inverter	23
3.4.3	Improved comparator	25
3.4.4	Area efficiency comparison	28
3.4.5	The low-leakage cell with ideal switch	29
3.5	Characterization of the AGC loop	31
3.5.1	The AGC Loop with an ideal comparator	32
3.5.2	The AGC Loop with the simple inverter comparator	34

3.5.3	The AGC Loop with the improved comparator	36
4	A New Modulable Neuron with Slower Dynamics	38
4.1	High-level description	39
4.2	Behavior and comparison with the previous design	39
4.3	AGC enabled neuromodulation	43
5	Conclusions and future works	50
5.1	Conclusions	50
5.2	Future works	51
	Bibliography	52

1 Introduction

1.1 Background Information

Neuromorphic engineering is a field at the intersection of neuroscience and electronics that aims to reproduce the efficiency and computational power of biological systems, especially the brain, in silicon. Conventional digital circuits, while powerful, cannot replicate the parallel and distributed processing of neurons because of their fundamentally different architecture. Neuromorphic circuits attempt to bridge this gap by emulating neural structures and functions in hardware, resulting in much more energy-efficient and faster processing systems. These circuits are particularly valuable for tasks that require real-time processing and adaptation, such as the interpretation of sensory data, where biological systems excel [1].

1.2 Problem Statement

Despite advances in neuromorphic hardware, current designs of neuromorphic neurons face limitations in achieving longer timescales, which is critical to accurately replicate biological neuronal behavior. This limitation is particularly evident in processes that require slow dynamics, such as neural bursting [2]. In addition, these designs suffer from area efficiency and integration density issues, which are critical to capture the complexity of large-scale neural networks. Extending the timescales of these circuits while increasing area efficiency and integration density is critical to advances in neuromorphic computing area [3].

1.3 Motivation

State-of-the-art neuromorphic circuits utilize complementary metal-oxide semiconductor (CMOS) technologies built with very-large-scale integration (VLSI) processes, where transistors can operate in two regimes: subthreshold and above-threshold. While conventional circuits typically operate transistors in the above-threshold regime, neuromorphic circuits exploit the subthreshold regime due to its lower power consumption, enabled by the reduced current flowing through the transistors.




Figure 1.1: I-V characteristic curve of an NMOS.

Figure 1.1 shows the typical relationship between the gate-to-source voltage and the drain current of an NMOS transistor, highlighting the two modes of operation. It should be noticed the difference in the drain current increase for similar gate-to-source voltage increase demonstrating the higher gain of transistors in subthreshold. In this regime, the relationship between the drain current and the voltages applied to the NMOS transistor is given by:

$$I_{DS} = I_0 e^{\frac{\kappa V_G}{U_T}} (e^{-\frac{V_S}{U_T}} - e^{-\frac{V_D}{U_T}}), \quad (1.1)$$

where I_0 is the leak current parameter, κ the subthreshold slope factor, U_T the thermal voltage ($\approx 26 \text{ mV}$ at $T = 300\text{K}$), V_G the gate voltage, V_D the drain voltage and V_S the source voltage.

The subthreshold regime’s logarithmic-exponential properties are particularly advantageous for replicating neuronal behaviors [4]. In this context, current-mode design is emphasized, where circuit state variables and signals are represented as currents [5]. The use of subthreshold CMOS devices ensures that the circuit operates effectively over a wide range of currents due to the low voltage swing, which keeps circuit voltages away from the supply rails. Additionally, current-mode design simplifies the integration of different circuit elements by leveraging Kirchhoff’s current law, allowing for seamless combination of circuits. Together, these features make current-mode design with subthreshold transistors a powerful approach for neuromorphic circuit design.

By leveraging these techniques, it may be possible to design neuromorphic circuits that not only match the temporal characteristics of biological neurons but also allow for higher integration densities, making neuromorphic hardware more practical for large-scale applications.

1.4 Methodology

This work focuses on the design and simulation of neuromorphic circuits using the Cadence Virtuoso platform, a widely-used tool in the semiconductor industry for analog and mixed-signal design. The circuits are implemented using the General-Purpose Development Kit 180 (GPDK180) technology, which allows the use models of transistor whose length can be scaled down to 180 nm , and the analogLib library which provides the necessary components for the simulation of subthreshold and current-mode circuits. The thesis is divided into two main parts: first, the reproduction of a reference circuit using the GPDK180 on Cadence, and second, an exploration of the filtering properties of these circuits, combined with the design of a new modifiable neuron circuit. The ultimate goal is to achieve a neuromorphic neuron circuit capable of ultra-slow dynamics suitable for replicating complex neural behaviors such as bursting while seeking for a high area efficiency.

2 A Modulable Neuromorphic Neuron

This chapter introduces the neuron circuit developed by L. Mendolia in the context of its PhD at the University of Liège [6]. This circuit has been the starting point of the work carried out during this thesis and it is crucial to understand its behavior as well as its limitations.

2.1 The basic building blocks

2.1.1 Differential Pair Integrator




Figure 2.1: DPI schematic.

The Differential Pair Integrator (DPI) is a fundamental building block for creating neuromorphic circuits, often used to mimic synaptic communication [7]. The DPI circuit acts as a log-domain linear temporal filter, meaning that: first, the input

current undergoes a logarithmic transformation achieved through the exponential I-V characteristics of CMOS devices operated in the subthreshold regime, as shown in section 1.3. This logarithmic conversion leverages a wide dynamic range of operation while ensuring low power consumption. Then, the logarithmic voltage is linearly integrated, effectively realizing a temporal filtering of the input. Finally, this filtered signal is converted back to a current, representing the filtered synaptic current.

Figure 2.1 shows a CMOS implementation of the DPI, utilizing 3 NMOS transistors, 3 PMOS transistors and a capacitor. It should be noted that the other part of the current-mirror that sets V_τ and thus I_τ is not shown in the figure for the sake of clarity but is built the same way as the one setting V_g . This approach ensures to avoid undesirable behaviors when the capacitor is fully discharged.

In the initial configuration, where the current source is directly connected to the DPI branch, if V_c approaches $0V$, the ideal current source will cause the capacitor to continue discharging and then start charging in the opposite direction, which is undesirable when using a polarized capacitor. The current mirror configuration mitigates this issue by shutting down the NMOS transistor replacing the current source when V_c approaches $0V$, as its V_{DS} will also approach $0V$.

Under the right assumptions, i.e., for similar sub-threshold slope factors between NMOS and PMOS: $\kappa_n = \kappa_p = \kappa$, and for $I_{in} \gg I_\tau$, the operation of the DPI can be modeled by the following equation [7]:

$$\tau \frac{d}{dt} I_{out} = -I_{out} + G I_{in}, \quad (2.1)$$

with time constant $\tau = \frac{CU_T}{\kappa I_\tau}$ – where U_T represents the thermal voltage – and

gain $G = \frac{I_g}{I_\tau}$. Hence, it can be seen that the filtering properties of the DPI can be independently and linearly tuned by the circuit parameters. From there, the transfer function of the circuit is derived:

$$\tau s Y = -Y + G X \hookrightarrow H(s) = \frac{Y}{X} = \frac{G}{1 + \tau s}, \quad (2.2)$$

which has the canonical form of a first-order RC low-pass filter with tunable gain G .

The results of the AC and transient simulations are presented hereunder, showing how the different parameter values affect the filtering characteristics of the circuit.




Figure 2.2: DPI AC simulation with fixed $C=10\text{pF}$.

Figure 2.2 illustrates the magnitude and phase response to small input variations (1 nA around $I_{DC} = 100\text{ nA}$) over a wide range of frequencies for different values of gain and time constants. The traces of the magnitude – expressed in dB – are represented by full lines, while those of the phase are represented by dashed lines.

Although the Bode plot of the DPI is visually analogous to that of a first-order low-pass filter, the results of these simulations demonstrate significant discrepancies with the theoretical expressions developed above. In particular, it can be observed that the two parameters of interest, namely G and τ , are not independent because the value of I_τ has a discernible impact on the gain, even when I_g is set to GI_τ .

Furthermore, the differences between the DPI model and the simulations can be observed numerically. The DC gain – i.e. the magnitude response at $f = 0\text{ Hz}$ – of the circuit with parameter values corresponding to the pink, gold, blue and green curves are 6, 2.73, -1.61 and -3.47 dB , respectively. Similarly, the cutoff frequency which is given theoretically by: $f_c = \frac{1}{2\pi\tau}$, is observed to be $\approx 6\text{ Hz}$ for $\tau = 0.087$ and $\approx 30\text{ Hz}$ for $\tau = 0.0087$ in the simulations instead of 1.83 Hz and 18.3 Hz ,

respectively.

These inconsistencies can be attributed to the fact that Equation 2.1 is based on two fundamental assumptions that are not fully satisfied within the simulation context.




Figure 2.3: DPI transient simulation

Figure 2.3 presents the DPI response to input current pulses, demonstrating the temporal filtering of these pulses. It can clearly be seen how the gain G and the time constant τ affect the dynamics of the circuit. The same observations concerning the fidelity of the simulation to the model made in Figure 2.2 are also applicable here.

2.1.2 Current-mode sigmoid

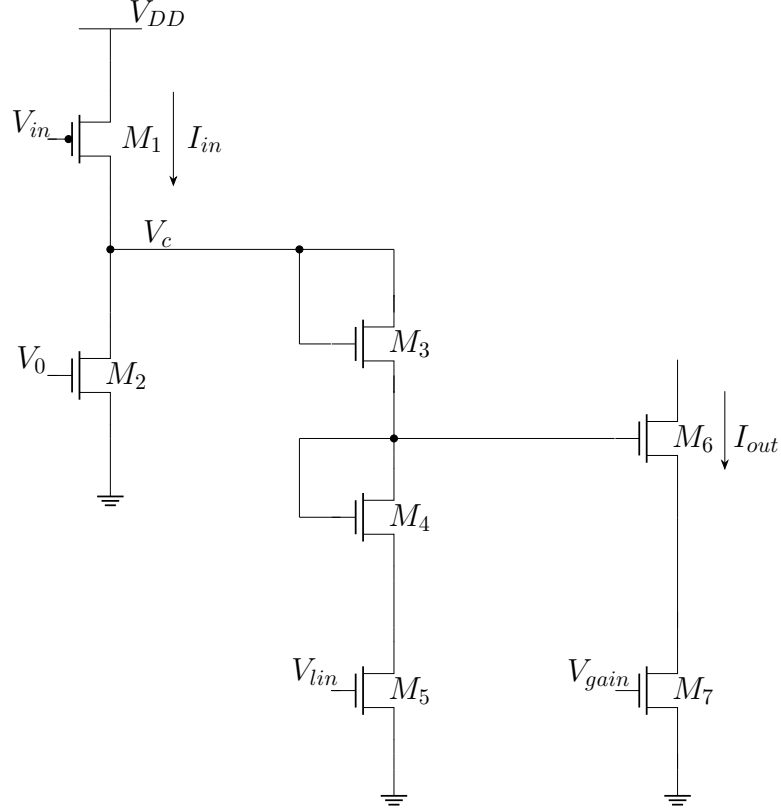


Figure 2.4: Schematic of the Sigmoid circuit designed by Chenxi Wen from the NCS group at the Institute of Neuroinformatics, University of Zurich and ETH Zurich.

This other current-mode circuit exhibits a sigmoidal input-output relationship. The behavior of the circuit is characterized by three key parameters: I_0 , I_{lin} and I_{gain} which control the transition between the low, linear, and high output states. It should be noted that here, I_0 is an independent parameter, different from the leakage current introduced in Equation 1.1. Figure 2.4 illustrates the implementation of the circuit, where V_{in} , V_0 , V_{lin} and V_{gain} are images of the corresponding currents. These voltages are set through current mirrors, which are not shown on this schematic for the sake of clarity.

In order to understand the origin of this sigmoid relationship, the circuit can be decomposed into three distinct branches. The initial branch, comprises the PMOS M_1 with its source connected to V_{dd} and the NMOS M_2 with its source connected to the ground. The drains of the two devices are then connected together. This configuration leads to a comparator. Indeed, a comparison is performed between

the gate voltages of the PMOS and the NMOS, which are proportional to their associated currents I_{in} and I_0 . If $I_{in} < I_0$, the output voltage V_c will be at a low state, close to ground. On the contrary, If $I_{in} > I_0$, the output voltage V_c will be at a high state, close to V_{DD} . However, this comparison is almost digital. Hence, a second branch is added, comprising M_3 , M_4 and M_5 . The purpose of this branch is to draw a controllable current I_{lin} from the first branch, effectively slowing down the comparison. This means that once the input crosses the threshold I_0 , V_c will not be directly pulled up to V_{DD} . Instead, V_c will linearly increase (in the log-domain) with I_{in} before saturating to V_{DD} when I_{in} becomes larger than $I_0 + I_{lin}$. Finally, a third branch is added, where transistor M_6 converts the voltage back to a current, which is capped to a maximum value determined by the current I_{gain} .




Figure 2.5: Sigmoid DC analysis

The results of the DC sweep of the input current are presented in Figure 2.5, which illustrates the influence of the key parameters on the output current.

2.2 Feedback structure

Spiking behavior in neural systems is a complex process that relies on the interplay of both positive and negative feedback mechanisms. Positive feedback plays a crucial role in rapidly driving the system to a high state once the input exceeds a certain threshold, effectively triggering the spike. However, if only positive feedback were present, the system would remain locked in this high state after crossing the threshold, preventing the generation of discrete spikes. To ensure the system returns to its resting equilibrium almost immediately after reaching the high state, negative feedback is introduced. This negative feedback allows the system to quickly reset, thereby enabling the creation of spikes – brief, transient events essential for neural signaling.

In real neural systems, these feedback mechanisms are even more complex. Hodgkin and Huxley have been pioneers in laying down the foundations of modern neuronal models. They proposed a widely-recognized model for describing the neurons electrical properties, taking into account multiple ionic currents and showing that these currents operate on different timescales due to the time-dependent evolution of their respective channel conductances [8]. More recently, research has highlighted the importance of both positive and negative feedback in various neuronal models to produce biologically plausible activity [9, 10, 11].

In the context of creating artificial neurons on chips, the neuron circuit presented in this section implements similar feedback loops – although simplified – acting on three different timescales by interconnecting multiple DPI and sigmoid circuits with varying parameter values. A block diagram providing a high-level description of the functioning of the circuit is shown in Figure 2.6.

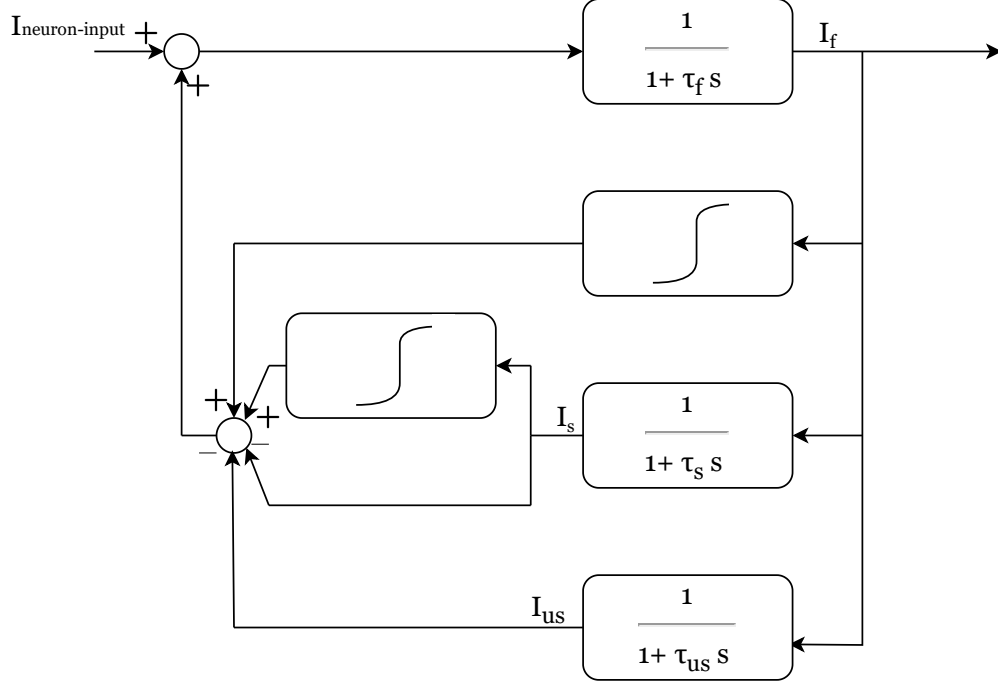


Figure 2.6: High-level description of the neuron circuit [6].

The neuron circuit designed in Cadence Virtuoso implements this block diagram, using three DPI circuits as the first-order low-pass filters and two current-mode sigmoid circuits to control the positive feedback. The DPIs are individually configured, with each one filtering its input on a distinct timescale: fast, slow, or ultra-slow. These timescales correspond to three key tunable parameters of the system: τ_f , τ_s and τ_{us} . Together with the parameters of the sigmoid circuits, they provide a versatile framework for the neuron circuit. The following section shows how tuning their values influences the dynamics of the neuron’s response.

2.3 Neuromorphic modulation

Translating from neuronal models to neuromorphic circuits, it has been demonstrated that it is possible to use similar methods with a mixed-feedback neuron circuit in order to modulate its dynamics [12].

In particular, the modulable neuromorphic neuron designed with the topology presented in the previous section can exhibit a wide range of complex behaviors, such as tonic spiking, tonic bursting, and other dynamic firing patterns observed in bi-

ological neurons. This is achieved by appropriately tuning the main parameters of the system: the fast, slow and ultra-slow times constants of the DPIs and the sigmoid thresholds, linear region and gain. As a result, the neuron circuit is highly adaptable for different computational tasks.




Figure 2.7: Neuron output dependency on input intensity and slow negative feedback gain. Figure taken from the work of L. Mendolia [13].

As an example, Figure 2.7 illustrates how changing the gain of the slow negative feedback, hence the gain of the sigmoid connected to the DPI with time constant τ_s , allows the neuron to switch from type-I excitability (green) to tonic bursting (orange) and type-II excitability (blue) for a given input intensity. Additionally, it is possible to modify the characteristics of a certain behavior, e.g., the spikes amplitude, the spiking frequency in tonic spiking, the intra/inter- burst spiking frequency.

2.4 Limitations of the current design

Despite showcasing high adaptability and already being able to perform several tasks [13], this neuromorphic neuron architecture could be further improved. The main drawbacks of the proposed design lies in the size of the circuit and in the maximum timescales achievable.

As in many areas of electronics, minimizing circuit size is crucial, particularly for enabling very large-scale integration (VLSI) on chips. Smaller circuits not only reduce costs but also allow these chips to be more easily embedded into systems,

making them more portable and less invasive which can be especially important in applications such as biomedical devices.

A key objective of this master thesis is to achieve longer timescales in the neuron circuit to enhance its fidelity to biological neuronal signaling. Current transient simulations in a bursting configuration result in bursting frequencies which are higher than what is typically observed in biological systems for given inputs [2]. For the ultra-slow current, realistic timescales should range from hundreds of milliseconds to minutes. Achieving these extended timescales is crucial for accurately replicating the dynamics of real neurons.

However, extending the timescales poses a challenge related to the circuit's physical layout. In the initial design by L. Mendolia, which occupies a total area of $180 \times 20 \mu m$, more than half of the area is taken up by the capacitors used in the DPI circuits. As discussed in subsection 2.1.1, the time constant τ is proportional to the capacitance C and inversely proportional to the current I_τ . Since I_τ cannot be scaled down beyond a certain point (on the order of picoamperes) due to transistor leakage currents, increasing C is the only way to achieve longer timescales, which inevitably leads to larger capacitors and a larger overall circuit area.

Thus, this thesis aims to explore efficient methods for achieving the necessary longer timescales while simultaneously striving to reduce the total area occupied by the circuit.

3 The Automatic Gain Control Loop

This chapter focuses on reproducing results from [14] using the GPDK180. In this work, the authors developed a circuit that reproduces a mechanism observed in real neural systems, called homeostasis, in order to achieve longer biological timescales. As new results, we also characterize the filtering properties of this system and compare it with the ultra-slow DPI of the neuromorphic neuron circuit detailed in the previous chapter.

3.1 Homeostasis

In 1929, Walter B. Cannon was the first to define homeostasis, building upon the concept of 'milieu intérieur' introduced by the French physiologist Claude Bernard in 1878. He described it as the physiological mechanisms through which living systems are able to maintain a sort of equilibrium or steady state, even when faced with external disturbances [15].

In their work, N. Qiao and colleagues demonstrate that it is feasible to reproduce the homeostatic plasticity of neural systems on a microchip, thereby ensuring that the computational abilities of neural systems are robust to long-lasting changes in their environment. The chip area occupancy is in the order of tens of micrometres, with a power consumption in the order of tens of nanowatts, using a standard $0.18\mu m$ CMOS process.

3.2 System Overview and Functionality

A circuit capable of achieving homeostasis on long timescales is achieved by adopting an automatic gain control (AGC) scheme of a DPI neuron. In [14], the authors showed that this technique allows the synaptic current flowing into the neuron to be maintained around a fixed value and that the AGC loop is capable of eventually restoring this equilibrium after that a perturbation has been introduced to the DPI input. Doing so ensure that the neuron response is robust to long-lasting changes in their environment. An overview of the system is shown in Figure 3.1.




Figure 3.1: Block diagram of the AGC Loop.

The AGC is composed of two blocks that together add feedback to the DPI.

The first block is a comparator circuit that compares the synaptic current – which is the DPI output current – to a reference current. If $I_{syn} < I_{ref}$, the comparator output is set to low, i.e., $V_{SW} = 0\text{ V}$. Otherwise, V_{SW} is set to V_{DD} .

The second block is the (ultra) low leakage cell (LLC). The DPI connected to the AGC is a p-type DPI, meaning that the PMOS and NMOS are swapped compared to Figure 2.1. As a result, the current I_{gain} can be expressed as:

$$I_{gain} = I_0 e^{\frac{\kappa(V_{DD} - V_{THR})}{U_T}}, \quad (3.1)$$

The LLC slowly adjusts V_{THR} based on the 'digital' comparison signal V_{SW} . When V_{SW} is low, the LLC slowly decreases V_{THR} , leading to an increase in I_{gain} and, consequently, bringing I_{syn} closer to I_{ref} . Conversely, when V_{SW} is high, the LLC slowly increases V_{THR} , decreasing I_{gain} and lowering I_{syn} closer to I_{ref} .

The variation of the output voltage V_{THR} is achieved by driving a current to charge or discharge a capacitor. The relationship between the current and the voltage across a capacitor yields that $\frac{dV}{dt} = \frac{I}{C}$. Hence, two options are possible to achieve small voltage variations leading to longer timescales: having a large capacitor or driving an extremely small current. However, as previously explained, large capacitors are undesirable because they inevitably lead to a large area occupancy. The Low-leakage cell thus utilizes multiple mechanisms to control the amplitude and direction of the current flowing through the capacitor.



16

3.3.1 The Low-Leakage PMOS

In the proposed circuit, the low-leakage PMOS is used as a controllable current source. The current I_{DS} flowing through it can be controlled both in direction and amplitude. In order to make the current flow into the direction that charges the capacitor and thus increases V_{THR} , the V_{DS} of the PMOS is set to a positive value. On the other hand, if the goal is to decrease V_{THR} , V_{DS} is set to a negative value. Both the source and the drain voltages are set using operational transconductance amplifiers (OTA) as depicted in Figure 3.2

The amplitude of I_{DS} is modulated by 2 factors: the drain-to-source voltage V_{DS} and the gate voltage V_G . Since, the aim of the LLC is to achieve ultra long timescales up to several minutes, the targeted orders of value for I_{DS} are from several tens of attoamperes to femtoamperes, i.e., approximately $10^{-17} - 10^{-15} A$.

Consistently driving such extremely small values from/into the drain of the transistor requires to minimize its leakage currents. The techniques used to reduce leakages are:

1. Finding the appropriate size, and in particular width-length ratio W/L , of the low-leakage PMOS to minimize the drain-to-gate current. Study [14] found that the optimal ratio is $W/L = 0.5 \mu m / 1 \mu m$.
2. Minimizing the leakage from the drain of the low-leakage PMOS to the connected gate of the input transistor of OTA2. This is accomplished by setting the W/L ratio of this transistor to $8 \mu m / 1 \mu m$, as found by [14].
3. Using an isolated well, i.e, ensuring that the bulk voltage V_B is the same as the drain voltage V_D to minimize the drain-to-bulk current.

It should be noted that the technology node used in [14] may be slightly different from the one used in this work, meaning that the physical models of the transistors may not match exactly. However, the simulations carried out in this work showed that the application of the methods described above does indeed result in negligible leakage currents. Therefore, the effects of varying the sizes of the transistors mentioned have not been further investigated.




Figure 3.3: Drain current of the low leakage PMOS for a DC sweep on V_s .




Figure 3.4: Drain current of the low leakage PMOS for different V_{DS} pulses.

Figure 3.3 and Figure 3.4 demonstrate how a drain current with controllable direction and amplitude in the range of attoamperes can be achieved. As it can be seen, reaching these extremely small current levels require setting the V_{DS} of the low-leakage PMOS to values below 1 mV . However, it is important to note that, unlike in these simulations where ideal voltage sources are used, the LLC sets these voltages through OTAs, which have limited accuracy.

The impact of the gate voltage V_G on the drain current I_D is also shown. In par-

ticular, it can be seen that setting a negative source-to-gate voltage, i.e. $V_G > V_S$, significantly increases the device resistance compared to a positive V_{SG} .

3.3.2 Operational Transconductance Amplifiers

The operational transconductance amplifiers OTA1 and OTA2 function as unity gain buffers. By incorporating a feedback loop between the output and the inverting input of the OTA, the output voltage is effectively clamped to the value applied at the non-inverting input. This configuration ensures that the output faithfully tracks the input voltage. This feedback mechanism is crucial for achieving precise control over the V_{DS} of the low-leakage PMOS, as discussed earlier, allowing the circuit to maintain extremely small current levels in the attoampere range.




Figure 3.5: Circuit implementation of the OTAs. Adapted from [14]

The accuracy of this mechanism is greatly influenced by the open-loop gain of the amplifiers. Indeed, starting from the basic equation of the operational amplifiers :

$$V_{out} = A_{OL}(V_{in(+)} - V_{in(-)}), \quad (3.2)$$

with A_{OL} the open-loop gain, $V_{in(+)}$ the non-inverting input and $V_{in(-)}$ the inverting input.

Then, replacing $V_{in(+)}$ by V_{in} and $V_{in(-)}$ by V_{out} (since they are connected together),

we obtain the formula for the closed-loop gain:

$$A_v = \frac{V_{out}}{V_{in}} = \frac{A_{OL}}{1 + A_{OL}}, \quad (3.3)$$

which tends to 1 as $A_{OL} \rightarrow \infty$.

In practice, however, the open-loop gain of any amplifier is always finite. Therefore, OTA1 and OTA2 have been designed to achieve the highest possible gain while keeping a reasonable circuit size. Their implementation, depicted in Figure 3.5 employs a pseudo-cascode scheme in order to enhance the open-loop gain, as suggested by [16].

Furthermore, in the case of OTA2, the input voltage remains fixed at V_{refM} during all the operation time of the LLC. However, concerning OTA1, its input varies depending on the output of the comparator circuit and the switch between the two possible values V_{refL} and V_{refH} can happen several times per second. It is thus important to examine the bandwidth of the OTAs.




Figure 3.6: Open-loop gain of OTA1 and OTA2.




Figure 3.7: Open-loop configuration of the OTAs

Figure 3.6 shows the open-loop gain of the OTAs described above. For this experiment, in order to perform simulations that illustrate the open-loop gain, the feedback loop is broken by introducing a large inductor ($1GH$). Indeed, adding such a large inductor in the feedback path allows to remove the AC component of the signal while ensuring that a DC operating point can be found.

The DC gain found in this configuration is $A_{OL} = 75dB$ which should result in a closed-loop gain very close to 1 as desired. The cutoff frequency is around $30Hz$ but as it is shown when analyzing the closed-loop gain, this is not representative of the behavior in a closed-loop configuration.

Re-establishing the feedback loop by removing the inductor and connecting the output back to the inverting input, a simulation showing the closed-loop gain is then performed.

Figure 3.8 shows the closed-loop gain of the OTAs which is indeed very close to 1 ($0dB$) for a wide range of frequencies, as it is desired, meaning that the output follows very closely the input.

3.4 Comparator circuit

As mentioned earlier, the role of the comparator circuit is to output a digital signal, low or high, depending on the value of I_{syn} compared to I_{ref} . Since the authors of [14] do not show the implementation of their comparator circuit, several designs




Figure 3.8: Closed-loop gain of OTA1 and OTA2.

have been tested in this work.

3.4.1 Ideal comparator

First, we observed the ideal behavior the circuit should have by simply using an ideal element: the voltage-controlled voltage source (VCVS). This dependent source outputs a voltage which is proportional to its input voltage with a tunable gain G_v , i.e.:

$$V_{out} = G_v V_{in}. \quad (3.4)$$

The VCVS compares voltages while the aim of the comparator is to compare currents and in particular, I_{syn} to I_{ref} . However, as it is the case with current-mirrors, diode-connected transistors provide voltages V_{syn} and V_{ref} that are images of the corresponding currents. Therefore, connecting V_{syn} to one of the input pins of the VCVS and V_{ref} to the other results in a faithful comparison between the two currents.




Figure 3.9: DC simulation of the ideal comparator using a VCVS.

Figure 3.9 shows a DC simulation of the VCVS where the gain has been fixed at a very large value (10^6) to ensure a digital behavior. The values of the input current I_{in} corresponding to I_{syn} are swept from 250 nA to 300 nA while keeping I_{ref} constant (250 nA). It can be noted that the linear region, which transitions from 0 V to V_{DD} is remarkably narrow ($\approx 2\text{ nA}$) thanks to the very high gain of the VCVS.

3.4.2 1-PMOS 1-NMOS inverter




Figure 3.10: Circuit implementation of the inverter.

The first tested implementation of a real comparator circuit that does not use any ideal element consist in a two-transistors inverter circuit as shown in Figure 3.10. The structure is the same as for the first branch of the sigmoid circuit presented in subsection 2.1.2. On the contrary to the case of the sigmoid circuit, it is desired for the comparator circuit the smallest possible linear region in order to produce an outputs which is as close as possible to a digital signal.

Figure 3.11 shows a DC simulation similar as the one performed with the VCVS. However, it has to be noted that tuning the sizes of the two transistors plays a crucial role in shaping the transition from one state to the other. In particular, small transistors lead in a very large linear region with a weak slope while larger transistor lead to a more digital behavior. In this case, large transistors refers to transistors that are both long and have a large width, as tuning the W/L ratio did not help to achieve a sharper transition.

Hence it can be seen that this design suffers from two major drawbacks. The first one is the requirement to have large transistors to achieve an input-output relationship resembling the one of a digital comparator, thus leading to a larger area occupancy. The second drawback is that even when using large transistors the linear region is ten times longer than in the case of the VCVS and is thus no longer negligible. It is shown in the later subsection 3.5.2 that this has a visible impact on the functioning of the circuit. It can also be noted than this design does not achieve a rail-to-rail output in the saturation regions but this has less impact on the functioning of the LLC.




Figure 3.11: DC simulation of the inverter including size effects.

3.4.3 Improved comparator




Figure 3.12: Full circuit implementation of the improved comparator. Adapted from [17].

The second design that have been considered, depicted in Figure 3.12, is a three-stages comparator circuit described in a book from R. Jacob Baker [17].




Figure 3.13: Pre-amplification stage of the comparator circuit. Adapted from [17].

The first stage, shown in Figure 3.13 is an input pre-amplifier circuit. Its role is, to amplify the differential input signal, which, in the case of the LLC consist in V_{syn} and V_{ref} , in order to improve the sensitivity of the comparison and reduce the impact of noise. The pre-amplification circuit is composed of a differential-pair amplifier that drives its two output currents i_{op} and i_{om} into active loads, i.e., the output currents flow through transistors M_x and M_y in the output branches. If the two input voltages are similar, $i_{op} = i_{om}$. However, the bigger the difference in input voltages the larger the corresponding current will be and conversely, the lower the current flowing into the other branch will be. Hence, if $V_p > V_m$, then $i_{op} > i_{om}$ while $V_p < V_m$ results in $i_{op} < i_{om}$.

It has to be noted that in the full circuit, the pre-amplifying stage is composed of both an n-type (as shown in Figure 3.13) and a p-type differential-pair amplifier whose outputs are summed together to feed the decision circuit.




Figure 3.14: Decision stage of the comparator circuit. Adapted from [17].

The second stage of the comparator, depicted in Figure 3.14 is the decision circuit. This circuit takes advantage of the cross-gate coupling between the transistors M_6 and M_7 to provide positive feedback in order to determine which input is the largest.

When $i_{om} \gg i_{op}$, meaning that $v_m \gg v_p$ at the pre-amplifier stage, then v_{om} is large enough to turn ON, through the gate connection, the transistors M_6 and M_7 . On the contrary, v_{op} turns OFF M_5 and M_7 . Therefore, in the left branch, the current i_{op} is almost entirely drawn by M_6 . Since, there is almost no current flowing through M_5 and that M_5 is in a diode-connected configuration, v_{op} is pushed to 0 V. Using a similar development for the right branch, almost all of the current i_{om} flows through M_8 and since it is diode-connected, v_{om} is pushed to a high value.

Then, if the difference between v_m and v_p decreases at the pre-amplifier stage, i_{op} and i_{om} get closer. Consequently, v_{op} increases, v_{om} decreases, thus, M_7 starts to draw more current away from M_8 . If all the NMOS transconductance are similar, the switching point is reached when $i_{op} = i_{om}$. Further increasing i_{op} and decreasing i_{om} effectively pushes v_{om} to 0 V and v_{op} to a high state.

In the full circuit, it can be seen that an additional diode-connected NMOS is connected to the sources of these transistors to ensure that the outputs of the decision circuit is a compatible input to the final stage of the full comparator circuit.

Finally, the third and last stage of this comparator circuit is the output buffer. Its role is to amplify the output signal of the decision circuit to provide a rail-to-rail digital output signal. It is composed of a p-type differential-pair amplifier followed by an inverter to further increase the gain.




Figure 3.15: Improved comparator DC simulation.

Figure 3.15 shows DC simulations of this improved comparator circuit. The blue curve refers to a circuit setup where all the W/L ratios are as indicated in [17]. As it can be seen, the input value for which the output is switched is shifted to the left by 10 nA compared to the expected value. By reducing the W/L ratio, which also allows to reduce the overall size of the circuit, the switching is re-centered and the transition becomes sharper as it is shown by the red curve. In this configuration, the size of the linear region is only 2 nA wide and the output is either at 0 V or V_{DD} outside of this region. This behavior is similar to the one observed with the ideal voltage-controlled voltage source shown in Figure 3.9 which shows that this improved comparator circuit is a good option for implementing the digital comparator of the automatic gain control loop.

3.4.4 Area efficiency comparison

Although the behavior of the comparator circuit described in subsection 3.4.3 is much better than the one of the simple inverter, it is interesting to compare the area occupancy of these circuits.

The inverter circuit is composed of only two transistors. However, as it has been discussed in subsection 3.4.2 the size of these transistor must be large enough to

provide a sufficiently digital input-output relationship. For this calculation and for the experiments of the full AGC loop using the inverter as the comparator, a width and a length of $6\text{ }\mu\text{m}$ is fixed, occupying an area of $12 \times 12\text{ }\mu\text{m}$. However, in order for the comparison to be correct, the transistors of the current-mirrors copying the currents I_{ref} and I_{syn} must also have the same size. Hence the total area considered increases to $24 \times 24\text{ }\mu\text{m}$.

As for the improved comparator circuit it has a combined width as high as $39\text{ }\mu\text{m}$ and a combined length of $21\text{ }\mu\text{m}$. In comparison, the total area occupied by the ultra-slow DPI of the previous neuron design is approximately $2000\text{ }\mu\text{m}^2$

Concerning the AGC, a clear trade-off appears between the quality of the circuit behavior and the total area it occupies. Some improvements could be made in this regard concerning the improved comparator. In particular, it could be further explored the effect of reducing the overall size of the transistors and it could be discussed the utility of having both an n-type and p-type differential-pair amplifier at the input stage of the comparator.

It is shown in the following sections that the AGC circuit allow to reach much larger timescales compared to the current design presented in chapter 2. In order to achieve these timescales with this design, it would be necessary to increase the circuit capacitance further, which would result in the use of capacitors of a considerable size. Therefore, in such cases, it may still be preferable to utilize the AGC loop from the perspective of the total circuit area.

3.4.5 The low-leakage cell with ideal switch

This section presents a first overview of the behavior of the LLC and shows that it is indeed possible to achieve ultra-long timescales by appropriately tuning the voltages applied to the low-leakage PMOS.

In the transient simulation shown in Figure 3.16, the input signal V_{SW} is set through an ideal voltage source to V_{DD} value for this first half of the simulation before dropping to 0 V for the second half. The gate voltage is set to $V_G = 1.5\text{ V}$ and it can be observed the effect of different V_{DS} on the dynamics of the LLC. The capacitor has a fixed value of 100 fF . The capacitor negative pin is connected to the drain of the LL-PMOS such that the drain current is always the equal to the capacitor current, i.e., $I_{DS} = I_{cap}$.




Figure 3.16: Temporal evolution of the LLC for different V_{DS} of the LL-PMOS.

When V_{SW} is high, the input transistor of OTA1 M_L is ON while M_H is OFF. V_S is then set at V_{ref_L} . Since $V_{ref_L} < V_{ref_M} < V_{ref_H}$, $V_{DS} = V_{ref_M} - V_{ref_L}$ is positive, hence $I_{DS} > 0$ and $I_{cap} < 0$. This current charges the capacitor C_f which leads to an increase in the output voltage V_{Thr} . On the contrary, when V_{SW} is low, $V_{DS} = V_{ref_M} - V_{ref_H}$ is negative, hence $I_{DS} < 0$ and $I_{cap} > 0$ which leads to a decrease in V_{Thr} .

The effect of the amplitude of V_{DS} on the speed at which the system evolves can be observed in Figure 3.16.

With a V_{DS} amplitude of $2mV$, the current drawn by the low-leakage PMOS is approximately $2.3fA$ and therefore $I_{cap} \approx 2.3fA$ (which cannot be observed in this

figure). This current charges the capacitor, resulting in an increase in V_{Thr} with a slope of approximately 23 mV/s . However, at this rate, V_{Thr} rapidly reach V_{DD} and the system stops with $I_{cap} = 0\text{ A}$ and V_{Thr} that saturates at V_{DD} . In practice, such a case has to be avoided as it would result in the AGC being unable to reach the reference current. Once, the input signal switched, the direction of I_{cap} changes and the capacitor starts to discharge, decreasing V_{Thr} with a similar slope.

Reducing V_{DS} reduces the amplitude of I_{DS} . As can be seen, a V_{DS} of $0.1 - 0.2\text{ mV}$ leads to a current in the order of hundreds of attoamperes. Lowering the current amplitude allows to slow down the output evolution. However, it should also be noted that the lower the V_{DS} the bigger the asymmetry in current amplitude between one switch state and the other and thus the bigger the difference between the increasing and decreasing slopes of V_{Thr} . Indeed, when $V_{DS} = \pm 0.2\text{ mV}$, I_{cap} is around $320/-160\text{ aA}$ and the increasing/decreasing slopes of V_{Thr} are $\approx +3.33/-1.54\text{ mV/s}$ while $V_{DS} = \pm 0.1\text{ mV}$ leads to $I_{cap} \approx +200/-30\text{ aA}$ and V_{Thr} slopes $\approx +2.24/-0.3\text{ mV/s}$.

Although it is also possible to increase the value of the gate voltage V_G to reduce the current drawn by the low-leakage PMOS, doing so further increases the amplitude difference between the positive and negative current I_{DS} .

In conclusion, it can be seen that the evolution of V_{Thr} reaches very large timescales. While it is not direct to deduce the timescale of a DPI whose gain would be controlled by this voltage, it is reasonable to assume that it would be in the order of several seconds. As a comparison, the timescale of the ultra-slow DPI of the neuron presented in chapter 2 is approximately 0.173 s . Also, It should to be noted that the capacitor value of 100 fF is much smaller than the capacitor that would be needed to achieve similar timescales using only the ultra-slow DPI of the previous neuron design.

3.5 Characterization of the AGC loop

With the individual components and main blocks of the circuit now thoroughly detailed, the characterization of the complete system is conducted. Transient simulations are first presented to observe the temporal evolution of the circuit's behavior, providing insight into how the different components interact dynamically. Following this, AC simulations are performed to characterize the circuit's frequency response,

examining its behavior as a filter.

All the simulations presented hereunder have been performed with similar circuit parameters and inputs. In particular, the input current is composed of a DC current of 60 nA and a large pulse of 100 nA is introduced from $t = 20\text{ s}$ to $t = 40\text{ s}$. The reference current of the comparator is maintained at $I_{ref} = 250\text{ nA}$. As for the main parameters of the low-leakage cell, $V_{ref_L} = 1.3800\text{ V}$, $V_{ref_M} = 1.3802\text{ V}$ and $V_{ref_L} = 1.3804\text{ V}$, ensuring that the V_{DS} of the low-leakage PMOS is always at $\pm 0.2\text{ mV}$, $V_G = 1.5\text{ V}$ and the value of the capacitor of the LLC is fixed at 100 fF .

3.5.1 The AGC Loop with an ideal comparator




Figure 3.17: AGC Loop with ideal comparator transient analysis.

First, the behavior of the AGC loop using the ideal voltage-controlled voltage source as the comparator circuit is shown in Figure 3.17.

At first, the synaptic current I_{syn} is lower than the reference current I_{ref} . Therefore, the comparator outputs a low digital signal. This low signal is transmitted to the LLC and OTA1 sets V_s to V_{ref_H} which results in a negative V_{DS} of the low-leakage PMOS. The low-leakage transistor thus pulls a current that discharges the capacitor which decreases the output voltage V_{Thr} . Finally, this voltage controls the gain of the DPI, causing an increase in the synaptic current. Hence, a full loop of the circuit has been performed. This behavior continues until I_{syn} becomes larger than I_{ref} .

Once the synaptic current reaches the level of the reference current, the comparator starts switching. In this 'locked region' (as described by [14]), V_s alternates between V_{ref_H} and V_{ref_L} . Although the capacitor current exhibits large spikes (several fA) which might be caused by the abrupt voltage changes, these transients are instantaneous and the average current flowing through the low-leakage PMOS becomes zero. Since the capacitor neither charges or discharges, V_{Thr} is maintained and the synaptic current stays constant around the reference current.

At $t = 20\text{ s}$ a perturbation is introduced by adding a step current to the input. Since the time constant τ of the DPI is very low, the synaptic current is nearly immediately pushed up. The comparator output now stays in a high state since $I_{syn} > I_{ref}$, the current flowing through the LL-PMOS starts charging the capacitor, leading to an increase in V_{Thr} and thus a decrease in I_{syn} .

The system enters once again the locked region around $t = 34\text{ s}$ before going back to the tracking mode when the input pulse is removed at $t = 40\text{ s}$.

3.5.2 The AGC Loop with the simple inverter comparator




Figure 3.18: AGC Loop with simple inverter transient analysis.

Figure 3.18 introduces the behavior of the AGC loop where the ideal comparator has been replaced by the simple inverter described in subsection 3.4.2.

In the tracking regions, the overall behavior of the AGC loop is very similar to the previous case. The main visible difference is found when I_{syn} approaches I_{ref} . Since the synaptic current approaches the reference very slowly, the inverter output enters its linear region and instead of reaching V_{DD} , it settles at an intermediate value. This is caused by the pair of transistors used for the input of OTA1. Indeed, these transistors are connected in a similar configuration that the inverter and outputs

either V_{ref_L} or V_{ref_H} depending on the digital signal it receives. However, in this case, the gate voltage of these transistors is an intermediate value. Consequently the value set at V_S by OTA1 is between V_{ref_L} and V_{ref_H} and eventually settles at V_{ref_M} . Therefore, V_{DS} is approximately zero, no current charges or discharges the capacitor and both V_{Thr} and I_{syn} become constants.

Even though this design exhibits a very similar behavior, regarding the evolution of V_{Thr} and I_{syn} , it should be noted that the value of I_{syn} for which the AGC stabilizes is slightly above the reference current ($\approx 10\text{ nA}$ in this case). This indicates that this design is less precise than the one using the ideal comparator.




Figure 3.19: Frequency response of the synaptic current to small variations of the reference current (with the simple inverter).

Looking at Figure 3.19, the filtering properties of this version of the AGC circuit are compared to the ones of the DPI responsible for the ultra-slow feedback in the neuron presented in chapter 2.

The plain lines are direct AC simulations from Cadence Virtuoso while the crosses are the measured synaptic current variations when introducing small sinusoidal variations of the reference current at different frequencies (every decades in this case). This 'handcrafted' AC analysis may be less precise but is still representative of the overall behavior for most frequencies and shows its usefulness in cases where Cadence AC analysis is irrelevant (which is shown later in subsection 3.5.3).

It can be seen that both the DC gain and the cutoff frequency are similar. After

the cutoff frequency, the slope of the AGC response is around -20 dB per decade. However, the AGC exhibits a rebound in its frequency response in the range of kHz to MHz before dropping again.

In the frequencies of interest of a neuron circuit, i.e., from DC to hundreds of hertz, the behavior of the AGC is very similar to the ultra-slow DPI which indicates the possibility of replacing it by the AGC circuit.

3.5.3 The AGC Loop with the improved comparator




Figure 3.20: AGC Loop with improved comparator transient analysis.

Replacing the simple inverter by the improved comparator of subsection 3.4.3, the switching behavior of the comparator output and of the current I_{cap} is restored as depicted in Figure 3.20. Furthermore, the accuracy of the tracking of the reference current is similar to the configuration that uses the ideal comparator.




Figure 3.21: Frequency response of the synaptic current to small variations of the reference current (with the improved comparator).

Figure 3.21 illustrates the frequency response of the AGC circuit using the improved comparator. In this case, it is evident that the direct AC analysis provided by the simulation tool misses the real behavior of the circuit. This completely different behavior may be explained by the fact that the fast switching nature of the comparator circuit induces instability and that the Cadence AC analysis may not be able to find the time at which the variations of synaptic current stabilizes. However, the 'handcrafted' AC analysis made by measuring the synaptic current variations after sufficient time, directly from a transient analysis allows to recover the low-pass filter behavior that is expected.

Although, the cutoff frequency is lower compared to the one of the ultra-slow DPI in this case due to the larger time scale of the AGC, the overall response is still faithful to the desired behavior with a similar slope in the frequencies of interest. Hence, this indicates that this version of the AGC design might also be a good replacement for the neuron ultra-slow DPI circuit. Further more, it should be noted that this difference in the cutoff frequency may appear because of the lower accuracy of the 'handcrafted' method.

4 A New Modulable Neuron with Slower Dynamics

The preceding chapter presented the automatic gain control (AGC) scheme and demonstrated that the output current of the DPI associated with the AGC circuit is a low-pass filtered version of the reference current. In particular, it has been demonstrated that the filtering properties of the DPI responsible for providing the ultra-slow negative feedback in the neuron circuit of chapter 2 are similar to those of the AGC loop.

Consequently, this chapter investigates the potential for replacing the ultra-slow DPI, which necessitates the utilization of a substantial capacitor, with the AGC circuit, which occupies a total area of approximately $960 \mu m^2$ compared to the $2000 \mu m^2$ of the ultra-slow DPI capacitors of previous design. However it should be noted that this size is only an estimation calculated with the transistors and capacitors sizes of the circuit and does not take into account the constraints set by a real physical layout. Hence, it is more likely that the total area occupied by the AGC is between 1500 and $2500 \mu m^2$. The new design provides a reduced overall capacitance and innovative tunable parameters that facilitate novel neuromodulation capabilities.

4.1 High-level description

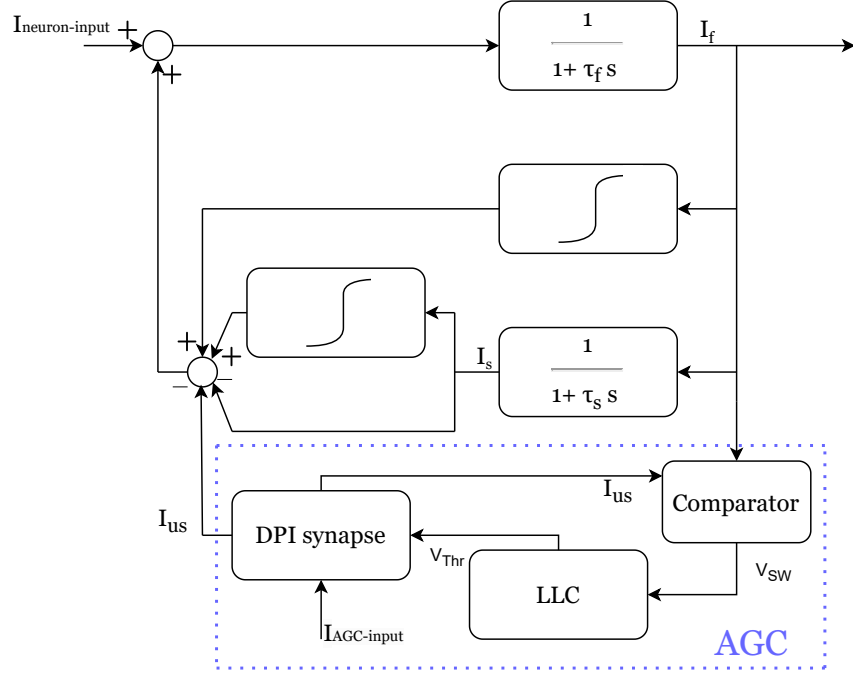


Figure 4.1: Block diagram of the new neuron design.

An overview of the new neuron design is depicted in Figure 4.1. As can be observed, the ultra-slow negative feedback is now provided by the AGC circuit. The copy of the neuron output is now directed to the comparator reference current pin and the low-leakage cell continuously adapts the gain of the AGC DPI so that its output current is a low-pass filtered version of the reference current. This effectively recreates the ultra-slow negative feedback current necessary for spiking and bursting.

4.2 Behavior and comparison with the previous design

Simulations of the new neuron design incorporating the AGC circuit have been conducted using both the simple inverter and the improved comparator circuits. The results indicate that both comparator configurations yield similar results in terms of the neuron's overall behavior. All subsequent simulations and analyses

presented in this section and in the following section have been carried out using the improved comparator but similar results are obtained with the simple inverter.




Figure 4.2: Neuron and AGC behavior with initial parameter set.

To evaluate the performance of the new neuron design with the AGC circuit, a first simulation is conducted where the parameters for the fast and slow feedback currents were kept identical to those used in the previous neuron design, which exhibited bursting behavior. Additionally, the AGC parameters are set to the same values as described in section 3.5. The results, depicted in Figure 4.2, show that under these conditions, the neuron produces a single spike followed by a prolonged period of approximately 28 seconds, during which the output current slowly decreases before eventually resetting. This behavior suggests that the timescale provided by the AGC

for the ultra-slow negative feedback is too large, resulting in the inability to produce realistic spiking. To achieve more plausible bursting patterns, the timescale of the AGC feedback must be reduced.

Following the observation that the initial AGC timescale was excessively long, the AGC parameters were carefully tuned to reduce this timescale to a more appropriate range. Figure 4.3 illustrates the results of this adjustment, demonstrating that with properly tuned AGC parameters, the new neuron design successfully exhibits bursting behavior. The figure highlights the combined dynamics of the circuit by displaying both the neuron output and the key signals from the AGC circuit.




Figure 4.3: Bursting neuron using the AGC Loop to provide ultra-slow negative feedback. $V_G = 1.4V$, $V_{ref_L} = 1.345V$, $V_{ref_M} = 1.365V$, $V_{ref_L} = 1.389V$, $AGC_{in} = 75nA$.

With the new neuron design now tuned to exhibit bursting behavior, a comparison with the previous design [6] is conducted using the same input and identical fast and slow feedback parameter values. Figure 4.4 presents this comparison, highlighting the differences in behavior between the two designs.

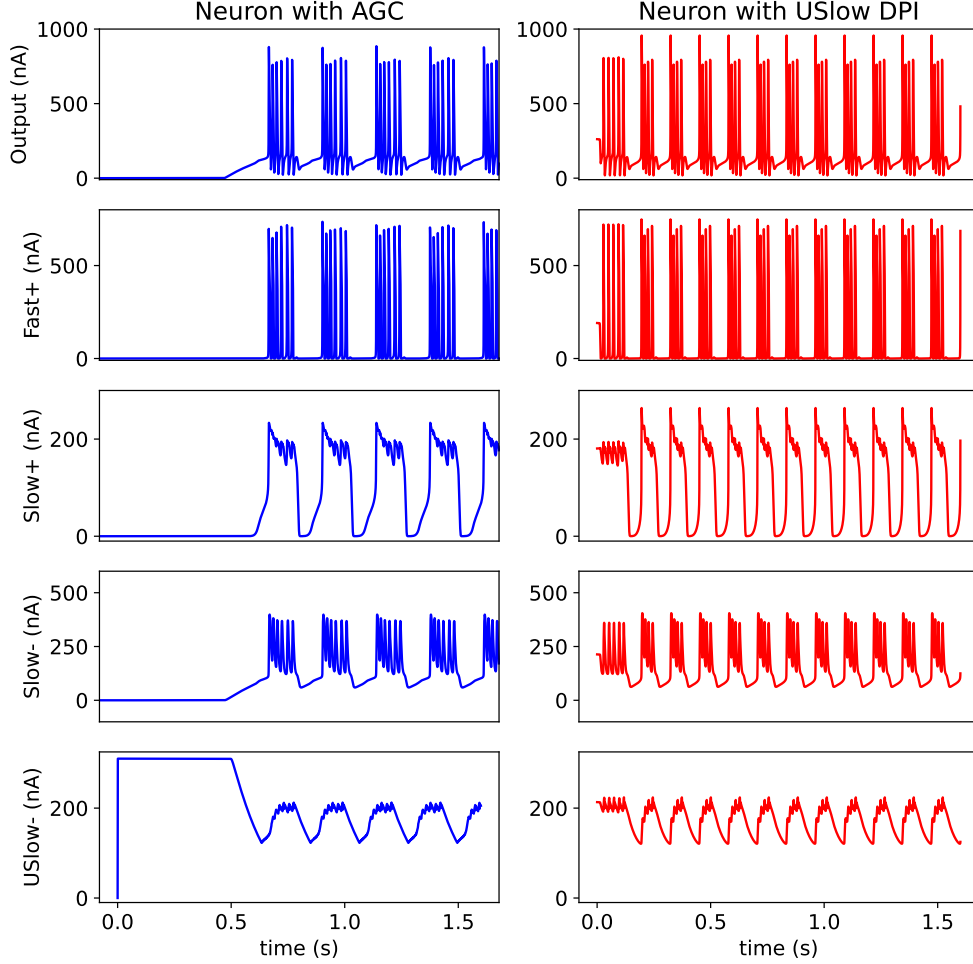


Figure 4.4: Comparison of the output and feedback currents of the two designs [13].

The figure reveals that the ultra-slow negative feedback current in the new design has overall slower dynamics compared to the old design. This slower dynamic has two significant consequences. First, the more gradual increase in the ultra-slow negative feedback current means that it takes longer to reach the reset threshold, which is the point at which the burst ends. As a result, the new design, using these parameter values, produces larger bursts, containing more spikes per burst. The second

consequence is that the slower decreasing slope of the ultra-slow feedback current extends the time between bursts, thereby reducing the overall bursting frequency.

4.3 AGC enabled neuromodulation

As demonstrated in the previous section, the behavior of the new neuron design can be completely altered by controlling the ultra-slow feedback current. This control is achieved simply by tuning the parameters of the Automatic Gain Control (AGC) circuit. By adjusting these parameters, it is possible to fine-tune the dynamics of the ultra-slow feedback current, allowing for a wide range of bursting patterns to be generated, enabling the neuron circuit to mimic various neural behaviors through parameter adjustments of the AGC alone.

As shown in Figure 4.4, there is an extended period where the neuron remains inactive due to the initial value of the ultra-slow feedback current being too high for the neuron to spike. To address this inhibition at the start, several approaches can be explored. One method is to increase the voltages V_{ref_L} , V_{ref_M} and V_{ref_H} equally, which decreases the initial value of I_{gain} from the AGC DPI, thereby lowering the initial ultra-slow negative feedback. Alternatively, a simpler approach is to reduce the constant input current of the AGC DPI, which also acts as a multiplying factor affecting the DPI output. Figures 4.5, 4.6 and 4.7 illustrate the effect of decreasing the AGC input on the initial value of the ultra-slow negative feedback current, demonstrating that this adjustment can significantly reduce the initial period during which the neuron is inhibited by the ultra-slow current.

It has to be noted that the initial oscillations that can be observed in these simulations are numerical instabilities which disappears once the low-leakage PMOS voltages V_D and V_S are correctly set by the OTAs.




Figure 4.5: Bursting neuron with $AGC_{in} = 100 \text{ pA}$

When the AGC input is set at 75 nA , the initial value of the ultra-slow current is around 300 nA which completely inhibits the neuron response. The gain adaptation of the AGC DPI takes around 0.5 s to increase V_{Thr} and thus V_{syn} to a value for which the output PMOS transistor of the AGC DPI operate in the subthreshold region. Afterwards, it takes an additional 0.1 s to lower the ultra-slow current to a level where the neuron can produce spikes.




Figure 4.6: Bursting neuron with $AGC_{in} = 100 pA$

Reducing the AGC input current to $100 pA$, it is shown in Figure 4.6 that the reduced input pushes V_{syn} up closer to V_{Thr} . Consequently, the output transistor of the AGC DPI enters the subthreshold regime way faster, which enables the neuron to burst already from $t \approx 0.2 s$.




Figure 4.7: Bursting neuron with $AGC_{in} = 25 pA$

A further reduction of the AGC input to $25 pA$ brings the initial value of the ultra-slow current down to approximately $100 nA$. At this current level, the neuron is not inhibited and immediately spikes. Subsequently, the AGC gradually reduces V_{Thr} , which in turn increases the ultra-slow current and allows for the rapid recovery of the stable bursting pattern.

It should be noted that modifying the input current of the AGC DPI has little influence on the neuron output behavior once it enters bursting as the difference in the AGC DPI integration is compensated by the V_{Thr} evolution.




Figure 4.8: Increased burst length by increasing V_G to 1.4005 V.

The simulation depicted in Figure 4.8 demonstrates the possibility to obtain a novel form of neural behavior. For this simulation, the gate voltage of the low-leakage PMOS has been slightly increased by 0.5 mV which has the effect of reducing both the ascending and descending slopes of V_{Thr} . As a result, the duration of the bursts is extended and the number of spikes per burst is doubled (from six to twelve).




Figure 4.9: Increased inter-burst period with two spikes per burst. $V_G = 1.4V$, $V_{refL} = 1.361V$, $V_{refM} = 1.365V$, $V_{refH} = 1.389V$

The final part of this section explores how to extend the inter-burst period, resulting in more realistic neural signaling, by independently adjusting the increasing and decreasing slopes of V_{Thr} through the tuning of V_{refL} and V_{refH} , respectively.

In Figure 4.9, it is shown that by bringing V_{refL} closer to V_{refM} , the decreasing slope of V_{Thr} is reduced. This change leads to the ultra-slow current reaching the burst reset threshold more quickly, thereby decreasing the number of spikes per burst. Additionally, the reduced slope extends the time before a new burst is triggered after the reset threshold is reached, lengthening the period between bursts, thus lowering the inter-burst frequency. In comparison, the burst frequency of the old neuron design is approximately nine bursts per second (cfr. Figure 4.4), whereas in this simulation, it is around 2 to 3 bursts per second for a similar input.




Figure 4.10: Increased inter-burst period with three spikes per burst. $V_G = 1.4V$, $V_{ref_L} = 1.361V$, $V_{ref_M} = 1.365V$, $V_{ref_H} = 1.387V$

Finally, Figure 4.10 illustrates that adjusting V_{ref_H} closer to V_{ref_M} reduces the increasing slope of V_{Thr} , which slows the rise of the ultra-slow current between spikes. As a result, the system takes longer to reach the burst reset threshold, effectively increasing the number of spikes per burst.

5 Conclusions and future works

5.1 Conclusions

This master’s thesis has introduced and developed a new neuromorphic neuron circuit capable of exhibiting long-timescale bursting behavior, bringing its dynamics closer to those observed in real neural systems compared to existing designs. The innovative approach of this new design allows for higher integration levels by eliminating the need for excessively large capacitors to achieve extended time scales, thereby enhancing its suitability for large-scale neuromorphic implementations.

The work began by presenting the neuromorphic continuous feedback neuron design developed by L. Mendolia, which served as the foundation and reference throughout this research. The key components of this design and its feedback structure were examined. The crucial role of fast, slow, and ultra-slow feedback loops in ensuring robust spiking behavior was highlighted, alongside the neuron’s capacity for neuro-modulation through appropriate tuning of its parameters. This initial analysis also identified significant limitations, notably the extensive area required by the capacitor associated with the ultra-slow feedback loop. The need for larger capacitors to mimic more realistic neural signaling, which necessitates longer time scales, further emphasized this limitation and motivated the search for alternative approaches to achieve ultra-slow feedback.

Subsequently, the thesis introduced the automatic gain control (AGC) scheme, drawing parallels to biological homeostasis mechanisms. The individual components of the AGC loop were presented, and the design of some of them was discussed. The interactions among these components were then observed, showing how they work together to achieve very long timescales while keeping the total capacitance low. The filtering characteristics of the AGC loop were studied, and the results demonstrated its low-pass filtering properties, making it a viable candidate to replace the ultra-slow DPI in the original neuron design.

The integration of the AGC circuit into the neuron design led to the creation of the new neuron model. It was shown that the area occupancy of the new design was reduced compare to the previous, especially when aiming at longer timescales. A comprehensive comparison between the behaviors of the new and previous designs under similar input conditions was conducted, showcasing the improvements and new capabilities introduced by the AGC. The new design not only matched the original in terms of spiking behavior but also expanded the range of achievable bursting patterns. By fine-tuning the AGC parameters, the neuron circuit exhibited a variety of longer timescales bursting behaviors, reinforcing the effectiveness of the AGC in replicating the ultra-slow feedback necessary for realistic neural dynamics, all while minimizing the circuit’s overall capacitance.

In summary, this thesis has successfully developed a novel neuromorphic neuron circuit that addresses the limitations of previous designs, particularly in terms of area efficiency and the ability to generate long-timescale neural behavior. The new design represents a significant step forward in creating more realistic and scalable neuromorphic systems, with potential applications in various fields requiring advanced neural modeling.

5.2 Future works

Future works could involve a deeper exploration of the circuit’s full parameter space, including the parameters associated with the fast, slow, and ultra-slow (AGC) feedback mechanisms. This exploration could lead to a deeper understanding of the neuron circuit’s behavior across different regimes and potentially uncover new operational modes. Additionally, optimizations at the transistor level should be pursued to enhance the circuit’s efficiency and integration. This includes refining transistor sizing to balance power consumption, area and targeted behavior, thereby improving the overall performance and scalability of the circuit for large-scale neuromorphic systems.

Bibliography

- [1] Carver Mead. Neuromorphic electronic systems. *Proceedings of the IEEE*, 78(10):1629–1636, 1990.
- [2] Eugene M Izhikevich. *Dynamical systems in neuroscience*. MIT press, 2007.
- [3] Giacomo Indiveri, Bernabé Linares-Barranco, Tara Julia Hamilton, André van Schaik, Ralph Etienne-Cummings, Tobi Delbruck, Shih-Chii Liu, Piotr Dudek, Philipp Häfliger, Sylvie Renaud, et al. Neuromorphic silicon neuron circuits. *Frontiers in neuroscience*, 5:73, 2011.
- [4] <https://capocaccia.cc/en/public/about/neuromorphic-engineering/>. Accessed: 15/08/2024.
- [5] Barrie Gilbert. Current mode, voltage mode, or free mode? a few sage suggestions. *Analog Integrated Circuits and Signal Processing*, 38:83–101, 2004.
- [6] Loris Mendolia, Chenxi Wen, Elisabetta Chicca, Giacomo Indiveri, Jean-Michel Redouté, and Alessio Franci. A neuromodulable current-mode silicon neuron circuit for robust and adaptive neuromorphic systems. In preparation, 2024.
- [7] Chiara Bartolozzi and Giacomo Indiveri. Synaptic dynamics in analog vlsi. *Neural computation*, 19(10):2581–2603, 2007.
- [8] Alan L Hodgkin and Andrew F Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, 117(4):500, 1952.
- [9] Julie Dethier, Guillaume Drion, Alessio Franci, and Rodolphe Sepulchre. A positive feedback at the cellular level promotes robustness and modulation at the circuit level. *Journal of neurophysiology*, 114(4):2472–2484, 2015.
- [10] Guillaume Drion, Alessio Franci, Julie Dethier, and Rodolphe Sepulchre. Dynamic input conductances shape neuronal spiking. *eneuro*, 2(1), 2015.
- [11] Alessio Franci, Guillaume Drion, and Rodolphe Sepulchre. Robust and tun-

- able bursting requires slow positive feedback. *Journal of neurophysiology*, 119(3):1222–1234, 2018.
- [12] Luka Ribar and Rodolphe Sepulchre. Neuromodulation of neuromorphic circuits. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 66(8):3028–3040, 2019.
- [13] Loris Mendolia, Alessio Franci, and Jean-Michel Redouté. Networks of low-power CMOS neuromorphic neurons with robust neuromodulation capabilities for intelligent and adaptable neuromorphic systems. In *SCIoI & ISAB Summer School 2023: “Embodied Intelligence – Perception and Learning in Nature and Robotics”*, 2023.
- [14] Ning Qiao, Chiara Bartolozzi, and Giacomo Indiveri. An ultralow leakage synaptic scaling homeostatic plasticity circuit with configurable time scales up to 100 ks. *IEEE transactions on biomedical circuits and systems*, 11(6):1271–1277, 2017.
- [15] Walter B Cannon. Organization for physiological homeostasis. *Physiological reviews*, 9(3):399–431, 1929.
- [16] Minhao Yang, Shih-Chii Liu, and Tobi Delbruck. Subthreshold dc-gain enhancement by exploiting small size effects of mosfets. *Electronics letters*, 50(11):835–837, 2014.
- [17] R Jacob Baker. *CMOS: circuit design, layout, and simulation*. John Wiley & Sons, 2019.