UCLouvain

École polytechnique de Louvain

Ultra-low-power frequency synthesizer for long-range 4G/5G radio IoT transceivers

Author: **Marco Antonio Gonzalez Gonzalez**
Supervisor: **David Bol**
Readers: **Denis Flandre**, **Danielle Vanhoenacker-Janvier**, **Maxime Schramme**
Academic year 2018–2019
Master [120] in Electrical Engineering

Abstract

One of the main challenges an Internet of Things (IoT) object has to overcome is to perform its functions of sensing, processing and transmitting information while operating with stringent power constraints. The transmission is the less energy efficient of these functions due to the physical limitations of the air as a communication channel. It becomes even more difficult for numerous IoT applications where the information needs to be transmitted over long distances. Thankfully, many of them do not have high data rate requirements and can exploit the Low-Power Wide Area Networks.

With the objective of optimizing the power of radio transceivers for this kind of IoT applications, this work studies the possible advantages that a back-bias-controlled ring oscillator (BBCO), as the central block of a phase-locked loop (PLL), could bring to long-range communications compared to other PLL types. Therefore, an ultra-low-power BBCO-based PLL is designed in the 28nm Fully-Depleted Silicon-on-Insulator (FD-SOI) CMOS technology and its power and phase noise performances are compared to state-of-the-art transceivers. A 30x reduction on the power consumption resulted, compared to the transceiver designed by Lachartre *et al.* which was used as a design reference in this work, in spite of an increase of 30dB of phase noise, but the specifications were still met. The advantage found of BBCO-based PLLs was thus the ability to go to lower power consumption than other PLL types. Further studies are recommended towards additional power savings and more in-depth studies of the PLL noise sources.

Table of Contents

Introduction

The Internet of Things (IoT) is a network of objects connected to the Internet, objects that are capable of sensing different physical properties of their environment and of processing and transmitting the information to a smartphone or a base station while operating with tiny batteries. These objects or smart sensors have as key attribute their ability of wirelessly connecting to the Internet thanks to the radio transceiver.

For numerous IoT applications, the smart sensor needs to communicate with a remote base station over a long distance. However, they still have a very limited power budget in order to ensure a long battery life. Fortunately, these applications do not usually require high data rates. This is the trade-off that Low-Power Wide-Area Networks (LPWAN) communications exploit, offering longer ranges than other standards like Bluetooth or Wi-Fi and lower power than cellular networks leveraging on the low data rate requirements of IoT smart sensor applications.

Unlike the sensing and the processing functions of these connected objects, the wireless communications are less energy efficient because of protocol overheads and the physical limitations of the air as a communication channel [1]. The impact of the radio transceiver on the power consumption of the object is thus very important, of which approximately between 20% to 30% corresponds to the power consumed by the frequency synthesizer block of the transceiver, as can be seen in Figure 1 where it is separated by the contribution of the oscillator (DCO) and the rest of the PLL. It is therefore of great importance to perform the frequency synthesis using very low power to reduce the total power consumed by the transceiver.

The goal of this thesis is to design an ultra-low-power frequency synthesizer in 28nm Fully-Depleted Silicon-on-Insulator (FDSOI) CMOS technology for RF transceivers implementing the wireless LPWAN communications in IoT smart sensors. To achieve ultra-low power, a special architecture of ring oscillator will be studied as the central block of the frequency synthesizer, a back-bias-controlled ring oscillator. This type of ring oscillator has already proven to have many advantages on short-range communications [2], such as ultra-low-voltage operation, low jitter, instantaneous startup and on-chip frequency tuning [1]. Here, its benefits for long-range communications will be studied.

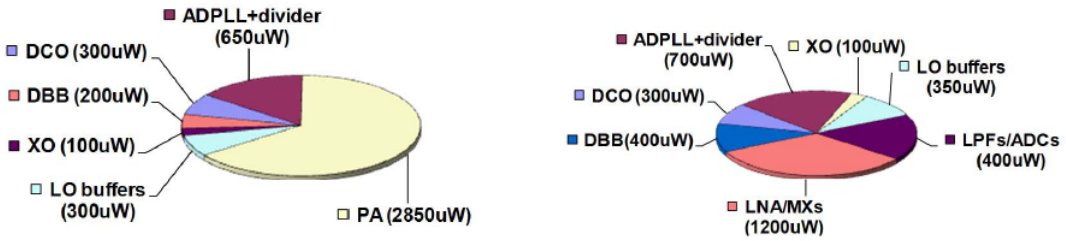

4.4mW TX @ -2dBm

(a)

3.7mW RX @ max. gain

(b)

Figure 1: Power breakdown at 1V of a Bluetooth Low-Energy (a) transmitter and (b) receiver (from [3])

The first contribution of this work is a complete study of the low-power techniques used in LPWAN, including a performance comparison of various commercial modules. The reason why Sigfox is the best adapted standard for the frequency synthesizer is shown, as well as giving the required specifications. A second contribution is the design methodology of both back-bias-controlled ring oscillators (BBCO) and BBCO-based phase-locked loops, along with a method to predict the performance of the latter. Finally, one last contribution is a frequency synthesizer in 28nm FD-SOI CMOS technology under 0.5V of supply voltage meeting the Sigfox specifications.

This master thesis is organized as follows.

- Chapter 1 gives an overview of the LPWAN main standards, comparing their specifications, advantages and energy performance. The main performance metrics are also presented along with the state of the art of transceivers using LPWAN.
- Chapter 2 discusses the design of a back-bias-controlled ring oscillator in the 28nm FD-SOI CMOS technology.
- Chapter 3 discusses the design of a PLL using the BBCO previously designed. Its performance is characterized and compared with state-of-the-art transceivers and the specifications of the Sigfox standard.

Chapter 1

Basics of LPWAN communications

Low-Power Wide Area Networks (LPWAN) have become a popular radio communication technology with the rapid growth of IoT [4] because the existing technologies do not support well long-range applications with low bandwidth requirements that are typical for IoT and machine-to-machine (M2M) scenarios [5]. Figure 1.1 shows a diagram with the location of the different radio communication technologies with respect to their data rate and range capacity. LPWAN technologies cover a previously vacant spot in this diagram.

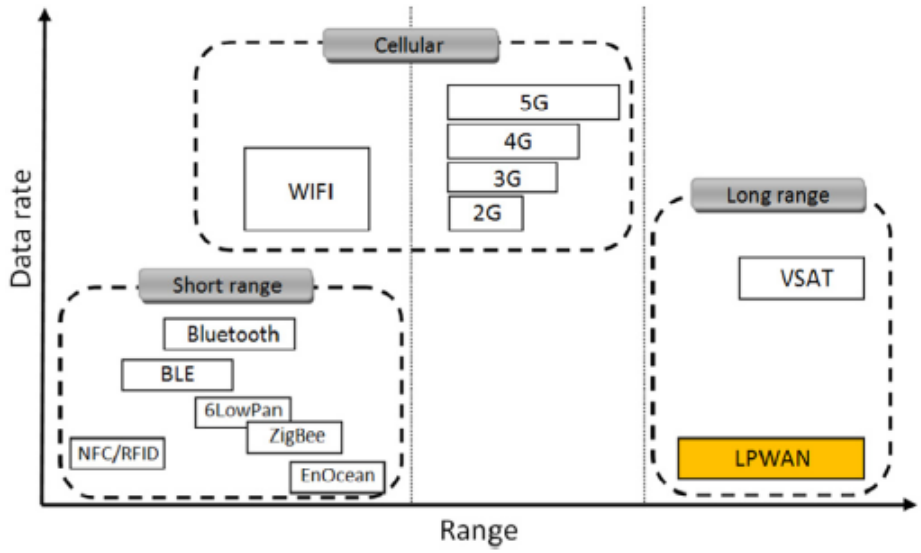

Figure 1.1: Required data rate versus range capacity of radio communication technologies: LPWAN positioning (from [4])

In this chapter, an overview of the main LPWAN characteristics and technical aspects will be covered, including a specification and energy comparison, along with some key performance metrics. This will allow to choose the best adapted standard for the frequency synthesizer that will be designed in further chapters.

1.1 Principles and key performance metrics

A summary of the main characteristics of LPWAN, the frequency band used in the communication and some key performance metrics are presented in this section.

1.1.1 LPWAN characteristics

Even though the applications of these networks are very diverse, there are some requirements that are common to all of them [6]. First, the connected object must be ultra low-power to prolong the battery life as much as possible and avoid costly battery change. High transmitted power is thus not acceptable to reach longer range. Also, the low cost of the object is important to allow a wide acceptance by the users. Finally, the data transfer between the object and the user must be fully secured.

1.1.2 Choice of frequency band

The frequency spectrum for wireless transmission is divided into two different kinds of frequency bands, those for which a license must be paid for transmitting in it and those license-free that are open for the public.

Each network selects, according to their own requirements, the type of band on which it will transmit.

Licensed bands

The main advantage of transmitting in licensed bands is the low congestion. The band of frequencies is reserved for the network so the level of interference is small. This guarantees a better Quality of Service (QoS) in the communications, metric which will be explained in Section 1.2.4. Also, the security and latency are assured.

The downside is the higher cost for the network operators, the little to none coverage in rural areas and the increased complexity, cost and power consumption in the connected object [6].

Unlicensed bands

Unlike the licensed bands, these frequency bands can be highly congested since they are free for any application to use it. The QoS is thus affected by the high level of interference and the latency is not guaranteed. As advantages, the coverage is better and the cost and power consumption of the object are lower.

The unlicensed frequency band is another important choice for the network to do. A too low frequency, around 169MHz, is not practical because of the large antenna size it implies. Going for the gigahertz frequencies solves the antenna size problem but the propagation is poor, the attenuation at one kilometer is 15dB higher than for the 863-870MHz band, the best choice with wider acceptance in Europe [6].

The output power in this band is mostly limited to 25mW and the duty cycle to 0.1%. Figure 1.2 shows the specifications of the 868-870MHz band in Europe. As it can be seen, the 868-868.8MHz band benefits from a higher duty cycle, which is why it is often used for uplinks (object to base station), while the 869.4-869.65MHz band benefits of a higher allowed output power, which is why it is often used for downlinks.

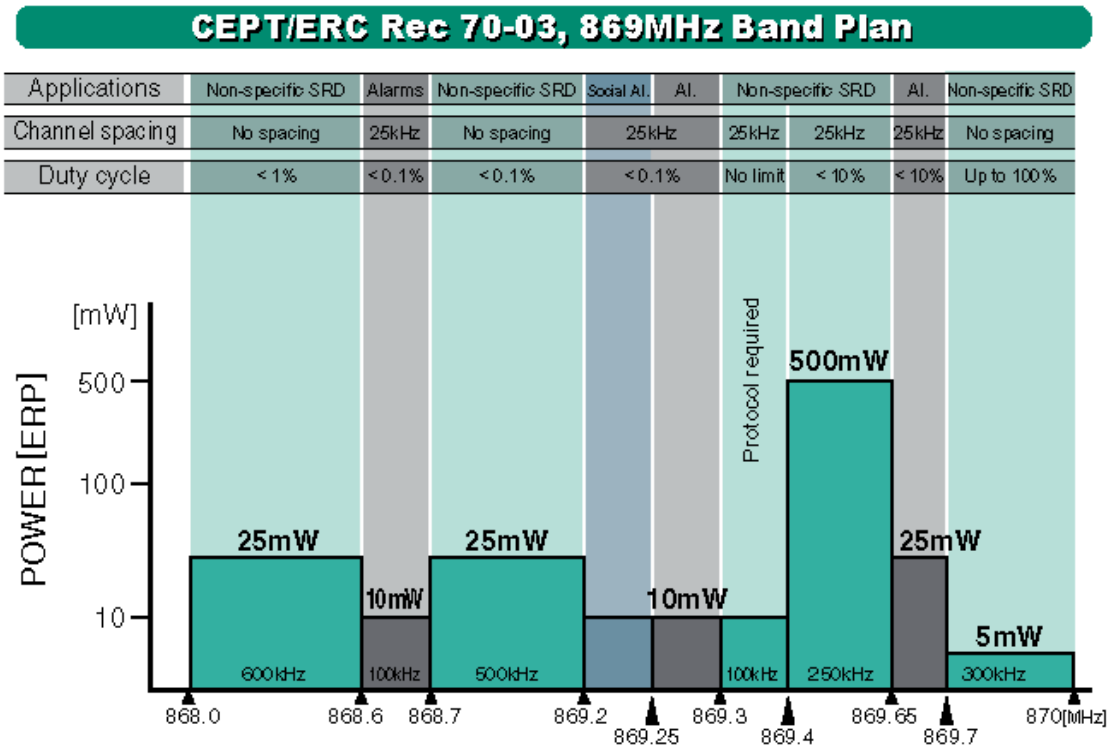

Figure 1.2: Unlicensed band in Europe in the 868-870MHz band (from [6])

1.1.3 Performance metrics

The main performance metric used to evaluate the performance of the circuits designed in the next chapters (oscillator and phase-locked loop) is the phase noise.

An ideal sinusoidal has an output with the general form $V_{out}(t) = Acos[\omega_0 t + \phi]$ [7] where A is the amplitude, ω_0 is the angular frequency and ϕ is a fixed phase reference. Its spectrum is thus a pair of impulses at the frequencies $\pm\omega_0$. However, in a practical oscillator, the output is generally of the form $V_{out}(t) = A(t) \cdot f[\omega_0 t + \phi(t)]$ where f is a function of period 2π. The spectrum of a practical oscillator has sidebands close to the frequency ω_0 [7].

This is usually characterized by the single sideband noise spectral density, measured in decibels below the carrier frequency per hertz (dBc/Hz). Furthermore, it is generally dominated by the phase variations, because the amplitude variations can be practically eliminated, which is why it is known as phase noise [7].

The phase noise affects differently the performance of the transmitter and the receiver in wireless communications.

Transmitter

In the case of the transmitter, there are two consequences of the phase noise in the frequency synthesizer, which are illustrated in Figure 1.3. The close-in phase noise, at low offset from the carrier frequency, reduces the Signal-to-Noise Ratio (SNR) of the output modulated signal because the signal energy is spread along a higher bandwidth than in the ideal case [8]. The far-away noise on the other hand, at high offset from the carrier frequency, creates spectral emissions outside the desired channel [8]. The out-of-band emissions are not allowed in the wireless protocols because it causes interference on other users. To avoid the problem, transmitters are required to meet the specifications of a spectral mask.

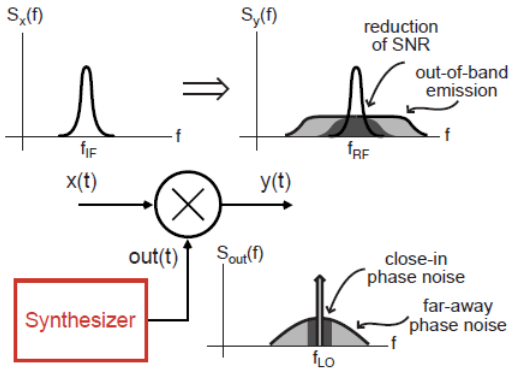

Figure 1.3: Impact on transmitters of the phase noise in the synthesizer (from [8])

Another important performance metric of a transmitter is its power efficiency, which is the percentage of the total power consumed during the transmission that is actually used in the transmitted signal.

$$P_{\text{eff}} = \frac{P_{out}}{V_{dd} \cdot I_{dd}} \times 100 \tag{1.1}$$

where P_{out} is the transmitter output power, V_{dd} is the supply voltage and I_{dd} is the current provided by this supply. The remaining percentage is the power consumed by the building blocks of the transmitter.

Receiver

The receiver is affected by the phase noise as shown in Figure 1.4. The receiver operates in the presence of multiple blockers that interfere with the desired signal by communicating in adjacent channels. The phase noise of the oscillator causes a part of the blocker's signal to come into the desired channel and increases the noise floor. Since the SNR must be increased to successfully receive the signal, the receiver's sensitivity gets deteriorated, i.e., it increases the minimum power required for a received signal to be successfully received.

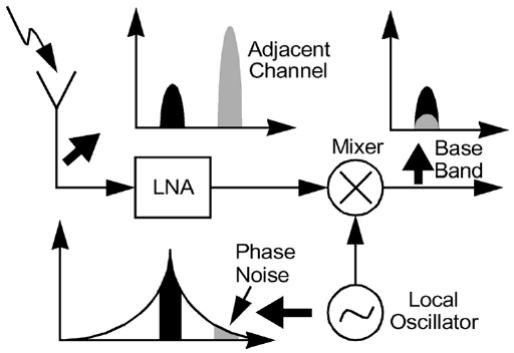

Figure 1.4: Effect of phase noise on RF systems, a simplified diagram (from [9])

1.2 Overview of LPWAN standards

In this section, a small introduction to three LPWAN standards will be presented, focusing primarily on their techniques to increase the range and decrease the power consumption. These standards are Sigfox and LoRa, that use unlicensed frequency bands, and NB-IoT that uses licensed ones.

1.2.1 Sigfox

The strategy used in the Sigfox standard is to communicate using ultra narrowband channels, typically 100Hz [10], and to provide long range and higher sensitivity, because the noise level is lower thanks to the very limited bandwidth, but with limited data rates [6], typically 100-600bps.

The modulation used in uplink is Differential Binary Phase-Shift Keying (DBPSK), a spectrally efficient modulation [10], that consists in changing the phase of the carrier wave depending on the sent bit. It is easily demodulated and it is not sensitive to absolute phase error. However, it has a limited QoS because it is not robust to fading [6].

Another advantage is the simplicity of the connected object. During uplinks, the object sends the same information using three different carrier frequencies [10] to leverage from the frequency diversity in the channel. The base station, on its side, takes all the complexity because it decodes multiple narrow-band channels without knowing the exact frequency of these channels [6]. For the downlinks, since higher power and duty cycle are authorized (see Figure 1.2), higher bandwidths and more traditional modulation schemes can be used, i.e. 600bps Gaussian Frequency-Shift Keying (GFSK) [6].

The main limitation of using such narrow channels is that the bandwidth is usually lower than the slow frequency drift of crystal oscillators [11]. The general solution for this problem is to use a Temperature Compensated Crystal Oscillator (TCXO) [6] [11], which increases the cost and power consumption of the connected object.

1.2.2 LoRa

The LoRa standard uses a different strategy to achieve long range. It is based on a custom modulation based on spread spectrum technology, Chirp Spread Spectrum (CSS) [6]. This modulation spreads the signal over a wider bandwidth, making it more robust to noise, which allows increasing the range of the communication. It is as well robust against multipath fading and jamming [4].

The spreading of the spectrum is achieved by the use of chirps, signals that vary continuously in frequency. Also, some redundancy is added to the transmitted signals to increase its robustness by using a variable error correction scheme. The data rate is given by the following equation [6]

$$R_b = SF \times \frac{BW}{2^{SF}} \times \frac{4}{4+CR}\ bits\ per\ second \tag{1.2}$$

where SF is the spreading factor, BW is the modulation bandwidth and CR is the code rate.

The spreading factor is used to control the trade-off between the data rate and the range [4]. A higher spreading factor allows longer range while a lower one allows a higher data rate and lower power consumption. A larger SF enhances also the sensitivity, becoming equivalent or even better than Sigfox, but in this case more bandwidth is used for a limited data rate [6]. This can be considered as the main limitation of LoRa because it decreases the network capacity. It can be easily overcome by adding more base stations as higher values of spreading factor will no longer be needed [6].

There are three different classes of LoRa end devices to address requirements like latency of the wide range of IoT applications [4]: Class-A devices have the lowest power because downlink messages can only be sent after an uplink message; Class-B devices have extra receive windows at scheduled times; and Class-C devices have almost continuously open receive windows.

1.2.3 NB-IoT

Narrowband Internet of Things (NB-IoT) is an LPWAN standard developed by 3GPP, optimizing the LTE protocol for long battery life [10]. The main power saving features implemented in NB-IoT are the Power Saving Mode (PSM) and the Extended Discontinuous Reception (eDRX), which can be combined to obtain a maximum of power saving in the device.

Figure 1.5 shows a typical power consumption diagram of a device using NB-IoT with PSM. First, in the active mode, the device transmits information to the base station. It is in this mode when the most power is consumed. It is followed by an idle mode, called discontinuous reception, where the device goes to reception mode each couple of seconds to listen for possible messages. Finally, if there are no messages to receive, it activates the PSM mode and goes to deep sleep. During this time, which can last from a couple of seconds up to a few days, the device is not reachable but it is still registered to the network. The device only needs to re-establish the connection if there is no data transfer after the time is up [12].

Figure 1.6 shows the same kind of diagram but of a device using eDRX mode. In this case, the idle mode of the device is more flexible. The periodicity of the reception mode activation can be extended from the usual two seconds up to a few hours. The difference from the PSM mode is that the device goes from time to time in the receiving mode to check for messages; whereas on PSM mode, the device must go to full active mode and make a new setup of the connection to check for messages [12].

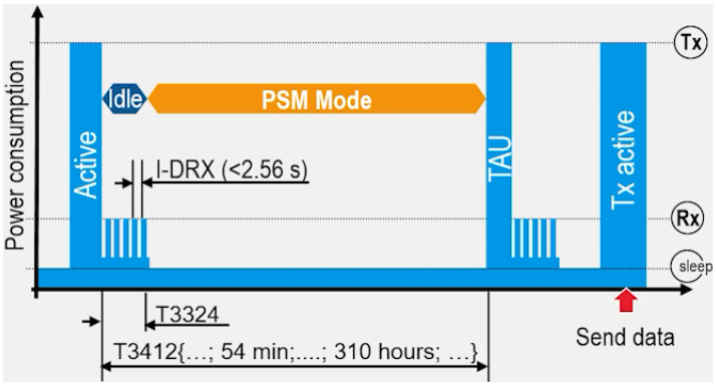

Figure 1.5: Power Saving Mode diagram (from [12])

1.2.4 Comparison

Figure 1.7 shows a comparison made in [4], in terms of some IoT factors, between the three standards that have just been introduced. Each IoT factor will be explained in the following.

Quality of Service

Quality of Service (QoS) is a measurement of the service performance that determines the degree of satisfaction of a user of the service, the communication. It can be quantified by parameters as data rate, delay, jitter and packet loss rate [13].

NB-IoT has a clear advantage over the others on the QoS because it transmits in licensed frequency bands, which guarantees QoS at the expense of the cost.

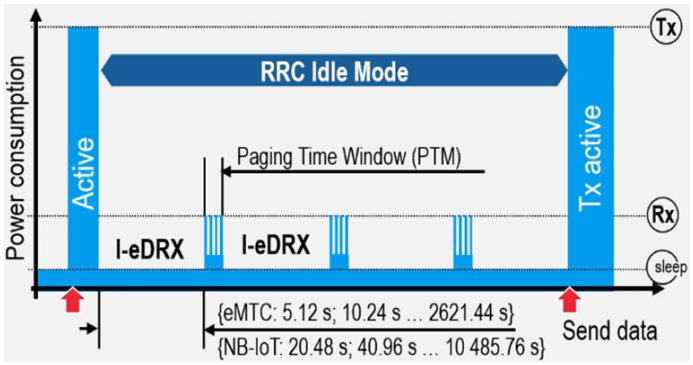

Figure 1.6: Extended Discontinuous Reception diagram (from [12])

Battery life and latency

Due to the synchronous communication, QoS handling and the higher current demand of the OFDM modulation, NB-IoT has a lower end-device battery life than Sigfox or LoRa [4], even though objects connected to these networks are in sleep mode most of time.

The lowest latency, i.e., time needed for the information to arrive from the transmitter to the receiver, is offered by NB-IoT. The class-C of LoRa provides a better latency performance than Sigfox at the cost of increased energy consumption [4].

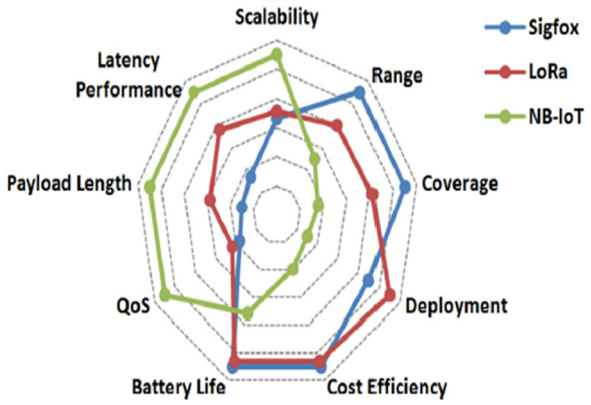

Figure 1.7: Respective advantages of Sigfox, LoRa, and NB-IoT in terms of IoT factors (from [4]).

Scalability and payload length

The scalability, the capability of handling an increasing number of connected devices, is also better in NB-IoT because it allows up to 100k connected devices per cell compared to 50k per cell for Sigfox and LoRa [4].

Another advantage offered by NB-IoT is the maximum payload length. It allows to transmit up to 1600 bytes of data. LoRa allows up to 243 bytes and Sigfox is the most limited because it allows a maximum of only 12 bytes [4].

Network coverage and range

The main advantage of Sigfox is the network coverage because the range can exceed the 40 km, needing only a couple of base stations to cover large areas. LoRa is second, offering a range of 20 km maximum, and NB-IoT is the last with a maximum of only 10 km [4].

Deployment model

Thanks to their maturity, LoRa and Sigfox can be already deployed in a large number of countries, unlike NB-IoT that will need more time to be established because its specifications were released only in 2016 [4].

Cost

Considering the spectrum, deployment and device costs, Sigfox and LoRa are more cost-efficient than NB-IoT [4].

These factors must be evaluated for each application in order to find the standard that fits best to its requirements. Applications where guaranteed QoS, high data transmissions and low latency are required should choose NB-IoT; whereas applications without these requirements that need higher range at lower cost should use LoRa or Sigfox.

1.3 Comparison of commercial LPWAN modules

To compare more in depth the LPWAN standards, the energy efficiency of some commercial modules will be analyzed. In [14] such a study was performed and their results are shown in Figure 1.8. It shows NB-IoT as clear winner because it has the lowest transmission energy per byte, followed by LoRa and finally Sigfox.

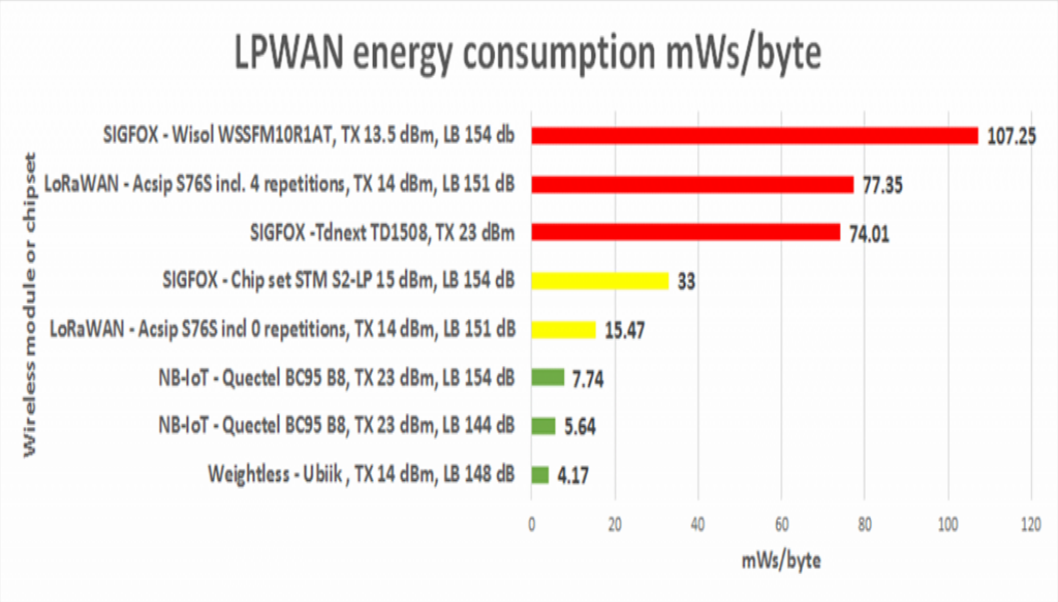

Figure 1.8: Commercial modules comparison (from [14])

To confirm these results and to perform a deeper analysis of these modules, the power consumption data shown in Table 1.1 were extracted from each datasheet.

	Wisol [15]	TDnext [16]	STmicro [17]	AcSiP [18]	Quectel [19]
Protocol	Sigfox	Sigfox	Sigfox	LoRa	NB-IoT
Supply voltage (V)	3.3	3.3	3.3	3.3	3.6
TX current (mA)	65	230	20	65	230
Output power (dBm)	13.5	23	14	14	23
Power efficiency (%)	10.44	26.29	41.86	11.71	24.1
Data rate	12 bytes in 6s 2bps	12 bytes in 1155ms 10.39bps	12 bytes in 6s 2bps	13.87bytes/s 13.87bps	147bytes/s 147bps
Energy/byte (mJ/byte)	107.25	74.01	33	15.47	5.64

Table 1.1: Commercial modules comparison

The values obtained by [14] were corroborated by these data after normalizing to the data rate, which allows a more fair comparison between the technologies. For the Sigfox modules in Europe (resp. in United States), each packet takes 2 seconds (resp. 385ms) to be transmitted, and each package is sent three times. The data rate for the LoRa modules was calculated using the LoRa Calculator software of Semtech [20]. No data about NB-IoT data rate was reported in [14], thus the data rate reported in the table is the value needed to obtain the corresponding energy per bit of this module.

It can be noticed that the power consumption is very similar for all the modules, but the data rate presents a significant variation. This represents a great advantage for NB-IoT, which has the highest data rate of the three protocols, but at the same time a great disadvantage for Sigfox whose data rate is the lowest.

The transmission power efficiency of these modules, computed using Equation 1.1, is shown in Figure 1.9. This value has also been reported in Table 1.1 from the datasheet information. All modules show the same tendency to increase the power efficiency as the transmitted power increases. As the power consumed by the building blocks of the transmitter remains more or less unchanged when the transmitted power is increased, a higher percentage of the total power consumed will belong to the signal power, increasing the efficiency. Comparing the different modules, the module from STmicroelectronics that uses Sigfox has a much better performance than the rest, which have a relatively close performance.

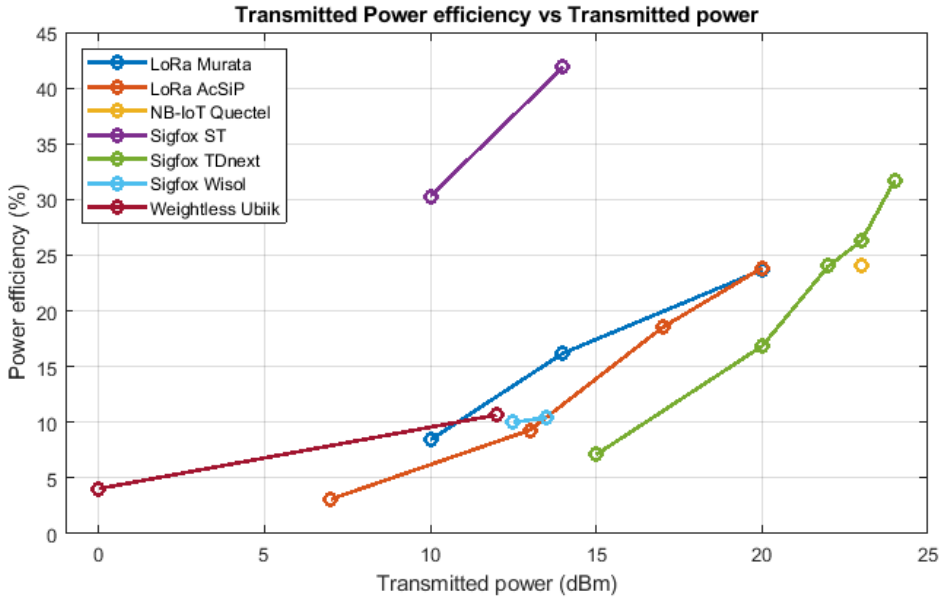

Figure 1.9: Transmitted power efficiency versus transmitted power

In order to compare the receiver performance, the receiver power was normalized to the data rate to obtain the energy per bit and have a more fair comparison between the standards, which have very different data rates. Figure 1.10 shows the obtained results versus the receiver sensitivity. As the sensitivity is closely related to the noise level of the receiver, it is logical that a better sensitivity is achieved at the expense of a higher energy per bit, which is what can be seen in the figure. NB-IoT has the best performance in this aspect.

It is interesting to notice how the spreading factor of LoRa can be used to control the trade-off between sensitivity, range and energy per bit. A lower spreading factor has a much lower sensitivity performance than a higher one because it is used for lower ranges to save power consumption, while the higher spreading factor is used for long range with a higher power consumption, thus having a better sensitivity.

From this comparison of commercial modules, NB-IoT proves to be the most energy efficient of the three LPWAN standards, which is the same conclusion reached by the study in [14]. However, the specifications for achieving such a performance are much too difficult to fulfill using ring oscillators at a low power consumption. Moreover, it is interesting to increase the energy efficiency of the other two standards given that their target applications are different from those of NB-IoT. Therefore, the latter is no longer considered from now on.

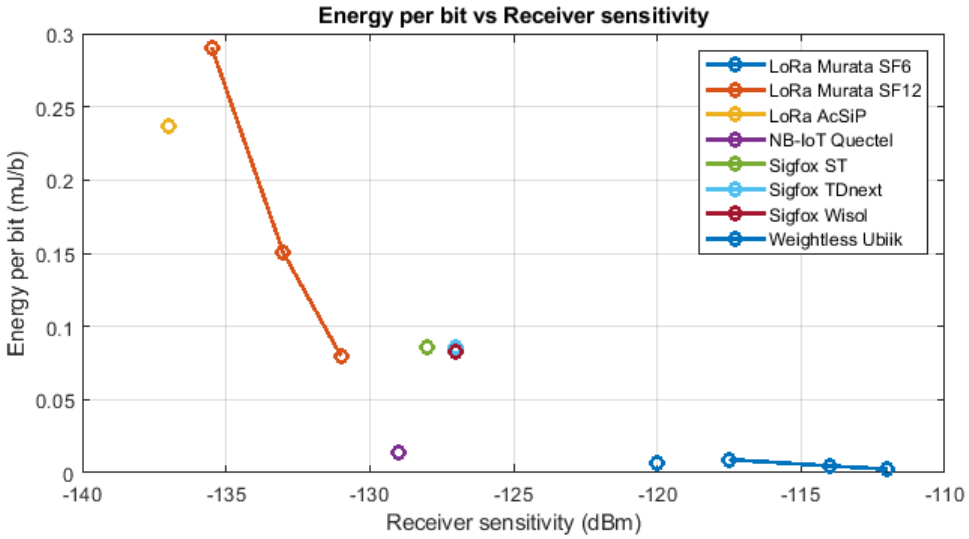

Figure 1.10: Energy per bit versus receiver sensitivity

1.4 State-of-the-art of LPWAN transceivers

Table 1.2 summarizes the state of the art of the LPWAN transceivers. Even though NB-IoT has a very good performance in the different transceiver aspects, it is not considered from this moment on, as explained before, because the transmission is made in the gigahertz band, where the ring oscillator that will be designed in Chapter 2 does not have a good enough performance to comply with all the specifications of this standard.

It is important to note that in the literature only LC resonant oscillators are used in LPWAN transceivers because they have superior phase noise performance compared to ring oscillators at the same power consumption [7]. But ring oscillators have the advantage of consuming less power, which is why the proposed frequency synthesizer is based on these oscillators. However, the values obtained in the literature for the phase noise are all much lower than what is expected to be obtained with a ring oscillator with the power consumption reported in the table.

The chip from Lachartre [11], which uses Sigfox, has the most achievable phase noise of the table and also it uses an integer-N PLL, which is simpler to design. Moreover, it has the interesting characteristic of functioning without a TCXO and is still able to achieve the frequency stability requirements of Sigfox. This decreases both the cost and the power consumption of the chip. This is why the Lachartre's chip is used as a reference for the design of the frequency synthesizer that uses the Sigfox standard described in Chapter 3.

	Lachartre[11]	**STmicro[17]**	**Wisol[15]**	**TDnext[16]**	**Semtech SX1272[21]**	**Semtech SX1276[18]**
Protocol	Sigfox	Sigfox	Sigfox	Sigfox	LoRa	LoRa
Technology	65 nm	-	-	-	-	-
Modulation	DBPSK/GFSK	GFSK	GFSK	GMSK/(G)FSK, OOK	(G)MSK/(G)FSK, OOK	(G)MSK/(G)FSK, OOK
Datarate	100bps	300bps	600bps	500bps	1.2kbps	1.2kbps
Sensitivity	-136dBm	-128dBm	-127dBm	-127dBm	-123dBm	-128dBm
RX energy/byte	18.13mJ	33mJ	107.25mJ	74.01mJ	3.2mJ	15.47mJ
TX power efficiency	12.63%	41.86%	10.44%	26.29%	21.59%	11.71%
Phase Noise @1MHz	-106.42dBc/Hz	-117dBc/Hz	-	-	-120dBc/Hz	-122dBc/Hz
Synthesizer Power	8.65mW	1.26mW	-	-	14.85mW	19.14mW
Type of PLL	Integer-N	Fractional-N	-	Fractional-N	Fractional-N	Fractional-N
Settling time	-	75μs	-	-	60μs	60μs

Table 1.2: State-of-the-art of LPWAN transceivers

1.5 PLL specifications for a ULP Sigfox transceiver

For the design of a frequency synthesizer that can be used in practice on a Sigfox chip, it is important to know the standard specifications that must be fulfilled. Figure 1.11 shows the transmission spectral mask, where TS_{UL} is the uplink symbol period, i.e., 10ms. As it was explained Section 1.1.3, this is the specification for the phase noise that a transmitter must fulfill to comply with the Sigfox standard.

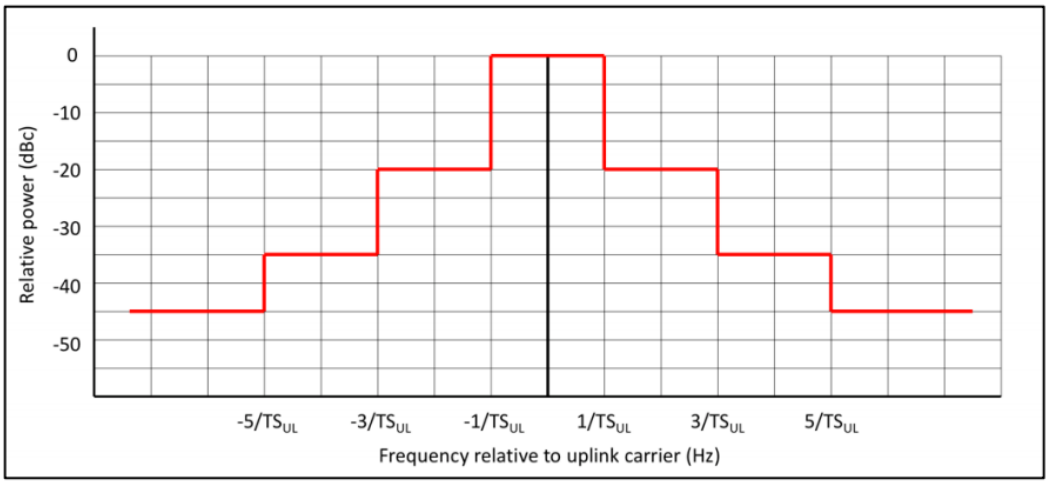

Figure 1.11: Transmission spectrum mask for Sigfox (from [22])

The mask shows the power relative to the carrier that the transmitter must fulfill. To convert this into a power spectrum, it must be divided by the resolution bandwidth used to measure the emissions of an object [23], which is 10Hz in this case. This means that the spectrum has the same form as this mask but it is placed 10dB lower than in the figure.

Chapter 2

Trade-off between power and phase noise of sub-GHz BBCOs

Ring oscillators are a cascade of an odd number of inverters whose output is fed back to the input. Due to the finite time it takes a single inverter to charge or discharge its output, the circuit oscillates at a frequency given by $f = \frac{1}{2Nt_d}$, where N is the number of inverters, or stages, that compose the oscillator and t_d is the delay of each stage. This delay depends on the supply voltage (V_{dd}), the inverter output current (I_{on}) and the total capacitance seen at the output (C_L), $t_d \propto C_L V_{dd}/I_{on}$.

The power consumed by such an oscillator is N times the power of each inverter, which is given by $P_{inv} = C_L V_{dd}^2 f$. Its phase noise, on the other hand, is inversely proportional to the number of stages [7]. To design a low-power oscillator, V_{dd}, N and C_L should be minimized, but it would result in a very poor phase noise performance.

In this chapter, this trade-off between power and phase noise will be studied through the design of a back-bias-controlled ring oscillator in the 28nm FD-SOI CMOS technology. This is a special ring oscillator architecture made available by this specific CMOS technology whose advantages, already proven for short-range communications [2], will be studied on a phase-locked loop for the frequency synthesizer of a Sigfox connected object, targeting a frequency of 868MHz. A key parameter for the design of PLLs, the gain of the oscillator, will be introduced and also used as guideline, along with the performance metrics explained in Chapter 1, for selecting the best performing oscillator to use in Chapter 3. Finally, a discussion on the different ways of implementing the oscillator will be presented to give a vision on how future works could further decrease the power consumption of BBCO-based PLLs.

2.1 Oscillator architecture

The frequency of a ring oscillator is controlled by changing the value of the inverter output current. Different architectures differ in the way this control is made possible. The usual architecture is the current-starved one, meaning that extra transistors are added to bias the inverters and their bias voltage is the control of the oscillator. A single delay stage of such a ring oscillator is shown in Figure 2.1, where transistors M1 and M2 form the inverter, biased by M3 and M4, while M5 and M6, the biasing branch, control and distribute the current for all the stages of the oscillator. Its main disadvantage is the power consumption that, being

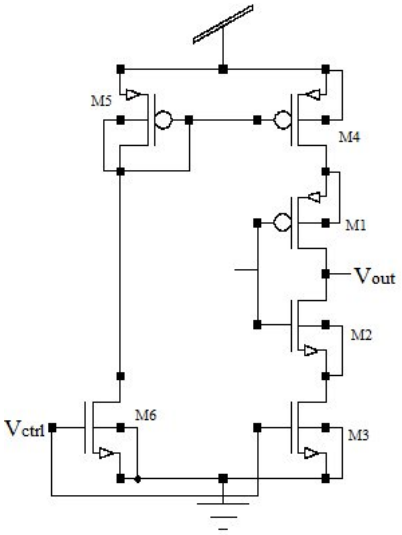

Figure 2.1: Delay stage for current-starved ring oscillator schematic (from [24])

computed as $P = V_{dd} \cdot I_{on}$, it is higher than in other architectures because of the limitation on the supply voltage. Due to transistors M3 and M4, the supply voltage cannot be lower than their drain-source voltage in order to keep them on. This higher supply voltage limits the down scaling of the power in the current-starved architecture. Furthermore, the biasing branch also consumes power, decreasing the power efficiency because this power is not actually used for the oscillations.

Another type of ring oscillator can be made by taking advantage of the forward back biasing (FBB) capability of the Fully-Depleted Silicon-On-Insulator (FD-SOI) CMOS technology [2]: the oscillator is controlled via the back voltage of the inverters instead of using additional transistors. Changing the back biasing modifies the threshold voltage, which in turn changes the current and thus the frequency. This configuration of ring oscillator is called back-bias-controlled ring oscillator (BBCO).

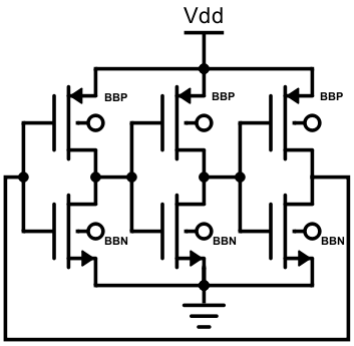

Figure 2.2: Back-bias-controlled ring oscillator schematic with three stages

2.2 BBCO design

To design such a circuit, whose schematic is shown in Figure 2.2, there are eight degrees of freedom: the width and length of both types of transistor; the nominal body voltages, BBN and BBP, for nMOS and pMOS, respectively; the supply voltage, Vdd; and the number of stages, N. In order to decrease the number of degrees of freedom, some design decisions are made. First, all lengths have the same value. Second, in order to match the rise and fall times of the oscillations, the pMOS transistors are 2.5x wider than the nMOS. Finally, the BBP voltage is the opposite of BBN. This last decision was made to follow the FBB specifications, depicted in Figure 2.3, for the 28nm low threshold voltage (LVT) FD-SOI CMOS technology in order to avoid turning on the parasitic diodes shown in the figure.

Additionally, the width of the nMOS (resp. pMOS) has a value of 200nm (resp. 500nm). These values were chosen because they are large enough to limit fabrication inaccuracies coming with minimal-width transistors, e.g., increased

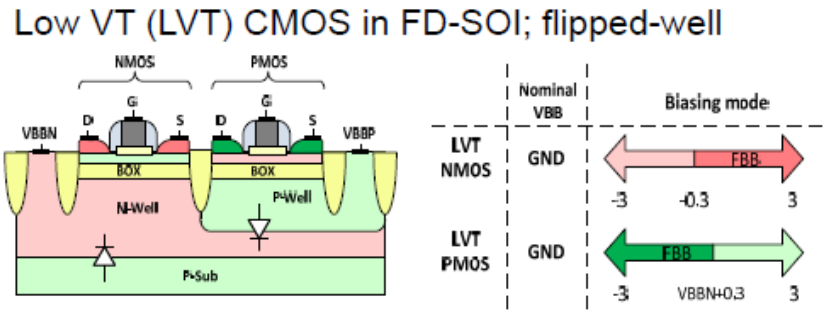

Figure 2.3: Cross-section of an LVT CMOS in FD-SOI and the Forward Body Bias ranges (from [25])

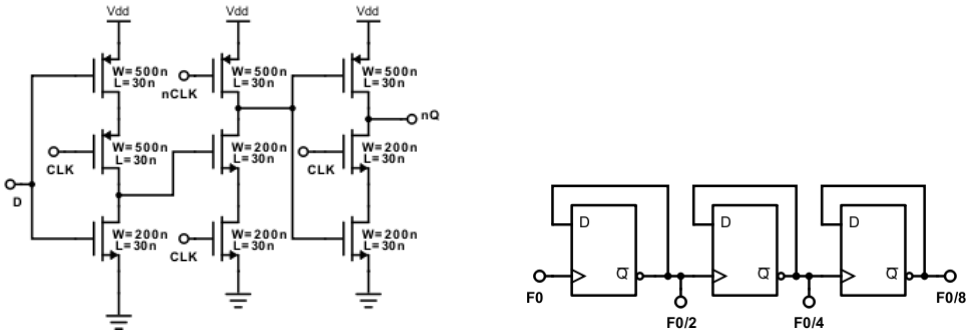

(a) **(b)**

Figure 2.4: Schematics of (a) TSPC flip-flop used for the frequency dividers and (b) frequency divider

mismatch and narrow-width effects where the threshold voltage depends on the width, but low enough to avoid having too large transistors to reach the target frequency of 868MHz. Also, BBN (resp. BBP) has a value of 1.5V (resp. -1.5V) to have a symmetrical control range from 0V to 3V (resp. -3V). All this reduces to three the degrees of freedom in the design: the length of the transistors, the supply voltage and the number of stages.

An additional degree of freedom appears if the use of frequency dividers is considered. The oscillator's frequency can be boosted, hoping to obtain a better phase noise performance as proposed in [26], and a frequency divider can follow the oscillator to maintain the desired output frequency. In this case, only multiples of the target frequency that are a power of two will be considered to avoid a complex division operation.

The frequency dividers used are composed of True Single-Phase Clock (TSPC) flip-flops, whose schematic is shown in Figure 2.4(a), where the transistors are minimal-length and the same values for the widths and the body biases as for the oscillator were taken. The advantage of such dynamic topology of a flip-flop resides on its high-speed capability as well as lower area and power [27].

To perform a frequency division by two, the output node of the flip-flop, the nQ node, is connected to the input, the D node, while the clock whose frequency is going to be divided is connected at the CLK node. To further divide the frequency by a factor two, more flip-flops can be connected in cascade, the output of one stage will be the CLK node of the next one. The schematic of a frequency divider is shown in Figure 2.4(b).

2.3 Results

In the rest of the chapter, the contribution of each of the four previously mentioned degrees of freedom will be studied to find the best trade-off between power and phase noise for the BBCO. Additionally, their impact on the gain of the oscillator, a key parameter for the design of PLLs, will also be studied to choose the BBCO that will be used for the PLL designed in Chapter 3.

In order to analyze more easily the influence of each one, we start by fixing the number of stages, with a value of 3, and the supply voltage, to 1V, while varying the frequency of the oscillator and the length of the transistors. These values are chosen because three is the minimum number of stages a ring oscillator can have and 1V is the nominal supply voltage of the CMOS technology used in the design. Then, the parameters that were set to fix values are progressively included to the results. The number of stages to evaluate are three, five, seven and eleven. The supply voltage takes the following values: 1, 0.75 and 0.5V. The frequency of the oscillator varies in multiples of the target frequency that are a power of two, i.e., one, two, four and eight times 868MHz, at each step increasing the number of frequency dividers by one. The length of the transistors changes in each case to reach the target frequency, with an accuracy of $\pm 0.5\%$, at the output of the last divider. The resolution of the length at fabrication being limited to 1nm, prevents the exact frequency from being precisely obtained.

2.3.1 Power

Figure 2.5 shows the power for each of the four cases of the oscillator's frequency, specifying the portions corresponding to the oscillator and to the frequency dividers. As expected, more power is consumed when the oscillator's frequency is higher

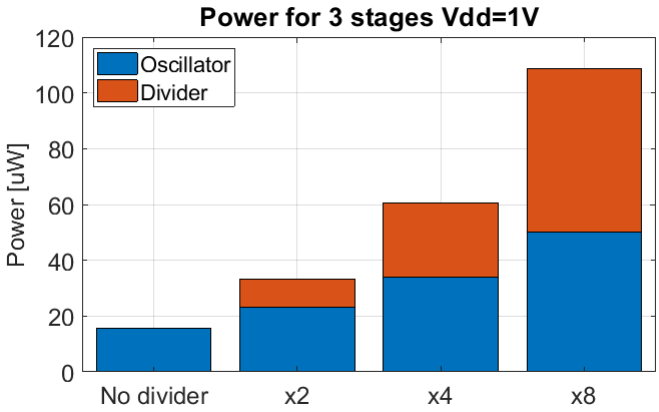

Figure 2.5: Power breakdown for three stages and 1V of supply voltage

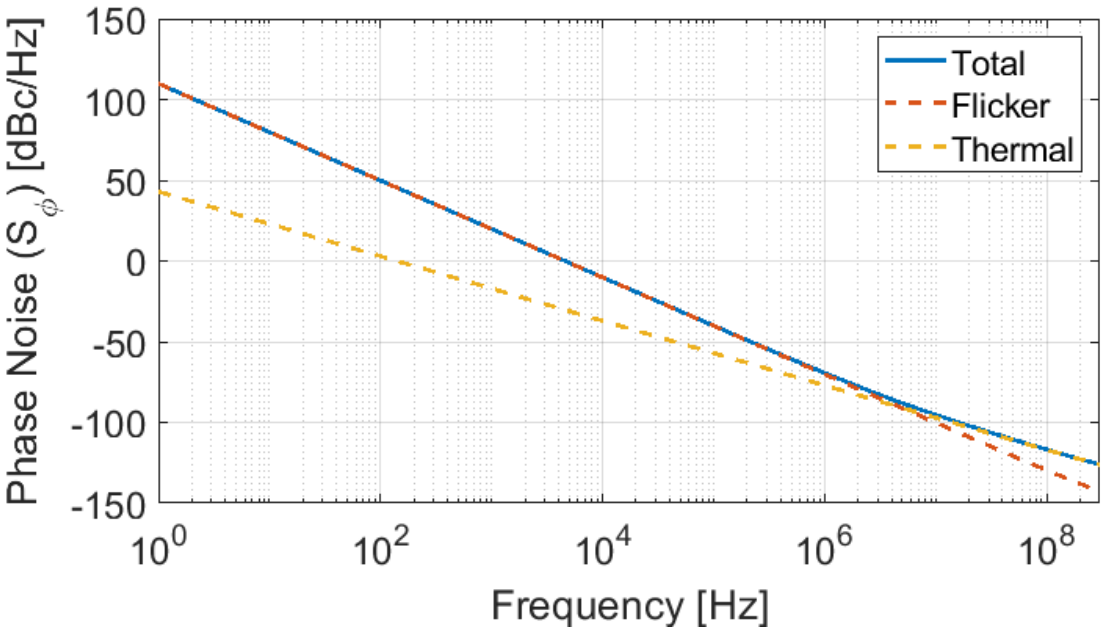

Figure 2.6: Typical plot of the phase noise of an oscillator versus offset from carrier

because more current is needed to reduce the delay of each stage. This relation is not proportional as doubling the frequency does not imply a doubling of the power consumed; instead, the increment is around 50%. Such behaviour can be explained by the fact that not only the current is being changed, but also the total capacitance, since the length of the transistors is being reduced to have a higher current. A higher current translates into higher power, while a lower capacitance reduces it.

On the other hand, the power of the dividers has a much higher increment, around 165% and 120% respectively. The reason for this is the different operating frequencies of each flip-flop, the power consumed by a digital circuit being directly proportional to the frequency at which it runs. Flip-flops further down the frequency division chain operate at a frequency that is half that of the previous flip-flop, the first one operating at the oscillator's frequency. As a consequence, for the last case, on top of having three dividers, they are functioning at higher frequencies than at the previous cases. Therefore, the dividers go from representing 30% of the total power to more than 50%.

2.3.2 Phase Noise

As already explained in Section 1.1.3, the phase noise is the spectral density of the oscillations' phase variations, measured in dBc/Hz. Figure 2.6 shows the typical plot of the phase noise of an oscillator versus the offset from the carrier frequency, obtained with the Steady-State Noise analysis of Eldo. It is important to notice the three main characteristics of this plot: the two noise regions and the corner frequency between them. The flicker-noise region, has a slope of -30dB per decade and the thermal-noise region, has a slope of -20dB per decade. The corner frequency is located at approximately 4MHz, which is the frequency where the total phase noise changes of slope because it passes from being dominated by the flicker noise to being dominated by the thermal noise.

The results for the phase noise of the same four cases studied before with the power are shown in Figure 2.7 at 1MHz and 100MHz of offset from the carrier frequency to evaluate its behaviour in the flicker and thermal-noise regions, respectively.

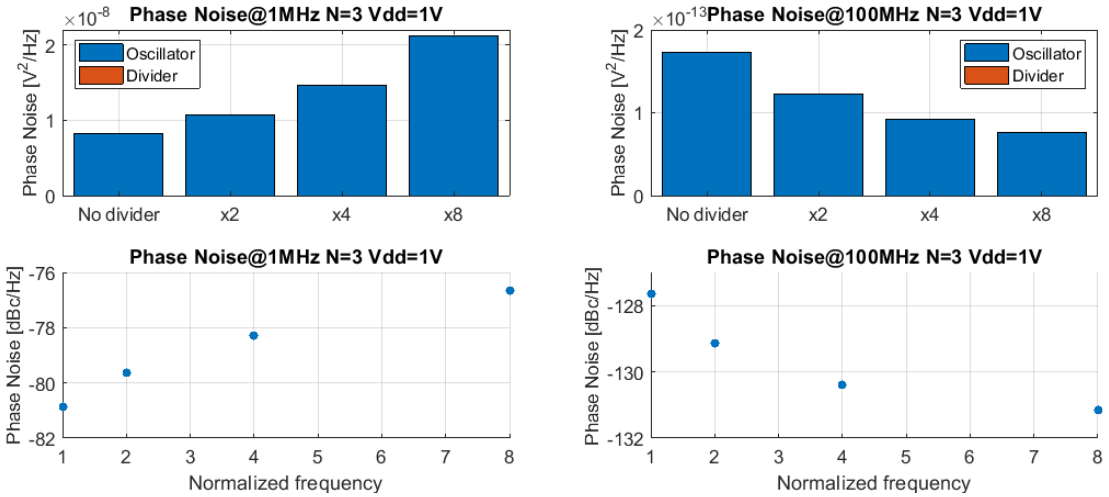

Figure 2.7: Phase noise breakdown for three stages and 1V of supply voltage

The top plots show the phase noise in natural numbers and squared in order to highlight its source. A negligible contribution from the noise sources of the frequency dividers can be noticed. The bottom plots, on the other hand, show the total phase noise in decibels. Theoretically, the phase resulting from a frequency division by N is N^2 times lower than at the input because the absolute jitter is not affected by this operation, as it was just shown, but the clock period is divided

by N [28]. This expressed in dB gives:

$$\mathcal{L}^o(f) := \mathcal{L}^i(f) - 20log(N)$$

A divide-by-2 operation lowers the phase noise by 6dB.

It can be noticed in the bottom plots of the figure that the phase noise in the flicker noise region increases as the operating frequency is increased, meaning that the frequency boosting worsened the phase noise of the oscillator more than the 6dB gain obtained from the frequency division, resulting in the contrary effect of what was intended; whereas, at the thermal noise region the results were as intended, a gain in the phase noise was still obtained from the frequency boosting. Nevertheless, in both zones, the difference between each multiple of the frequency is lower than 2dB, which is not a significant variation.

2.3.3 Figure of Merit

To better analyze the trade-off between power and phase noise and to include the rest of the parameters of the oscillator in the study, the figure of merit described in Equation 2.1 and proposed by [26] was used to quantify it.

$$FoM[dBc/Hz] = L(df) + 20log(df/f_0) + 10log(P/1mW) \tag{2.1}$$

The ideal case is to have a figure of merit which is flat because a higher power consumption enhances the noise performance of the oscillator. Also, it should be as low as possible because it would imply a good performance both in power and phase noise. Figure 2.8 shows its variation for the cases that have been studied until now.

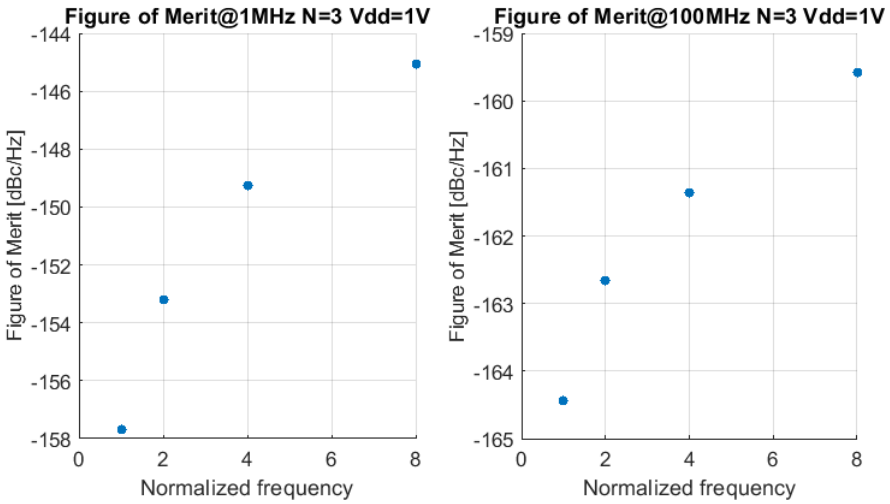

Figure 2.8: Figure of Merit for three stages and 1V of supply voltage

For both noise regions, the same trend can be noticed: as the frequency gets higher, the worse the figure of merit becomes. The gain in power when going to higher frequencies counteracts the gain in phase noise obtained in the thermal noise region, while at the flicker noise region the difference is worsened. The best performance is thus obtained with the oscillator running at the target frequency, no frequency dividers used.

If the number of stages is also varied, the power and phase noise follow the same trend that was already discussed for Figure 2.5 and Figure 2.7, with the difference that the total power is increased while the phase noise is decreased as more stages are included, which is why the results are not shown. For the figure of merit, the results are shown in Figure 2.9.

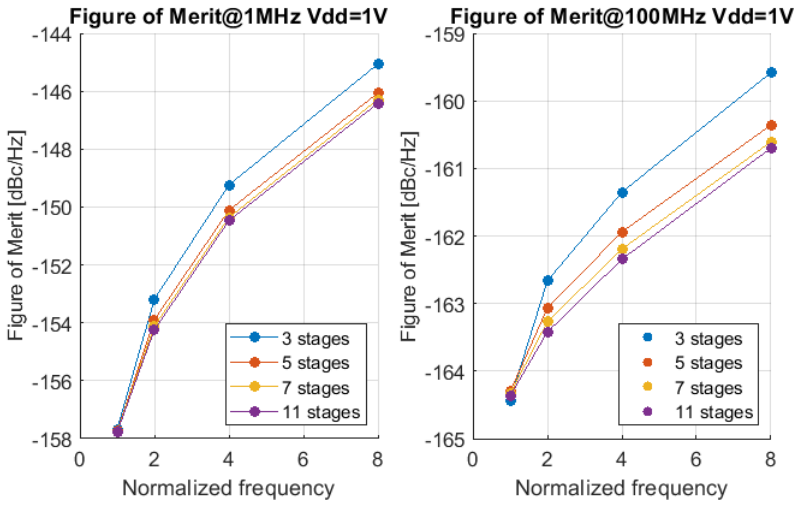

Figure 2.9: Figure of Merit for 1V of supply voltage

The same trend can be observed as in Figure 2.8, but an improvement for higher number of stages can be noticed. The improvement in the phase noise compensates the higher power consumed in the oscillator as the number of stages is increased. The best in this case is to maximize the number of stages, but the difference with regards to the oscillators with less stages is not significant, as it is always lower than 2dB.

Finally, when decreasing the supply voltage to $0.75V$ and $0.5V$, a significant decrease in the power is obtained but the phase noise worsens. Again, these results show the same trend as Figure 2.5 and Figure 2.7 and are not shown. With a reduced supply voltage, higher frequencies are harder to achieve because the flip-flops start failing. Figure 2.10 shows the evolution of the figure of merit including all the degrees of freedom of the oscillator. The same trend as before can be seen for all the different Vdd values, the lowest value having the worst performance of

all but with a difference lower than 3dB between them.

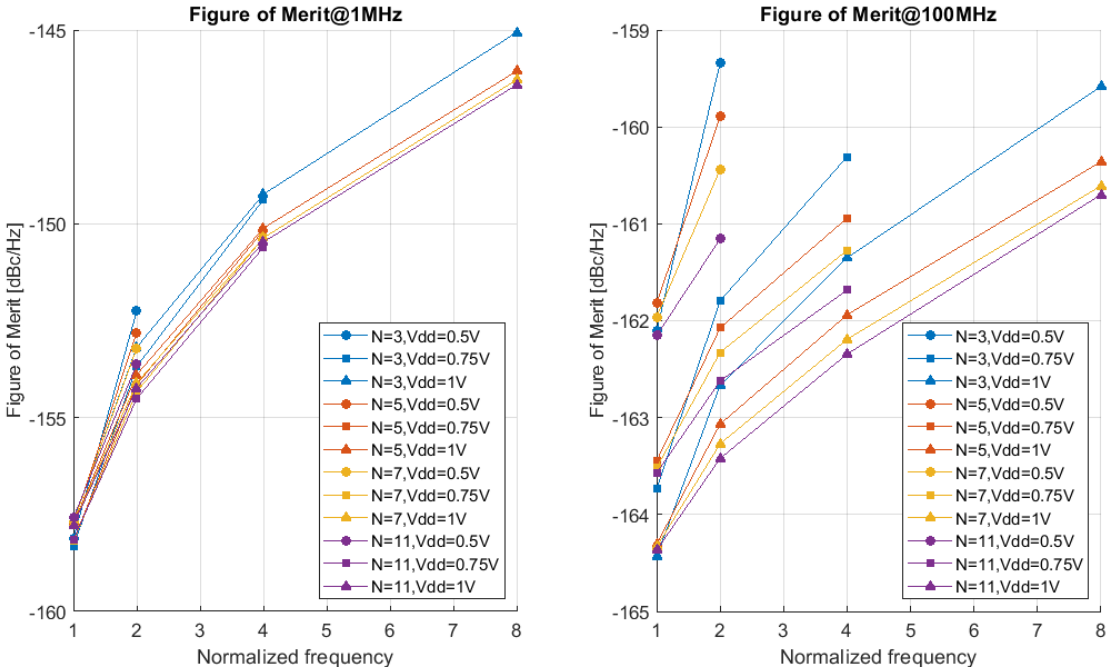

Figure 2.10: Figure of Merit for all cases

2.3.4 Oscillator gain

A voltage-controlled oscillator (VCO) has the typical response shown in Figure 2.11, the nominal frequency is normally achieved at the center of the control voltage range. Between the minimum and maximum values, the response is linear and characterized by the equation

$$F_{out}(t) = K_{VCO}v(t) \tag{2.2}$$

where K_{VCO} is the oscillator gain and the slope of the curve, $F_{out}(t)$ is the output frequency and $v(t)$ is the control voltage. This modelling of the VCO facilitates the design of Phase-Locked Loops (PLL), as it will be seen in Chapter 3.

When this gain is extracted from the VCO at twice the nominal frequency, i.e. with one frequency divider, while varying the supply voltage and the number of stages, the results shown in Figure 2.12 are obtained.

The reason for using twice the frequency is because a divide-by-2 oscillator should be used to create the two phase components (0° and 90°) needed for the quadrature modulation, as inputs to the I and Q paths. Higher frequencies are not desired because, as already seen before, they do not provide a better power/phase noise trade-off than lower frequencies.

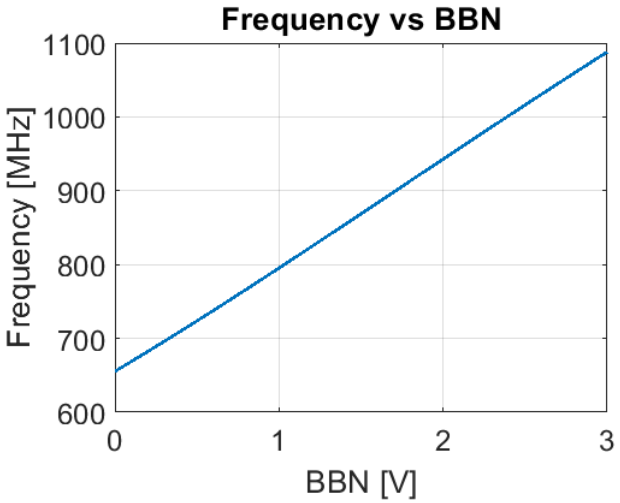

Figure 2.11: Output frequency of BBCO versus input control voltage when BBP is the opposite of BBN

The gain does not vary significantly when the number of stages is increased; whereas, when decreasing the supply voltage a large variation can be noticed, the gain at 0.5V is almost two and a half times the gain at 0.75V.

Since a large K_{VCO} is important for the stability of PLLs, as will become evident in Chapter 3; also given the significant increase on the gain when the supply voltage is 0.5V, while the figure of merit is only deteriorated by a couple of decibels, the oscillator that will be used for the PLL designed in the next chapter will be composed of three stages and biased with 0.5V with one frequency divider.

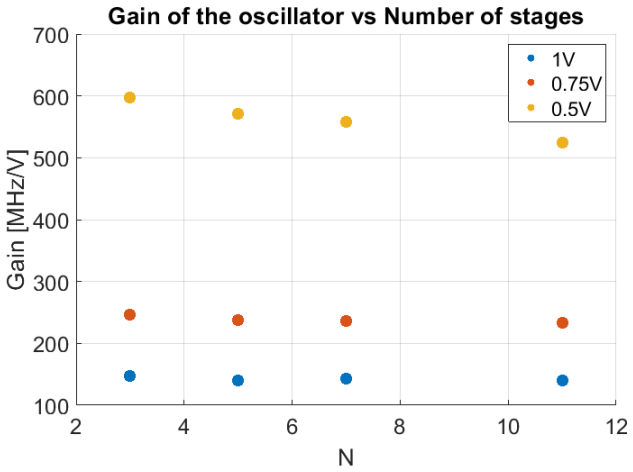

Figure 2.12: Oscillator gain versus number of stages and supply voltage

The phase noise spectrum for this oscillator is shown in Figure 2.13. The two phase noise region can be distinguished from each other by the difference in slope. The flicker noise has a slope of -30dB per decade while the thermal noise is -20dB/decade. The corner frequency is located approximately at 1MHz.

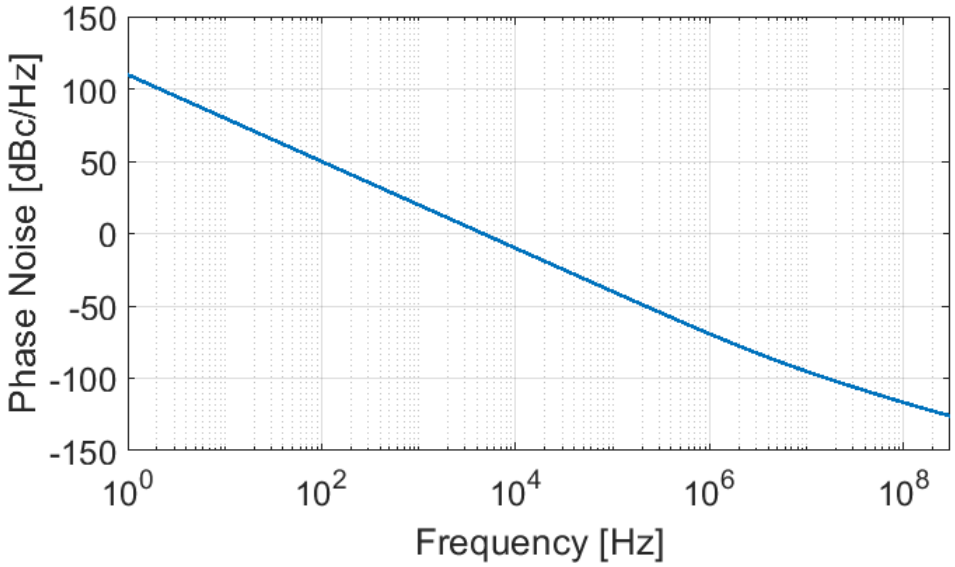

Figure 2.13: Phase Noise for the selected oscillator

2.4 Alternative implementations

The fact that negative voltages have to be generated to bias the back of the pMOS LVT transistors brings higher complexity and power consumption to the frequency synthesizer, as it will be explained in the next chapter, due to the additional block needed, i.e., the BBP driver.

One way to overcome this issue is to simply ground the back of the pMOS transistors. The disadvantage of this method is the reduced gain. Indeed, as shown in Figure 2.14, the gain decreases approximately by a factor of three compared to the oscillator selected in the previous section, i.e., from 600MHz/V to 184MHz/V. The figure also shows the impact of the process corners on the BBCO response, the worst case scenarios are in the crossed corners. When the nMOS is slower, the BBN must be increased almost to the maximum value to reach the target frequency; whereas, when the pMOS is slower, the BBN has to be greatly decreased. Both these effects occur so the nMOS compensates for the effect of the process variation. If temperature variations are included to these scenarios, the target frequency can no longer be reached.

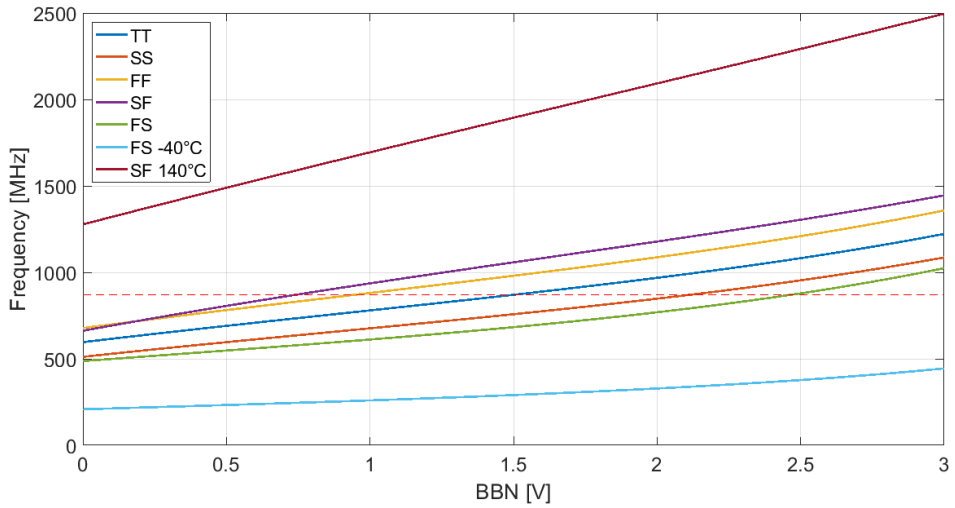

Figure 2.14: Output frequency of BBCO versus input control voltage with grounded BBP

Another way of avoiding the use of negative voltage for the oscillator is by using regular V_t transistors (RVT) thanks to their Reverse Back Bias (RBB) capabilities, whose specifications are shown in Figure 2.15. Unlike LVT transistors, the back of the pMOS transistors is the one that can be biased with positive voltage.

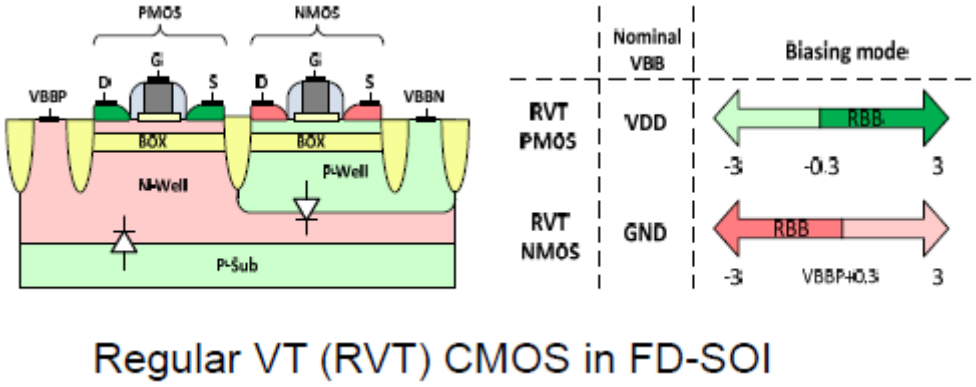

Figure 2.15: Cross-section of an RVT CMOS in FD-SOI and the Reverse Body Bias ranges (from [25])

If the BBCO is implemented by using RVT transistors for both nMOS and pMOS, keeping this time BBN at zero voltage, the response of the BBCO is the one in Figure 2.16, shown for all the process corners. A first difference with respect to the response for the BBCO implemented with LVT transistors is the slope of

the curve. When increasing the value of BBP, the V_t of the pMOS transistors is decreased instead of being increased, as in the previous case. As a consequence, the current, and thus the frequency, decreases. This BBCO can still be used for the PLL with minor alterations to the architecture of the latter. Another difference is the fact that the oscillator gain was only affected by a reduction lower than 10%, resulting in 557MHz/V.

The slow and fast corners are the worst-case scenarios for this case, for which the temperature variations were included. The fast corner at high temperature fails at reaching the target frequency.

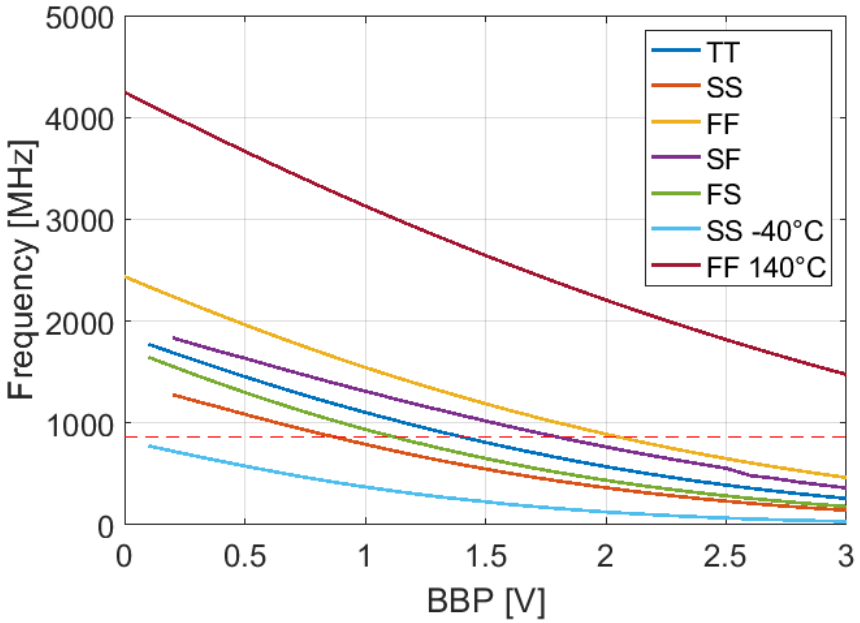

Figure 2.16: Output frequency of BBCO, implemented with RVT transistors, versus input control voltage with grounded BBN

Finally, a last alternative implementation of the BBCO can be done with a combination of both types of V_t transistors. If the RVT pMOS transistors are used to take advantage of the higher gain observed with respect to BBP in the previous case and the LVT nMOS transistors are used to take advantage of their speed and remove its dependence on the output frequency, a better performing oscillator should be obtained. Indeed that is the result obtained in the simulations, shown in Figure 2.17, as the gain increased slightly compared to the original BBCO, 610MHz/V that represents an increase of 3%. The same impact as in the previous case of the process and temperature variations can also be observed in the figure.

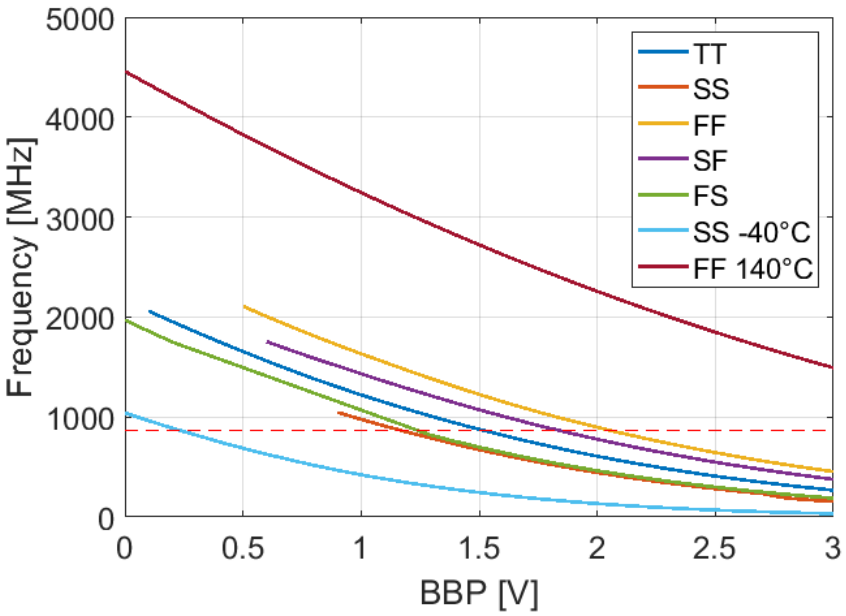

Figure 2.17: Output frequency of BBCO, implemented with RVT pMOS transistors and LVT nMOS transistors, versus input control voltage with grounded BBN

2.4.1 Power and phase noise comparison

Table 2.1 shows the power and phase noise performance comparison between all the different BBCO implementations. The original implementation has the best phase noise performance but the highest power consumption; whereas the alternative implementations have all very similar power and also phase noise performance, differing from one another less than 2dB. It is important to add to the comparison the fact that the supply voltage of the implementations using RVT transistors had to be increased from 0.5V to 0.55V in order to reach the target frequency with the length range of the technology. While this benefits the phase noise, it impacts the power of the oscillator. Including this factor on the comparison, the implementation using LVT transistors and grounded BBP is the best among the alternative ones. This is confirmed using the figure of merit, in the flicker-noise region this implementation has a lower value for the FoM compared to the other alternative implementations, while in the thermal-noise region the values are very similar. The disadvantage of this implementation is that it has the lowest gain. There is thus a trade-off between the phase noise performance and the frequency range of the BBCO. However, the gain on frequency range is much more significant than the deterioration of the phase noise. Therefore, the

implementation with the highest gain is preferred.

However, the original implementation is the one used in the designed PLL. The alternative implementations introduced in this section should be considered for achieving a further reduction on the power consumed in the frequency synthesizer, as will be thoroughly discussed in Chapter 3.

Transistor type	Power	Phase Noise @ 1MHz	Phase Noise @ 100MHz	FoM @ 1MHz	FoM @ 100MHz	Gain
LVT (BBP=-BBN)	2.23μW	-69.58dBc/Hz	-116.7dBc/Hz	-154.87dBc/Hz	-161.99 dBc/Hz	600MHz/V
LVT (BBP=0V)	1.66μW	-65.53dBc/Hz	-113.2dBc/Hz	-152.1dBc/Hz	-159.77dBc/Hz	184MHz/V
RVT (BBN=0V)	1.99μW	-64.73dBc/Hz	-114.1dBc/Hz	-150.51dBc/Hz	-159.88dbC/Hz	557MHz/V
RVT + LVT (BBN=0V)	1.87μW	-64.22dBc/Hz	-113.5dBc/Hz	-150.27dBc/Hz	-159.55dBc/Hz	610MHz/V

Table 2.1: Power and phase noise performance comparison of the BBCO implementations

Chapter 3

Design of a ULP PLL based on a BBCO

A Phase-Locked Loop (PLL) is a feedback system that synchronizes the BBCO oscillations to those of an input reference. Figure 3.1 shows a schematic representation of the building blocks of this system, whose key blocks are the oscillator designed in Chapter 2 and the Phase Frequency Detector (PFD).

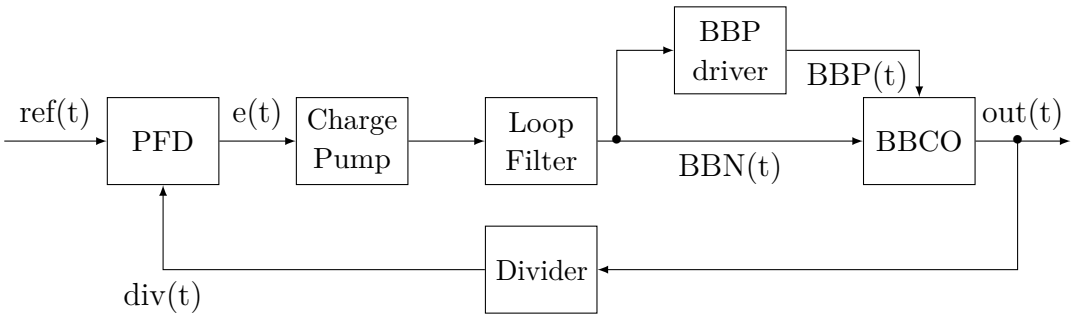

Figure 3.1: PLL block diagram

In this chapter, a detailed design of this system will be presented. The design is based on the phase-domain model of the system and its performance, in terms of the phase noise coming from the oscillator and its non-idealities, will be compared to the specifications of the Sigfox standard selected in Chapter 1. Finally, the global power consumption as well as the total phase noise will be compared to state-of-the-art transceivers.

3.1 Phase-domain model

The PLL is a feedback system whose state variable is the phase of the signal. This is the actual characteristic the PFD measures, as will be explained next. In order to design and analyse the PLL, the transfer function $\frac{\Phi_{out}(s)}{\Phi_{in}(s)}$ must be determined [29], for which a model of each of the building blocks has to be built.

3.1.1 PFD

The PFD is a circuit capable of sensing both phase and frequency differences between two input signals, in this case between the reference and divided clocks.

A typical implementation is shown in Figure 3.2(a), consisting in two resettable flip-flops and an AND gate. It has three states, both outputs at zero or Qa (resp. Qb) at one when a rising edge of clock A (resp. clock B) comes first. The width of the pulses created at the outputs provide information on the phase and frequency difference between the two clocks. The signals Qa and Qb are more commonly named UP and DOWN since that is their effect on the BBCO frequency, to increase or decrease it according to the phase difference between the reference and divided clocks.

Figure 3.2(b) shows the input-output characteristic of the PFD, where the output is defined as the difference between the average values of Qa and Qb when the frequency of both input clocks is the same [29]. The key attribute of this phase error characteristic is that it is asymmetric about zero phase, which enables the frequency detection [8]. The model of this circuit is a subtraction with a gain of $\frac{1}{2\pi}$, which is the slope of the PFD characteristic.

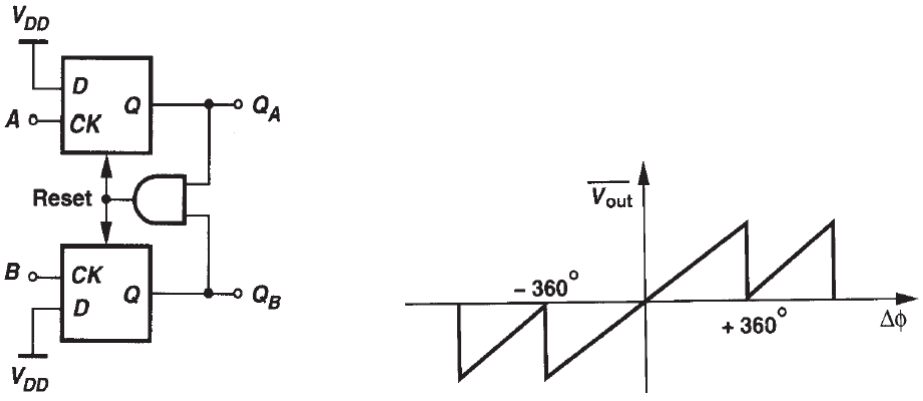

(a) (b)

Figure 3.2: PFD (a) implementation and (b) input-output characteristics (from [29])

3.1.2 BBCO

The BBCO transforms the input voltage into a frequency by Equation 2.2. Since the phase is the integral of the frequency, the model of the oscillator is an integrator with the gain described in Section 2.3.4, which translates into $\frac{K_{BBCO}}{s}$.

3.1.3 Charge-pump

The most common way to sense the difference between the average values of the UP and DOWN signals is to use a charge pump [29]. A circuit made of two switched current sources that pumps charge into or out of the loop filter according to the logic signals coming from the PSD. The model for this circuit is a gain equal to the current of the sources, I_{cp}.

Figure 3.3 shows a typical implementation of a charge pump in CMOS technology, where transistors Q1 and Q3 act as switches controlled by the UP and DOWN signals and transistor Q2 (resp. Q4) is the source for the down (resp. up) current.

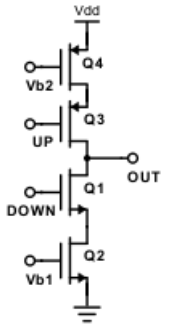

Figure 3.3: Implementation of the charge-pump

3.1.4 Loop filter

The loop filter has a low-pass behaviour to perform the averaging of the PFD pulses. The most common architecture used with the charge pump [8], [29], [30] is depicted in Figure 3.4. The main advantage of this architecture is that the additional integrator allows to achieve full range span of the oscillator [8]. The resistor is added to create a zero and enhance the phase margin of the system, while the capacitor C_2 is added to filter the large voltage jumps that occur when current is injected in the other branch [29].

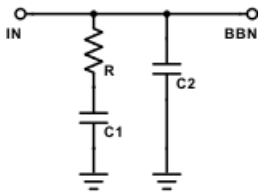

Figure 3.4: Implementation of the loop filter

The transfer function of this filter is

$$G_1(s) = \frac{1 + sRC_1}{s(C_1 + C_2) + s^2RC_1C_2(C_1 + C_2)}$$

but if the capacitance of C_2 is chosen five to ten times lower to that of C_1, it can be neglected [29] and thus the transfer function is simplified to

$$G_1(s) = \frac{1 + sRC_1}{sC_1}$$

which will be the one considered from now on.

3.1.5 BBP driver

The output of the loop filter is a voltage that is used to bias the n-wells of the BBCO, i.e. the BBN. As discussed in Section 2.2, the bias of the p-wells (BBP) has to be the opposite of that of the n-wells. This implies that a block needs to be added between the filter and the oscillator in order to generate this voltage, the BBP driver.

Ideally, the input-output relation of this block is $BBP = -BBN$, but in reality the circuit that performs this function behaves like a low-pass filter because it cannot follow the fast changes of the input voltage. Therefore, the transfer function of this block is

$$G_2(s) = \frac{-1}{1 + \frac{s}{\omega_c}} \tag{3.1}$$

where ω_c is the cut-off frequency of the circuit. The choice and impact of this cut-off frequency will be discussed in Section 3.3.4.

An implementation of this block can be done by the use of continuous-time current-matching loops [2], as the one shown in Figure 3.5. The idea of such a circuit is to balance the current, between the nMOS and the pMOS, of a replica of an inverter of the BBCO in order to have the correct back bias voltage on the pMOS, the opposite of the BBN, and input the value to the oscillator.

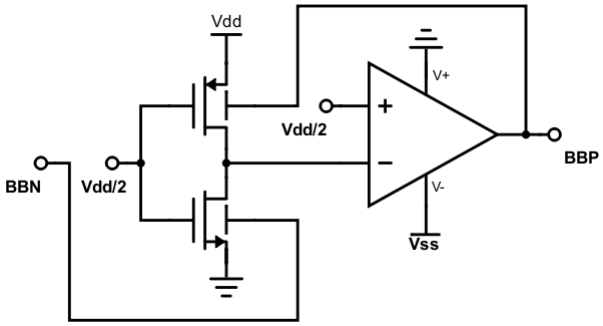

Figure 3.5: BBP driver schematic using a continuous-time current-matching loop

3.1.6 Divider

The frequency divider used in the PLL behaves as the ones described in the previous chapter, it divides the input signal frequency by an integer number, N. Thus, its model in the frequency domain is simply $\frac{1}{N}$.

Figure 3.6 shows the implementation of the frequency divider for a divide value that is not a power of two, and cannot be achieved by cascading flip-flops but by a counter. In this case, that the frequency must be reduced from 868MHz to 250kHz, the frequency of the reference clock used in [11], the divide value must be 3472.

In order to reduce power consumption, the counter needs to be operated at a frequency lower than the output frequency. Therefore, a maximum of divide-by-two frequency dividers has to be used before the counter. The maximum power of two that is also a divisor of 3472 is 16, which is why four of those devices where used, three before the counter and one after the comparator. The output of the comparator resets the counter and toggles the last frequency divider when the counter reaches $\frac{N}{16}$, i.e., 217. The 8-bit counter was implemented using a register bank and a ripple carry adder and the 8-bit comparator was implemented using standard AND and NOR cells.

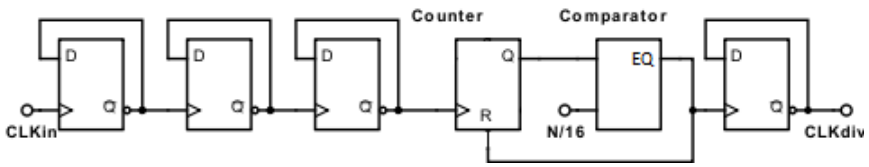

Figure 3.6: Frequency divider schematic for higher divide values

3.1.7 System response

Putting all these individual models together results in the model shown in Figure 3.7.

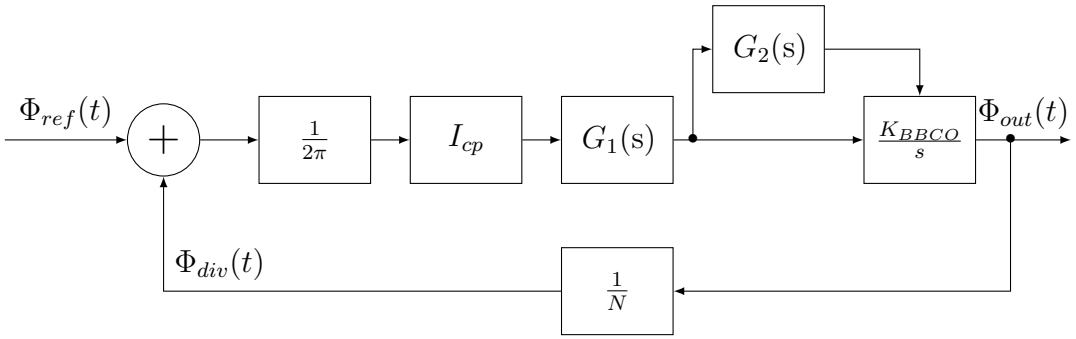

Figure 3.7: PLL linearized model diagram

The resulting closed-loop transfer function is given by the following equation

$$H(s) = \frac{N\omega_n^2(1 + sRC_1)}{s^2 + 2\zeta\omega_n s + \omega_n^2} \tag{3.2}$$

where the parameters ω_n and ζ are given by Equation 3.3 and Equation 3.4 [30].

$$\omega_n = \sqrt{\frac{I_{cp}K_{BBCO}}{2\pi C_1 N}} \tag{3.3}$$

$$\zeta = \frac{R}{2}\sqrt{\frac{I_{cp}C_1K_{BBCO}}{2\pi N}} \tag{3.4}$$

3.2 Design

Now that the transfer function is characterized, the ω_n and ζ parameters can be used to design the PLL, more specifically the charge pump and the loop filter.

The natural frequency ω_n must be below the reference clock's frequency, 250kHz, and above the channel bandwidth. In the case of Sigfox, this bandwidth is either 100Hz or 600Hz. For reasons that will be clarified later when the phase noise of the PLL is analysed, a natural frequency closer to the reference is preferred to decrease as much as possible the phase noise. The parameter ω_n is set to 1×10^5rad/s.

In the case of the attenuation ζ, in order to get a good trade-off between low settling time and low ringing in the transient response, this parameter is usually set to 0.71 [29].

Knowing these values, that K_{BBCO} is equal to 597MHz/V, N is equal to 3472 and that C_2 is five times lower than C_1, only three parameters are left to be determined using Equation 3.3 and Equation 3.4. To be able to solve them, the parameter I_{cp} was set to 2.5μA, a low value to have a low power consumption in the charge pump. This results in C_1 equal to 10pF, C_2 equal to 2pF and R equal to 1.7MΩ.

Figure 3.8 shows the transient response of the frequency at the output of the PLL obtained implementing its behavioral model on Verilog-AMS, based on its phase-domain model. The settling time, i.e., the time needed for the PLL to lock at the target frequency of 868MHz, is approximately 80μs. A negligible value when compared to the symbol time of Sigfox (10ms) and very close to the settling time of the state-of-the-art PLLs described in Table 1.2.

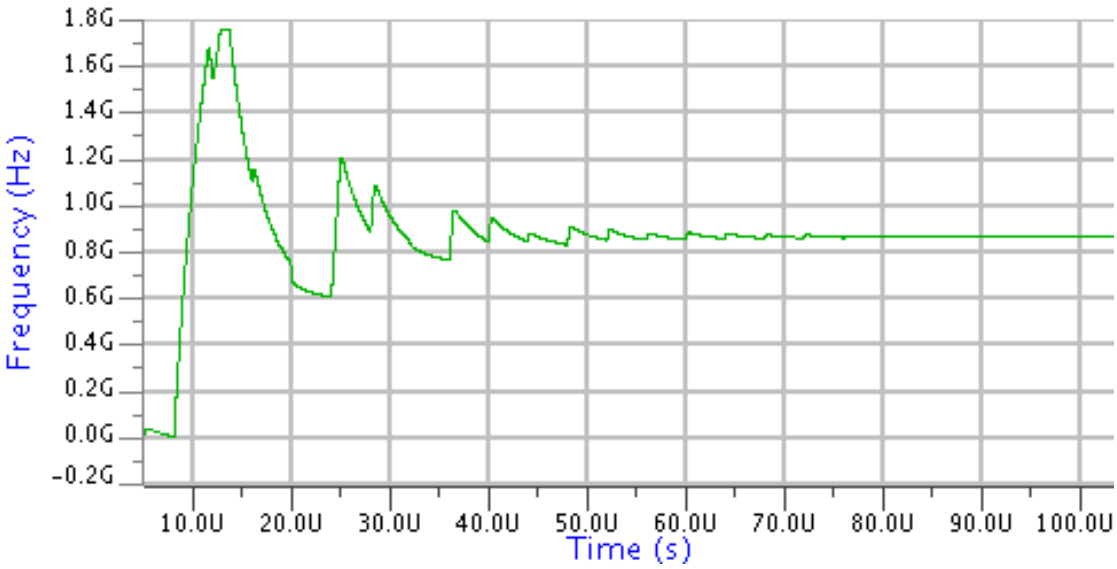

Figure 3.8: Transient behaviour of the output clock's frequency

The bode plot of the transfer function $H(s)$ described in Equation 3.2 is shown in Figure 3.9, where the low-pass nature of the loop can be noticed.

3.3 Non-idealities

Even if the PLL is in lock condition, imperfections in the PFD and charge pump lead to ripple in the control signal of the oscillator, producing a signal that is not longer periodic [29].

These non-idealities are mainly the leakage currents and the current mismatch in the charge pump, the non-ideal behaviour of the BBP driver and three different

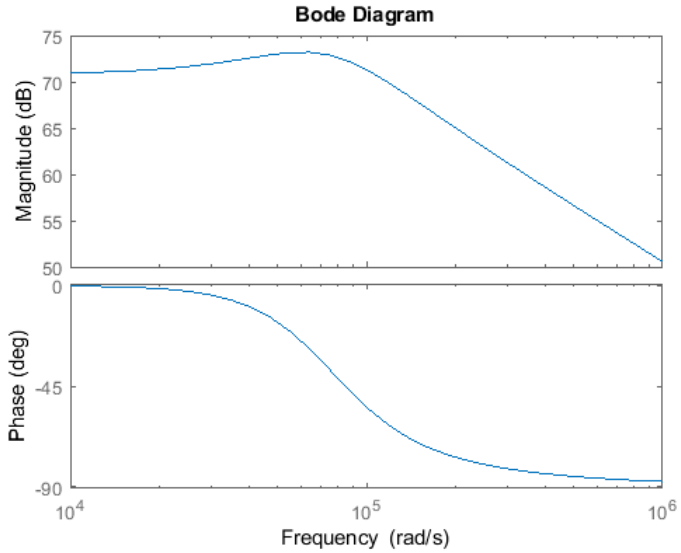

Figure 3.9: Closed-loop transfer function of the PLL

delays: the clock-to-Q delay, the reset delay and the skew between the UP and DOWN signals.

To measure the individual impact of these non-idealities on the PLL performance, the period jitter of the output clock was measured. This jitter is the variation on the clock's period from cycle to cycle. A histogram was made with the period of each cycle, from which the mean and standard deviation were extracted.

The period jitter can be converted to phase noise by means of the equation [31]

$$\mathcal{L}(\Delta f) = 10log\left(\frac{J^2 f_0^3}{\Delta f^2}\right) \tag{3.5}$$

where J is the period jitter, f_0 is the carrier frequency and Δf is the offset from the carrier frequency. This offset frequency is equal to the reference clock's frequency because this is the frequency at which the blocks that introduce the non-idealities work.

3.3.1 Leakage

Due to the leakage current in the transistors of the charge pump, a small amount of current is injected to the loop filter even when both switches are off. This causes a phase error, introduced by the loop to compensate for this extra current.

Figure 3.10 shows the mean and standard deviation of the clock's period versus the leakage current of the charge pump, normalized to the nominal current of the charge pump, 2.5μA.

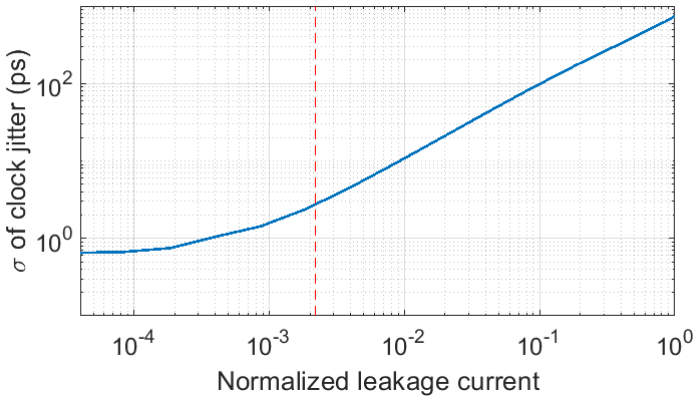

Figure 3.10: Jitter performance versus the charge pump leakage

An increase in the period jitter can be noticed as the leakage current gets closer to the nominal value. The mean value of the clock period stays always at a close value to the target value of 1.1521ns.

A typical value of this leakage, shown in red in the figure, was measured on a charge-pump implementation using I/O LVT transistors of the 28nm FD-SOI CMOS technology, shown in Figure 3.11, supplied at 3V to cover the whole BBN range of the BBCO and biasing the BBN at 1.5V while the BBP at 0V to reduce the parasitic capacitance seen at the output of the BBP driver, an important measure for the power consumption of this block, as will be explained later.

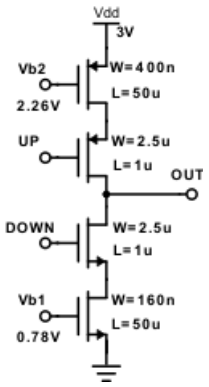

Figure 3.11: Implementation of the charge pump, including transistor sizes and bias values

The leakage obtained for this implementation is 5.5nA, resulting in a normalized current of 2.2×10^{-3}. The period jitter at this value is approximately 2.4ps, which converted to phase noise, using Equation 3.5, gives -72.2dBc/Hz at an offset of 250kHz.

3.3.2 Current mismatch

If the down current is not exactly equal to the up current in the charge pump due to, for example, PVT variations, the PLL introduces a phase error in order to compensate for it and inject zero net charge at each cycle. This causes a periodic ripple in the control voltage of the oscillator. The impact of this non-ideality is studied in conjunction with the non-zero delays explained next because their effect is dependent on each other, as will become clear soon.

3.3.3 Delays

There are three different delays that occur in the PLL that affect its performance. The first two involve the PFD implementation: the non-zero clock-to-Q delay of its flip-flops and the reset delay, i.e., the time between when both UP and DOWN signals are up and when they are both reset. The latter is represented by T_P in Figure 3.12.

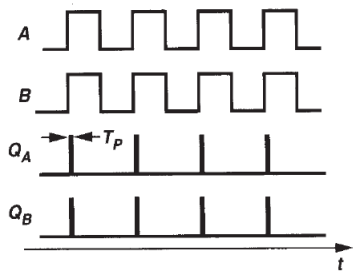

Figure 3.12: Timing diagram depicting the reset delay in the PFD (from [29])

The last delay comes from the charge pump implementation. Since the up switch is made with a pMOS transistor, which is active low, but the UP signal is active high, an inverter must be inserted between them to have a functional system. This additional gate introduces a skew between the UP and DOWN signals, causing, as in the previous case, a periodic ripple in the control voltage of the oscillator.

According to [29], these three delays can be related. The skew between the PFD signals (T_D) is equal to the delay of a single gate. The clock-to-Q and reset delays, depending on the flip-flop implementation, can be approximated to the delay of two and ten gates, i.e., two and ten times T_D, respectively.

Figure 3.13 shows the evolution of the period jitter versus the value of the skew between the UP and DOWN signals for different values of the current mismatch at the charge pump. The first thing to highlight is that the delays have almost no impact on the PLL performance if the up and down currents are perfectly matched.

Also, no matter the level of mismatch, if the delays are small enough the impact is also limited.

It is when both non-idealities combine that an increase is observed in the period jitter, the highest mismatch having the highest increase rate. The mean of the clock period, as in the previous case, stays at a close value from the target period.

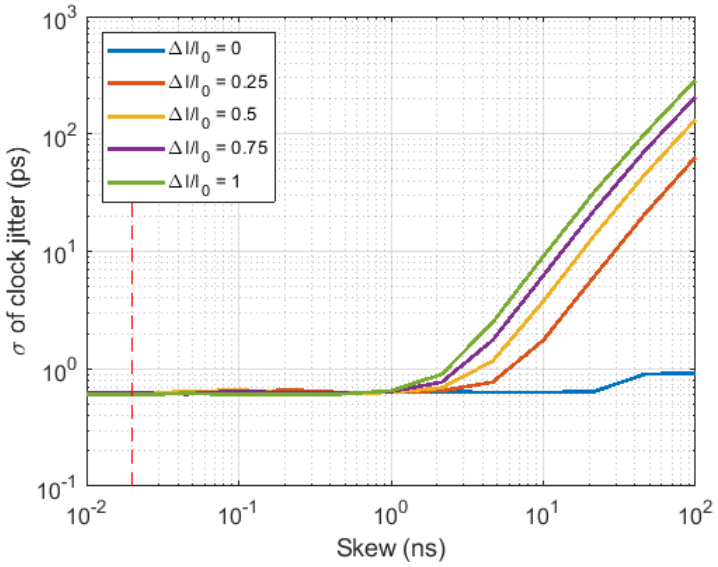

Figure 3.13: Jitter performance versus the charge pump current mismatch and the delays

The typical gate delay was determined by measuring the reset delay of a PFD implemented with 8-track post-layout standard cells of the 28nm LVT FD-SOI CMOS technology supplied at 0.5V, 1.5V of BBN and 0V of BBP. The measured delay was 200ps, meaning that the typical gate delay in this technology is approximately 20ps. This typical delay is shown in red in the figure. The period jitter at this value is approximately 0.6ps for any current mismatch, which converted to phase noise using Equation 3.5, gives -84.2dBc/Hz at an offset of 250kHz.

3.3.4 BBP driver

As already described with Equation 3.1, the BBP driver has a low-pass behaviour. Figure 3.14 shows how the cutoff frequency impacts the jitter performance of the PLL. The period jitter increases as the cut-off frequency gets closer to the cut-off frequency of the PLL. For lower values, below 10kHz, the locking of the PLL is affected.

The BBP driver designed in [2] using a continuous-time current matching block has a cut-off frequency of 1MHz, value shown in red in the figure. The period jitter at this value is approximately 1.4ps, which converted to phase noise using Equation 3.5, gives -77.1dBc/Hz at an offset of 250kHz.

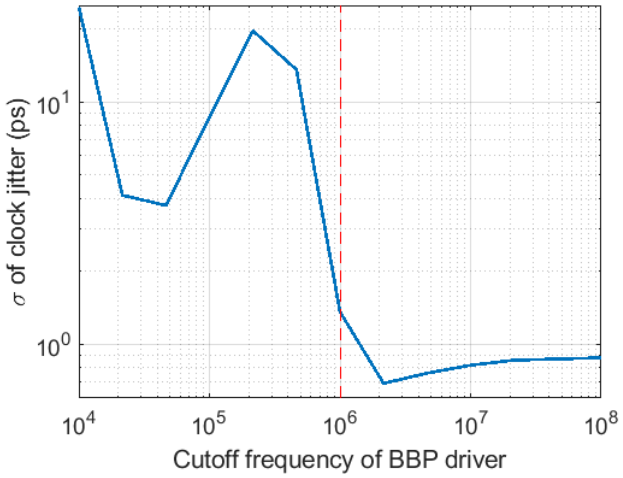

Figure 3.14: Jitter performance versus the BBP driver cutoff frequency

3.4 Power

To evaluate the total power consumption of the PLL, the PFD and the frequency divider were implemented using 8-track post-layout standard cells of the 28nm LVT FD-SOI CMOS technology, supplied at 0.5V, 1.5V of BBN and 0V of BBP. The charge pump was implemented using I/O LVT transistors and the BBP driver was implemented as in [2], its design trade-offs will be explain next. The results are summarized in Table 3.1.

Block	Power	Percentage of total power
PFD	233.9nW	0.13%
BBCO	3.16μW	1.69%
Charge pump	101.31μW	54.32%
Divider	6.4μW	3.43%
BBP driver	75.4μW	40.43%
Total	186.5μW	100%

Table 3.1: Power consumption of the PLL decomposed in the different building blocks

It is important to notice the big impact the BBP driver and the charge pump have on the total power of the PLL. This is because they need to be supplied with 3V in order to generate the complete range of the back bias. This is why the alternative implementations of BBCOs, presented in Section 2.4, are important. Also, a BBCO with a reduced back bias range, e.g., 1.8V instead of 3V to match the nominal supply voltage of I/O transistors, should be considered to reduce the power consumption and simplify the design of these blocks at the expense of a reduced frequency range in the oscillator.

3.4.1 BBP driver

The BBP driver, as implemented in [2], consists in a replica of an inverter of the oscillator and an amplifier, as shown in Figure 3.5. The amplifier can be implemented with a simple architecture, as the one proposed in [32], with back-gate-driven nMOS input pair as proposed in [2] to allow voltage range compatibility, without the use of level-shifting, between the amplifier implemented with thick-oxide 1.8V I/O transistors and the 0.5V inverter implemented with thin-oxide core devices. Indeed, without this special architecture, the maximum of 1.8V that I/O transistors support between their electrodes could be exceeded due to the -3V supply needed in the amplifier to generate the complete back bias voltage range.

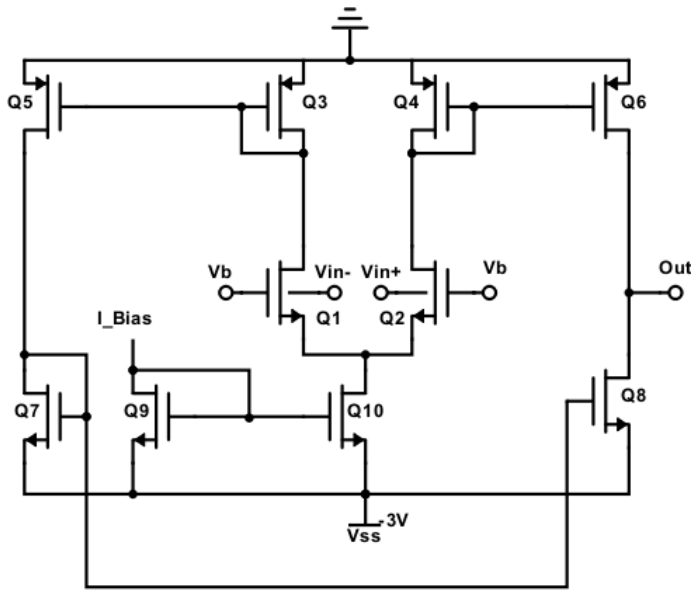

Figure 3.15: Schematic of operational amplifier

This simple architecture, shown in Figure 3.15, can be used to perform a quick approximation of the power needed in the amplifier to achieve a cut-off frequency

of 1MHz, in function of the total load capacitance. This capacitance corresponds to the total parasitic capacitance of the back gates of all the pMOS transistors biased with the voltage generated by the BBP driver. Since this value depends on the number of transistors connected to this node, three different values will be taken for the power study: 10fF, 100fF and 1pF. The higher value is the one used in Table 3.2.

The equation relating the power consumption of the amplifier architecture shown in Figure 3.15 and its cut-off frequency is the following

$$f_T = \left[\frac{g_{mb}}{I_D}\right]_1 \frac{I_{bias}}{4\pi C_L} \tag{3.6}$$

where f_T is the cut-off frequency, targeted at 1MHz, $\left[\frac{g_{mb}}{I_D}\right]_1$ is a property of the transistor Q1 that indicates its region of operation, C_L is the load capacitance and I_{bias} is the bias current.

From this equation, the bias current can be computed, the total current of the amplifier being twice this current. Taking the supply voltage as 3V and the $\frac{g_{mb}}{I_D}$ of the transistor Q1 as 1, keeping in mind that the transconductance of a back-gate-driven input pair, g_{mb}, is approximately ten times lower than that of the usual input pair, g_m, results in the power shown in Table 3.2 for the three values of the output capacitance. The total power is obtained by adding the power consumed by the current-matching inverter, 167.6nW.

C_L	Amplifier power	Total power
10fF	754nW	921.6nW
100fF	7.54μW	7.71μW
1pF	75.4μW	77.24μW

Table 3.2: Power consumption of the BBP driver block versus the load capacitance

It is therefore important, as another technique for reducing power consumption in the PLL, to limit as much as possible the load capacitance of the amplifier, i.e., to minimize the number of back gates supplied with negative voltages.

3.5 Phase Noise

The noise of all the building blocks of the PLL, as well as the reference jitter and the reference feedthrough due to the non-idealities contribute to the phase noise seen at its output.

To simplify the noise model, the usual approach is to refer all PLL noise sources, besides the BBCO, to the PFD output [8]. Like this, the noise only has two components: the PFD-referred noise and the BBCO-referred noise, each of which is filtered differently by the PLL loop.

The BBCO-referred noise is filtered by the PLL by the following transfer function

$$H_3(s) = \frac{s^2}{s^2 + 2\zeta\omega_n + \omega_n^2} \tag{3.7}$$

where ω_n and ζ are the parameters given by Equation 3.3 and Equation 3.4. This means that the phase noise of the BBCO is high-pass filtered, the flicker-noise region being attenuated while the thermal-noise region is left unchanged. The bode diagram of this filter, obtained with the parameters used in the design, is shown in Figure 3.16.

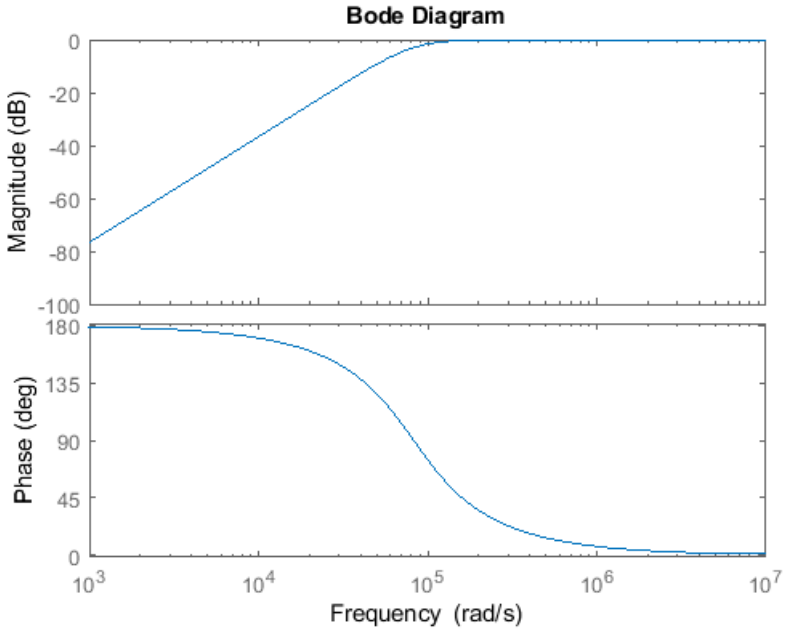

Figure 3.16: Bode diagram of the transfer function of the filter for the BBCO-referred noise

The result of filtering the BBCO phase noise, shown in Figure 2.13, using this filter is that the slope of the flicker noise, up to ω_n, is now of 10dB/decade, while the rest of the flicker noise and the thermal noise is left more or less unchanged. The results are shown in Figure 3.17 where they are compared to the transmission spectral mask shown Figure 1.11, already converted to a power spectrum by dividing

by the resolution. It can be noticed that the phase noise coming from the non-idealities discussed in Section 3.3 is below the mask at the frequency offset of 250kHz but the BBCO-referred noise violates the mask.

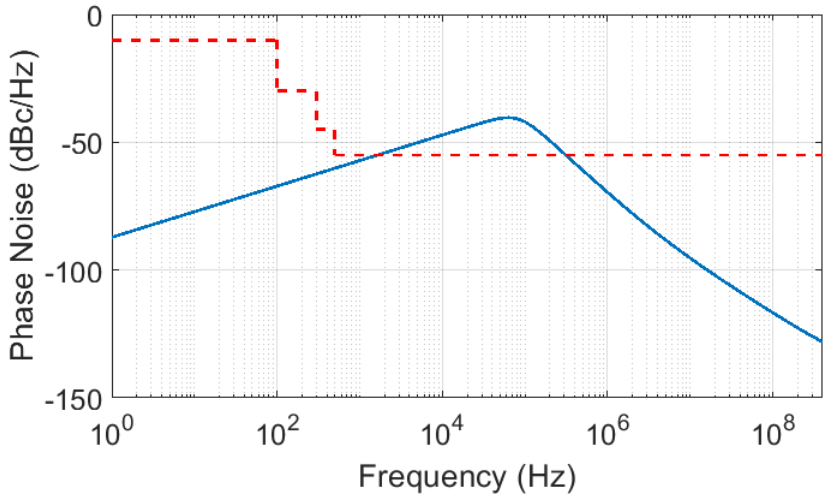

Figure 3.17: BBCO-referred phase noise compared to the Sigfox spectral mask

The PFD-referred noise is filtered by the following transfer function

$$H_2(s) = \frac{2\pi N\omega_n^2(1+sRC_1)}{s^2+2\zeta\omega_n+\omega_n^2} = 2\pi H(s) \tag{3.8}$$

which is simply the transfer function of the PLL scaled by a factor 2π, it thus has a low-pass behaviour. However, obtaining the PFD-referred noise is not as straightforward as the BBCO-referred noise, for which a special simulation is available in the circuit simulator Eldo.

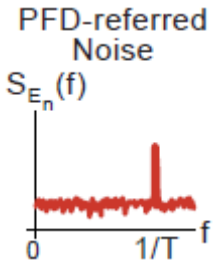

Figure 3.18: Typical plot of PFD-referred noise (from [8])

The PFD-referred noise has the typical form shown in Figure 3.18, where the peak at the frequency $\frac{1}{T}$, the frequency of the reference clock, corresponds to the

non-idealities discussed before. One way to obtain an approximate of this power spectrum, ignoring for the moment the contributions from the divider and the reference clock, is to simulate the PFD and charge pump together, using level-shifters between these two blocks due to their different supply voltages (0.5V for the PFD and 3V for the charge pump), and measure the output current. The approximate PFD-referred noise would then be the power spectrum of the output current divided by the square of the nominal current, 2.5μA. Unfortunately, due to lack of time, this simulation could not be performed to complete the phase noise study of the PLL. Nevertheless, a specification for its value will be given next.

In order to comply with the spectral mask, the BBCO-referred noise shown in Figure 3.17 must be reduced by a factor of 15dB. Since the phase noise is inversely proportional to the oscillator's power, this reduction can be obtained by increasing the power consumed by the BBCO by a factor of 30, which is possible if more inverters are used in parallel in the oscillator.

	This work	**Lachartre[11]**	**STmicro[17]**	**Semtech[21]**	**Semtech[18]**
Protocol	Sigfox	Sigfox	Sigfox	LoRa	LoRa
Technology	28nm FD-SOI	65 nm	-	-	-
Datarate	100bps	100bps	300bps	1.2kbps	1.2kbps
Phase Noise @1MHz	-70dBc/Hz	-106.42dBc/Hz	-117dBc/Hz	-120dBc/Hz	-122dBc/Hz
Synthesizer Power	278μW	8.65mW	1.26mW	14.85mW	19.14mW
Type of PLL	Integer-N	Integer-N	Fractional-N	Fractional-N	Fractional-N
Settling time	80μs	-	75μs	60μs	60μs

Table 3.3: State-of-the-art of LPWAN transceivers compared to the designed PLL

Assuming the PFD-referred noise, after filtering, is below -55dBc/Hz, the PLL total power could be increased to 278.14μW, an increment of 50%, to successfully fulfill the Sigfox specifications. Even after this increase, the power consumption is more than 70% lower than the power consumed by the frequency synthesizers of the state-of-the-art transceivers discussed in Table 1.2. More specifically, it is 96% lower than the power consumed by the synthesizer of the transceiver designed in [11], used as inspiration for this design. The downside is the poor phase noise obtained, which was expected as a ring oscillator has a higher phase noise than LC oscillators, the type of oscillator used in the state-of-the-art transceivers. All these results are summarized in Table 3.3, the synthesizer power considered for the state-of-the-art transceivers corresponds to their TX-ready mode of operation, in which other blocks of the transceiver, on top of their PLL, could impact this power; however, it is assumed that this impact is limited. Therefore, the advantage of a BBCO-based PLL is the possibility to go to lower power in spite of an increase in the phase noise, which fortunately still fulfills the specifications.

Conclusion

To conclude this work, a critical look will be taken at the different steps we carried out, bringing also some future work perspectives.

The main question of this work was to analyze the possible advantages a BBCO-based PLL could bring to long-range communications compared to other PLL types, such as current-starved ring oscillator or LC-VCO. An important ingredient for an objective comparison between the different PLL types is to define the specifications required for their performance in objects connected through Low-Power Wide Area Networks. The relevant specifications for a PLL are related to its phase noise, in the form of an emission spectral mask for the transmitter and of blocker-tolerance performance for the receiver.

As the specifications are directly related to the LPWAN standard, we first introduced in Chapter 1 the three main ones: Sigfox, LoRa and NB-IoT, along with their specific strategies to reduce the power consumption and to reach a long range in the communication. NB-IoT, as an optimized version of LTE technology for low power, proved to be the most energy efficient standard of the three in the study of commercial modules. However, Sigfox and LoRa present advantages in terms of lower complexity and cost for applications needing longer range with lower data rate requirements. The Sigfox standard was finally selected for our frequency synthesizer. The main reason for this choice was the high phase noise, at the same power consumption, expected from a ring oscillator compared to the LC VCOs used in state-of-the-art transceivers. Moreover, a particular transceiver stood out in the analysis of the state of the art. The design presented in [11] solves a major issue of Sigfox by eliminating the need of a TCXO, which used to increase the cost and power of transceivers. This design was used as inspiration for the frequency synthesizer.

The BBCO, central block of the PLL, was designed having two figures of merit in mind: power and phase noise, whose trade-off was the main guideline in the design methodology presented in Chapter 2. A ring oscillator with minimal number of stages and low supply voltage, 0.5V, running at twice the target frequency and followed by a divide-by-2 frequency divider was proposed as the best candidate for the PLL. Its low power and high gain, in spite of an increase in the phase noise,

were the main reasons for this selection.

Another question was raised regarding the BBCO implementation, more specifically, the type of transistors to use in order to reduce the power consumed by the PLL. Indeed, if the back gates of some transistors need to be biased at the opposite value of the back gates of the other ones, a special block must be added to the PLL. Due to the large voltage range of the back bias, this block must be supplied with 3V, resulting in a high power consumption. Therefore, the use of negative voltages in the BBCO should be avoided, if possible. It was found that a combination of RVT pMOS transistors and LVT nMOS transistors, keeping BBN at ground, increased the oscillator gain by 3% while it slightly decreased the consumed power by the oscillator and the phase noise performance. This implementation is thus preferable for BBCO-based PLLs instead of using a single transistor flavor, since it would result in a 40% power saving in the PLL. A possible enhancement of this circuit would be to reduce the range of the back bias to 1.8V because it would further reduce the power consumption as the supply voltage of the charge pump would be reduced. Its trade-off against the oscillator frequency range would have to be studied.

Finally, the design methodology for the PLL was presented in Chapter 3, with the figure of merits this time being the stability and settling time, as well as the power and phase noise. For the latter, a method to approximate it at the output of the PLL using the system's transfer functions was detailed and used to evaluate the final result. When comparing the designed frequency synthesizer against state-of-the art transceivers, a similar settling time was obtained and can be neglected when compared to the symbol period. Also, a large power saving was achieved at the price of a higher phase noise, although it still meets the Sigfox specifications of emission spectral mask. Therefore, the advantage found for BBCO-based PLLs compared to the other types is the possibility to go to lower power consumption. A pending task on the design is to extract the total phase noise at the output, either by approximating the PFD-referred noise or by extracting it from a whole-system simulation, and meet the transmission spectral mask with the same power budget. Moreover, the addition of a BBCO without negative bias should be done to further reduce the power consumption of the frequency synthesizer.

Bibliography

[1] D. Bol and G. de Streel, "An 802.15.4 IR-UWB transmitter SoC with adaptive-FBB-based channel selection and programmable pulse shape," 2019.

[2] G. de Streel, F. Stas, T. Gurne, F. Durant, C. Frenkel, A. Cathelin, and D. Bol, "SleepTalker: A ULV 802.15.4a IR-UWB Transmitter SoC in 28-nm FDSOI Achieving 14 pJ/b at 27 Mb/s With Channel Selection Based on Adaptive FBB and Digitally Programmable Pulse Shaping," *IEEE Journal of Solid-State Circuits*, vol. 52, no. 4, pp. 1163–1177, Apr. 2017.

[3] Y.-H. Liu, C. Bachmann, X. Wang, Y. Zhang, A. Ba, B. Busze, M. Ding, P. Harpe, G.-J. van Schaik, G. Selimis, H. Giesen, J. Gloudemans, A. Sbai, L. Huang, H. Kato, G. Dolmans, K. Philips, and H. de Groot, "13.2 A 3.7mw-RX 4.4mw-TX fully integrated Bluetooth Low-Energy/IEEE802.15.4/proprietary SoC with an ADPLL-based fast frequency offset compensation in 40nm CMOS," in *2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers*. San Francisco, CA, USA: IEEE, Feb. 2015, pp. 1–3. [Online]. Available: http://ieeexplore.ieee.org/document/7063013/

[4] K. Mekki, E. Bajic, F. Chaxel, and F. Meyer, "A comparative study of LPWAN technologies for large-scale IoT deployment," *ICT Express*, vol. 5, no. 1, pp. 1–7, Mar. 2019.

[5] T. Almholt, "Emerging standards for low power IoT from BLE to Wide Area Networking solutions," in *2017 IEEE International Solid-State Circuits Conference (ISSCC)*. San Francisco, CA, USA: IEEE, Feb. 2017.

[6] J.-P. Bardyn, T. Melly, O. Seller, and N. Sornin, "IoT: The era of LPWAN is starting now," in *ESSCIRC Conference 2016: 42nd European Solid-State Circuits Conference*. Lausanne, Switzerland: IEEE, Sep. 2016, pp. 25–30.

[7] A. Hajimiri, "A General Theory of Phase Noise in Electrical Oscillators," *IEEE Journal of Solid-State Circuits*, vol. 33, no. 2, 1998.

[8] M. H. Perrott, "Integer-N Frequency Synthesizers," Aug. 2008. [Online]. Available: http://www.cppsim.com/pll_lectures.html

[9] R. Navid, T. Lee, and R. Dutton, "Minimum achievable phase noise of RC oscillators," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 3, pp. 630–637, Mar. 2005.

[10] C. Bernier, "Ultra-low-power wireless connectivity," in *2018 IEEE International Solid- State Circuits Conference - (ISSCC)*, 2018.

[11] D. Lachartre, F. Dehmas, C. Bernier, C. Fourtet, L. Ouvry, F. Lepin, E. Mercier, S. Hamard, L. Zirphile, S. Thuries, and F. Chaix, "7.5 A TCXO-less 100hz-minimum-bandwidth transceiver for ultra-narrow-band sub-GHz IoT cellular networks," in *2017 IEEE International Solid-State Circuits Conference (ISSCC)*. IEEE, Feb. 2017, pp. 134–135.

[12] Rohde&Schwarz, "NB-IoT/eMTC power saving features: eDRX vs. PSM - Solution." [Online]. Available: https://www.rohde-schwarz.com/us/solutions/test-and-measurement/wireless-communication/iot-m2m/iot-m2m-webinars-videos/nb-iot-emtc-power-saving-features-edrx-vs.-psm_234108.html#media-gallery-4

[13] J. Gozdecki, A. Jajszczyk, and R. Stankiewicz, "Quality of service terminology in IP networks," *IEEE Communications Magazine*, vol. 41, no. 3, pp. 153–159, Mar. 2003.

[14] H. Naumann, "NB-IoT versus SIGFOX, LoRaWAN, and Weightless – power / energy the inconvenient truth," Feb. 2018. [Online]. Available: http://www.gsm-modem.de/M2M/iot-university/nb-iot-power-consumption/

[15] Wisol, "WISOL / WSSFM10r1at," Wisol, WSSFM10R1AT datasheet, Mar. 2017.

[16] TDnext, "High-performance, low-current US Sigfox Gateway," TDnext, TD1508 datasheet, Jul. 2016.

[17] STmicroelectronics, "Ultra-low power, high performance, sub-1 GHz transceiver," STmicroelectronics, S2-LP datasheet, May 2018.

[18] Acsip, "Product specification," Acsip, S76S datasheet, Mar. 2017.

[19] Quectel, "BC95 Hardware Design," Quectel, BC95 datasheet, Jun. 2017.

[20] Semtech, "LoRa Calculator," 2015.

[21] ——, "SX1272/73 - 860 MHz to 1020 MHz Low Power Long Range Transceiver," Semtech, SX1272 datasheet, Jan. 2019.

[22] Sigfox, "Sigfox connected objects: Radio specifications," Sigfox SA, Tech. Rep., Feb. 2019.

[23] A. Kay, "Introduction to Noise Measurement," in *Operational Amplifier Noise*. Elsevier, 2012, pp. 63–87.

[24] S. Suman, K. G. Sharma, and P. K. Ghosh, "Analysis and design of current starved ring VCO," in *2016 International Conference on Electrical, Electronics, and Optimization Techniques (ICEEOT)*. Chennai, India: IEEE, Mar. 2016, pp. 3222–3227.

[25] P. Cathelin and A. Cathelin, "FD-SOI Technology Advantages for Analog/RF & Mixed-Signal Designs," STMicroelectronics, Crolles, France, Oct. 2016.

[26] T.-H. Lee and P. A. Abshire, "Frequency-Boost Jitter Reduction for Voltage-Controlled Ring Oscillators," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 24, no. 10, pp. 3156–3168, Oct. 2016.

[27] B. Razavi, "TSPC Logic [A Circuit for All Seasons]," *IEEE Solid-State Circuits Magazine*, vol. 8, no. 4, pp. 10–13, 2016.

[28] N. Da Dalt and A. Sheikholeslami, *Understanding Jitter and Phase Noise*. Cambridge, United Kingdom: Cambridge University Press, 2018.

[29] B. Razavi, "Phase-Locked Loops," in *Design of Analog CMOS Integrated Circuits*. New York, NY: McGraw Hill, 2001, pp. 532–578.

[30] R. J. Baker, "Digital Phase-Locked Loops," in *CMOS: Circuit Design, Layout, and Simulation*, 3rd ed. Hoboken, NJ: John Wiley & Sons, Inc., Aug. 1997, pp. 551–612.

[31] K. Kundert, "Predicting the Phase Noise and Jitter of PLL-Based Frequency Synthesizers," Mar. 2012.

[32] J.-P. Eggermont, D. De Ceuster, D. Flandre, B. Gentinne, P. Jespers, and J.-P. Colinge, "Design of SOI CMOS operational amplifiers for applications up to 300°C," *IEEE Journal of Solid-State Circuits*, vol. 31, no. 2, pp. 179–186, Feb. 1996.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
École polytechnique de Louvain
Rue Archimède, 1 bte L6.11.01, 1348 Louvain-la-Neuve, Belgique | **www.uclouvain.be/epl**