Leveraging a Digital Chirp Spread Spectrum Detector for LPWAN Wake-Up Receivers

Pol Maistriaux, Marco Gonzalez, Jérôme Louveaux, David bol
ICTEAM Institute, *Université catholique de Louvain*, Louvain-la-Neuve, Belgium
{pol.maistriaux, marco.gonzalez, jerome.louveaux, david.bol}@uclouvain.be

***Abstract*—Wireless sensor networks are emerging as serious candidates to detect critical events in ecosystems and remote areas. To maximize their coverage in those environments, mesh networks can be used. However, they require more complex medium access control, at the cost of higher power consumption for the sensor-node RF receiver. The alternative paradigm of wake-up receivers (WuRXs) offers a possible solution to this issue. The objective is to offer an ultra-low-power detection mode to wake up the main receiver, placed in deep sleep, only when an incoming RF signal is found. Previous WuRXs designed for short-range communications fail to reach the sensitivity level required for longer distances, partly due to their simple digital baseband (DBB) processor. In this work, we investigate the use of chirp spread spectrum (CSS), already well-known thanks to LoRa, as an efficient modulation scheme for long-range WuRXs, leveraging its high processing gain and acceptable complexity. To validate this approach, we propose a digital hardware implementation and provide synthesis results in 65-nm CMOS. Compared to state-of-the-art long-range WuRXs, the proposed DBB improves the sensitivity by 11.8 dB with a limited 23-% power consumption overhead.**

***Index Terms*—Wake-up Receiver, LPWAN, CSS, DBB, DSP**

I. Introduction

Natural ecosystem degradation is one of the challenges of our modern societies, climate change being both a cause and a consequence. Indeed, ecosystems act both as carbon sinks and as biodiversity reservoirs. They can be protected by continuously monitoring their state with environmental data [1] but also by trying to detect critical events as early as possible to quickly react to various threats, e.g., wildfires or landslides. The Internet of Things (IoT) has introduced wireless sensor networks (WSNs) as a potential solution for the early detection of such events [2]. As shown in Fig. 1(a), WSNs rely on the deployment of a large number of low-cost sensor nodes in the monitored area [3]. Due to the potential long distances between the sensor nodes and the need for long battery lifetime, low-power wide-area networks (LPWANs) are of particular interest [4]. This is especially true when using a mesh topology which offers a theoretically unlimited coverage as long as each node has a neighbour able to relay its packets. As they must either be synchronized or always-on to receive the packets to be relayed, there is considerable pressure on the reception mode of the nodes' transceivers [5]. Thus, the reception mode becomes the bottleneck of the system in terms of power consumption, which is exacerbated by the limited data traffic as critical events are rare. This currently limits the use of WSNs as early warning systems in remote areas.

This work was supported by the Fonds de la Recherche Scientifique – FNRS under FRIA Grant n°40021516.

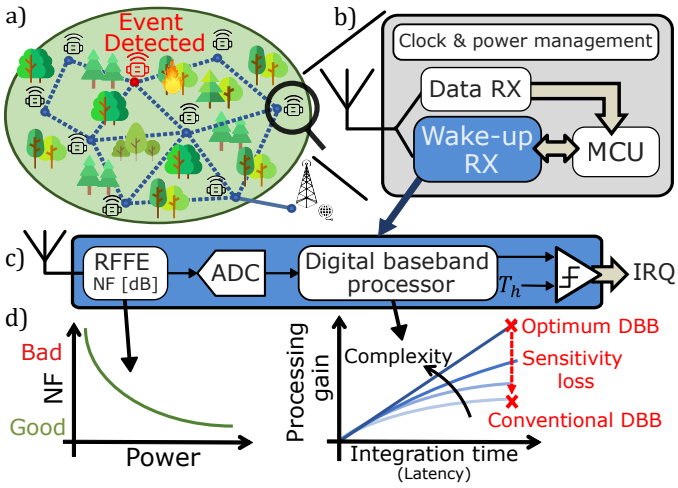


Fig. 1. (a) Wireless sensor network for critical-event detection. (b) Architecture of the sensor node and (c) of the wake-up receiver. (d) Trade-offs of its analog and digital blocks.

In the meantime, a new paradigm of communication has emerged with wake-up receivers (WuRXs) providing an ultra-low-power (ULP) reception mode for the transceivers [6]. This is enabled by an additional receiver as depicted in Fig. 1(b), in parallel with the main receiver, with the only objective of detecting the presence of an incoming wake-up signal (WuS). The main receiver and the rest of the node can be powered off or placed in deep sleep, effectively saving energy, while only the WuRX is active. Upon detection of a WuS, it wakes up the microcontroller unit (MCU) and the main radio so that the incoming packet can be decoded. Several papers have already highlighted the significant energy-saving potential of WuRXs [7] but these savings are application-specific [8]. Current WuRXs, with their general architecture depicted in Fig. 1(c) and presented in Section II-A, fail to enable long-range communication with an acceptable power consumption and latency due to the inherent trade-offs between these performance metrics [6]. For the RF front ends (RFFEs), their fundamental noise–power trade-off, depicted in Fig. 1(d), limits the achievable range of WuRXs because of their constrained power budget, as explained in [9]. The other block of WuRXs, the digital baseband (DBB) processor, can be used to trade latency for range, providing a certain processing gain

by integrating the signal over time. It has been far less studied in the literature and is usually kept as simple as possible to avoid power consumption overheads. The simplicity of these DBB processors leads to a performance loss compared to the optimal one, as explained in Section II-B. As represented in Fig. 1(d), this loss increases for longer integration time, which accentuates the difficulty of communicating at long distances.

This work highlights the limitation of the DBB processors used in current WuRXs and proves that more advanced ones can provide much higher processing gain while only slightly increasing the power consumption. This is even more promising with CMOS technology scaling achieved in integrated circuits, which greatly improves the energy efficiency of digital blocks unlike for analog ones. To prove this, we leverage the moderate complexity and the near-optimal processing gain of the chirp spread spectrum (CSS) modulation, used in LoRa, with Section III providing a detailed analysis of its benefits on the communication range. Preamble detection, through which LoRa receivers detect incoming packets, has already been studied in the literature [10], but not as a ULP feature for WuRXs. As most of the digital blocks could be shared with a LoRa receiver, WuRX capabilities could be added to this LPWAN technology at a limited cost in terms of silicon area. To validate this approach, we then propose a CSS-DBB hardware implementation in Section III-B, with synthesis results evaluating both its power consumption and area provided in Section IV-A. Based on a state-of-the-art design from [11], we evaluate in Section IV-B the potential performance of a WuRX employing the proposed CSS-DBB processor. Finally, Section V highlights several future perspectives to improve and make the most of the proposed approach.

II. State of the Art

A. Overview of WuRX Architectures

WuRXs monitor the channel to detect the presence of an incoming WuS. They are mainly evaluated based on their power consumption P, their latency L, i.e., the time between the arrival of the signal and its detection, and their sensitivity S expressed in dBm [9]. The latter is the minimum input signal power required to ensure a reliable detection of the WuS. A robust communication can be specified as having a missed-detection ratio (MDR) of 0.1 % for a false-alarm rate (FAR) of less than one per hour. The sensitivity can be calculated with the following expression:

$$\begin{aligned} S\ [dBm] = &-174 + 10\log(B) + NF(P_{RFFE}) \\ &+ SNR_{\text{min}}(p_{FA}, p_{MD}), \end{aligned} \tag{1}$$

where B is the bandwidth, NF is the noise figure of the RFFE, i.e., the degradation of the signal-to-noise ratio (SNR) caused by the analog circuits, which is a function of its power consumption P_{RFFE}. SNR_{min} is the minimum SNR required at the DBB processor input to achieve a given MDR expressed as p_{MD} and a probability of false alarm p_{FA}. The latter is derived from the FAR with the following expression, where f_d is the rate of detection attempts:

$$p_{FA} = \frac{FAR}{3600 f_d}. \tag{2}$$

WuRXs were first developed for short-range communications, using envelope-detection (ED) schemes to achieve sub-µW power consumption, offering sensitivities limited to –70 dBm [12]. To reach such low power consumption, mostly passive circuits are used, with limited gain and efficiency, resulting in poor NFs [13]. Different RFFE architectures have been studied to improve the NF, e.g., efficient downconversion using mixers or the addition of low-noise amplifiers, but all these solutions imply an increased power consumption [9], [11]. This non-linear relation between NF and power consumption is expressed in (1) through the term $NF(P_{RFFE})$ and is qualitatively shown in Fig. 1(d).

Most WuRXs rely on absolute values of the received signal strength indicator (RSSI) to make at least a partial detection decision at the output of the RFFE, using a comparator with a fixed threshold [6] [14] [15]. Intuitively, RSSI-based decisions are inefficient in the case of negative SNRs [11]. With the signal buried in the noise, the RSSI is almost unchanged whether or not a signal is present. This makes conventional WuRX architectures unsuitable for long-range communications, the key characteristic of LPWAN technologies.

B. WuRX DBB Processors in Negative SNR Conditions

In standard detection schemes, e.g. the ones used in GPS technologies, spread-spectrum (SS) techniques are often employed to improve the system's resistance to noise and interference, enabling reliable communications even below the noise floor [16]. With SS, the transmitter sends a determined code spread out over a certain bandwidth. At the receiver, the signal is despread, i.e., correlated with the used code. Longer codes result in longer correlations and therefore higher latencies. The result can then be compared to a threshold T_h, to assess the presence of a signal. The correlation between the signal x, affected by an additive white gaussian noise (AWGN) w, and the reference code r is expressed as follows

$$y[k] = \sum_{n=0}^{N-1} (x[n] + w[n]) r[n-k]. \tag{3}$$

When the code generated at the receiver and the transmitted one are aligned, i.e., k is zero, the correlation integrates coherently the signal and randomly the noise. As already expressed in [16], this results in an improvement in SNR, also called processing gain, by a factor of N. Increasing N would therefore scale linearly both the latency and the sensitivity. Power-of-two values, called spreading factors (SF), are often used for N with $N = 2^{SF}$. Increasing the SF by one doubles N and the WuS length, potentially improving the sensitivity by 3 dB. The code employed can be of various kinds. To provide addressing capabilities, binary codes representing the identifier of the sender or the receiver have often been used.

Performing the complete correlation is a costly digital signal processing (DSP) operation. In (3), one correlation required

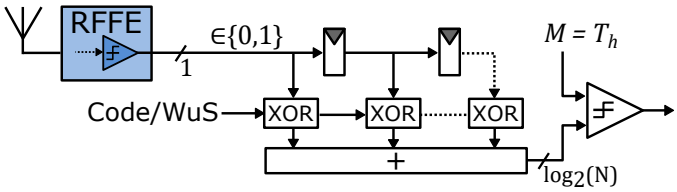


Fig. 2. M-out-of-N detection scheme in WuRx from [9], [11], [13], based on a 1-bit correlator.

for a single detection trial implies N multiplications, and must be repeated for each new sample. N useless trials are thus performed before the transmitted and generated WuSs are aligned. Overall, this results in a computational complexity of $O(N^2)$. Moreover, multiplications are resource-intensive DSP operations, both in terms of power and area. The conventional DBB processors of current WuRXs therefore simplify the correlation by performing the processing on a single bit. The costly multi-bit multiplications are reduced to trivial exclusive-or (XOR) operations, and memory elements to 1-bit registers [13], i.e., flip-flops. This implies quantizing the RFFE output to a single bit with a comparator replacing the analog-to-digital converter (ADC). For example, binary frequency-shift keying (BFSK) relies on the transmission of zeros and ones at two distinct frequencies. Bit decision is thus made at the RFFE output by extracting the RSSI at these two frequencies with EDs. They are then compared and the bit value associated with the highest one is selected. It should be noted that the two compared RSSI values are affected by two different noise realizations, resulting in twice the amount of noise. BFSK is thus inherently less efficient than other modulations.

The simplified 1-bit correlation scheme is shown in Fig. 2. It is equivalent to continuously trying to demodulate a hypothetical input signal and comparing the resulting N-bit sequence to the code. This correlator can be referred to as an M-out-of-N (MoN) detector since M bits out of N must match the WuS to declare a detection. The simplicity of this implementation leads to a loss of processing gain. Since MoN analyses the number of bits successfully identified in a sequence of length N, the correlation result follows a binomial distribution. We can thus express p_{FA} and p_{MD} as,

$$p_{FA,MoN} = 1 - \sum_{k=0}^{M-1} \binom{N}{k} (0.5)^N \tag{4}$$

$$p_{MD,MoN} = \sum_{k=0}^{M-1} \binom{N}{k} p_b^{N-k} (1-p_b)^k, \tag{5}$$

where M is equal to the detection threshold T_h, and p_b is the bit error rate (BER) of the underlying modulation scheme, here BFSK, with respect to the SNR. Indeed, the signal is continuously demodulated with a bit decision, explaining the dependence on the BER. In the absence of signal, noise is demodulated as either a 0 or a 1 with the same 0.5 probability, appearing in (4). Simulations were used to calculate and verify these results for an SF of 12 (SF12), i.e., $N = 2^{12} = 4096$,

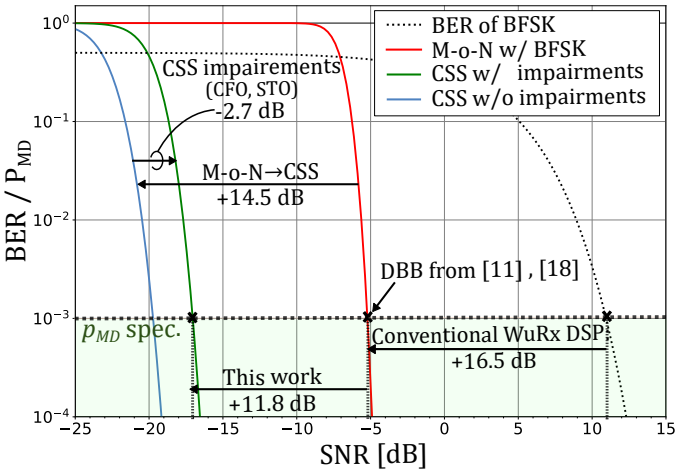


Fig. 3. Probability of missed detection with respect to the SNR for a false-alarm rate of one per hour and different detection schemes.

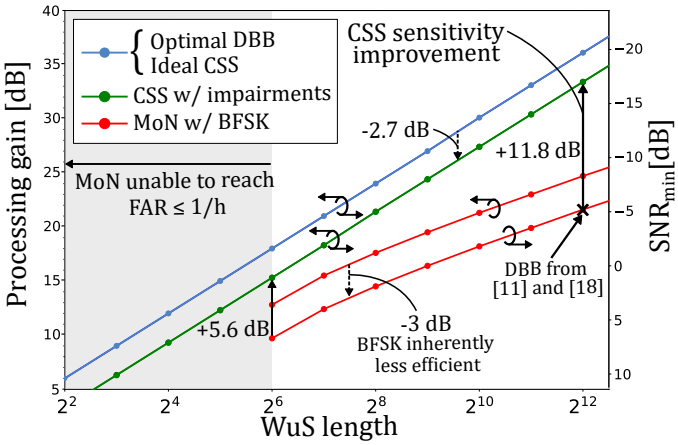


Fig. 4. Processing gain and SNR_{min} with respect to the WuS length for a false-alarm rate of one per hour and a missed-detection probability of 0.1 %.

which closely replicates the implementation results of [17] and [11]. In Fig. 3, p_{MD} results of MoN with respect to the SNR are shown for a FAR of 1/h. It can be seen that a minimum SNR of −5.2 dB is achieved. In Fig. 4, we show the processing gain with respect to the WuS length, i.e., the code length. First, we see that the optimal DBB processor implementing a correlation exhibits an ideal linear relationship between the processing gain and latency, whereas MoN becomes less efficient for long WuSs. Also, when comparing the SNR_{min} depicted on the right axis of Fig. 4, MoN suffers from a 3-dB loss due to BFSK being inherently less efficient than other modulations. Moreover, the MoN processing gain is not defined for small WuS lengths since no value of M allows to meet the FAR specification.

The complexity of (3) can be reduced while maintaining its processing gain leveraging that a time-domain correlation is equivalent to a multiplication in the frequency domain. Fast Fourier transforms (FFTs), with their complexity of $O(N \log N)$ can be used to efficiently convert to and from

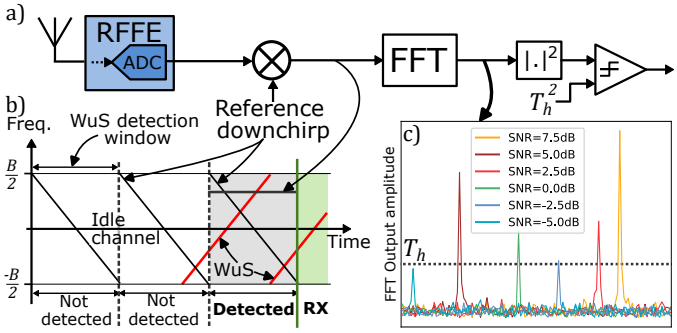


Fig. 5. (a) CSS detection scheme based on a fixed threshold. (b) Overview of signals at different point of the WuRX, up to the dechirped signal. (c) Output amplitude of the FFT for different SNR and with random STO.

the frequency domain. However, this approach requires at least two FFTs. CSS on the other hand, only requires a single FFT to correlate the signal.

III. CSS Detection

Let us now consider CSS, used by LoRa, as an efficient scheme to more effectively harness the processing gain of long WuSs while performing the correlation cyclically. A CSS transmitted signal, called chirp, is a linearly increasing frequency tone, spread over the whole considered bandwidth, here 125 kHz. The receiver scheme is shown in Fig. 5(a). The incoming signal is despread, also called dechirped, by multiplying it by the opposite chirp, i.e., with a decreasing frequency. This is shown in Fig. 5(b). The resulting signal is a single tone which can be integrated with a single FFT. Examples of the FFT output are shown in Fig. 5(c) for input chirps with different SNRs and random sample offsets between the transmitted WuS and the receiver. It is observed that this lack of synchronization only shifts the FFT maximum to a different bin. This means CSS inherently performs a cyclic correlation. Therefore, a cyclic shift of the reference chirp, used here as a WuS, by any number of samples does not affect the magnitude of the FFT outputs. Even though the CSS detection scheme works on a buffer of N samples to perform the cyclic correlation rather than continuously, repeating successively the WuS at the transmission ensures a complete WuS will be acquired. This repetition is often already required by modern WuRXs, which use duty-cycling to reduce their average power consumption while increasing their latency. Moreover, the cyclic correlation inherent to CSS only requires the acquisition of N samples with the RFFE. Its duty-cycling scheme is thus twice as efficient as the one of a linear correlator which must acquire $2N$ samples [11].

The false-alarm and missed-detection probabilities were already derived in the literature in the context of LoRa preamble detection [10]. In the absence of a CSS WuS, the FFT output follows a Rayleigh distribution and the probability of a false alarm is thereby expressed in (6). In the opposite case, it follows a Rician distribution as shown in (7), where Q_1 is the Marcum Q-function of order 1, T_h is the detection threshold and N_0 is the noise power at the DBB processor input.

$$p_{FA,CSS} = \exp\left(-\frac{T_h^2}{NN_0}\right) \tag{6}$$

$$p_{MD,CSS} = Q_1\left(\sqrt{2N \cdot SNR}, \frac{\sqrt{2}T_h}{\sqrt{NN_0}}\right) \tag{7}$$

A. CSS-based Detection Performance Results

We can now evaluate the performance of CSS and compare it to the simpler MoN scheme. Fig. 3 shows the major sensitivity improvement of CSS, with an initial SNR_{min} of -19.7 dB compared to -5.2 dB for MoN. This 14.5-dB gain corresponds to a range in free-space path loss (FSPL) which is five times longer. Two points must be noted. First, due to carrier frequency and symbol timing offsets (CFO and STO), the dechirped signal can be located in between two FFT bins, distributing its energy between them. When subject to the worst-case CFO and STO, a 2.7 dB degradation is observed, reducing the gain to 11.8 dB. In this work, the magnitude of this impairment has been observed empirically in simulations, but it should be proved mathematically in future works. Second, since the WuS is the reference chirp, common to all receivers, the addressing capabilities are lost, even though further work could exploit multiple CSS symbols to encode an address. As shown in Fig. 4, the processing gain of CSS is equivalent to the correlation one in the ideal case. Even with CFO and STO impairments, the linear relationship with N is preserved. Fig. 4 clearly highlights the increasing benefits of CSS compared to the conventional MoN.

B. CSS DBB Processor

To evaluate the usage of more efficient DBB processors in long-range WuRXs, we propose a digital hardware implementation in Verilog HDL of the CSS detection. The main focus of past research on CSS and LoRa was directed towards software implementation of efficient algorithms for packet detection and synchronization running on unconstrained hardware. In this paper, we focus on accelerating the part required for packet detection, which is also the bottleneck of the energy consumption of LoRa demodulation [18], i.e., the dechirping followed by the FFT. The proposed DBB processor architecture is represented in Fig. 6, with the dechirping and FFT stages performed sequentially. Since the dechirping is performed continuously while the FFT operates on a buffer of N samples, a ping-pong approach has therefore been used for the FFT, with two local memories (LMEMs). N complex samples, represented with I and Q components, come from the ADCs and are multiplied by the reference downchirp, stored in a read-only memory (ROM). The results are written in one of the two LMEMs. When completely stored, the FFT is performed while the new incoming dechirped samples are being written in the other LMEM.

The FFT is the most complex block of the DBB processor. Its efficiency relies on reusing common intermediate results of the discrete Fourier transform (DFT) when performing it on a power-of-two number of samples. The implementation

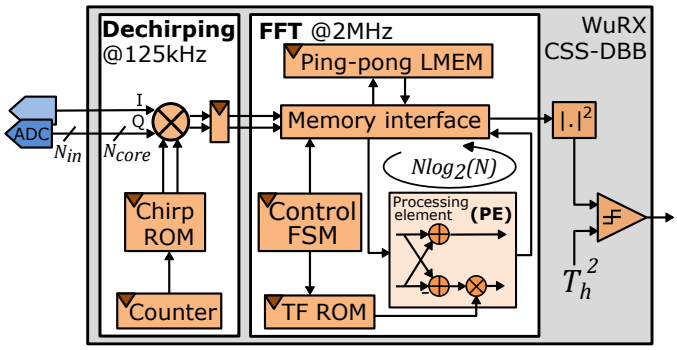


Fig. 6. Architecture of the CSS DBB processor. The main blocks, respectively the dechirping and the FFT are detailed.

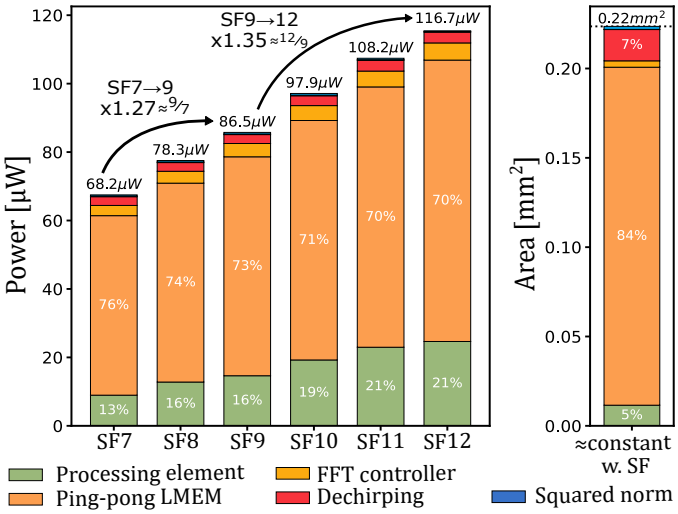


Fig. 7. Power and area of the DBB processor for different SFs. Synthesis was conducted in 65-nm CMOS technology with HVT transistors and a 1-V supply. Inputs are 12-bit IQ samples and 16 bits are used for the processing.

used here is a memory-based FFT using a single processing element (PE) which performs a butterfly operation on two samples, similarly to [19]. This operation consists in adding and subtracting the two inputs separately and rotating the complex output of the subtraction, which is achieved with a multiplication by a twiddle factor (TF). $N\log_2(N)/2$ butterfly operations are performed sequentially and $N\log_2(N)$ read and write operations are required. Dual-port SRAMs are used for the LMEMs, which allows the completion of the whole FFT in $N\log_2(N)$ cycles. The memory-based implementation has been chosen to limit the silicon area overhead since we do not require the throughput advantage of a pipeline implementation.

The dechirping stage requires one cycle per sample and can thus run at 125 kHz, while the FFT needs to run at least $\log_2(N)$ times faster. In the case of CSS symbols up to SF12, a clock running at 2 MHz was used to only require a simple clock divider to generate both clocks. An important point of the DBB processor is the number of bits used. Using more bits results in lower quantization noise, which is key when working with negative SNR values. LoRa receivers usually work with 12 bits [18]. Since the objective is to make parts of the DBB usable by the main receiver later on, we choose 12 bits. This moderate number of bits also results in a higher ADC power consumption than for the simplified 1-bit correlator. As fewer than 12 bits could be used for WuS detection, this DBB configuration represents a worst-case scenario. With SRAMs traditionally working with byte-size data, we choose to work internally with 16 bits for both I and Q samples, increasing the resolution in computations. As the FFT integrates and accumulates the dechirped frequency tone, overflow might still occur. To prevent this, a conditional scaling approach as introduced in [19] is employed by applying a bitshift of one to the PE inputs if one of its previous outputs was close to overflow.

IV. Implementation Results

A. Power and Area Characterization

The DBB processor was synthesized in TSMC 65-nm CMOS technology with high-threshold-voltage (HVT) transistors and a 1-V supply for the logic and 1.2 V for the SRAMs. The ability to switch between SFs is very important at the network level, with higher SFs resulting in a longer range and lower SFs resulting in shorter WuSs and thus lower transmission energy. Therefore, multiple spreading factors were characterized but with a constant LMEM size of 4096 complex samples, i.e., the memory space needed for the highest SF of 12. Fig. 7 shows synthesis results for the DBB processor, which consumes from 68.2 to 116.7 µW, corresponding to SF7 and 12. As the FFT is the most complex block, we expect the energy consumption to scale in $N\log_2 N$, both for the logic and for the memory. However, since the WuS duration increases linearly, the power consumption should only scale in $\log_2 N$. This is confirmed in Fig. 7 with the scaling between SFs being highlighted and compared to the theoretical trend. As expected, the FFT, and more specifically its memory accesses, dominates the power consumption with at least 70 %. For the area, the DBB processor occupies 0.22 mm^2, also dominated by the memory, which represents 84 %.

B. System-level Analysis

This section provides an estimation of the potential benefits of a CSS-based WuRX. As discussed in Section III-A, a CSS detection scheme allows a WuRX in SF12 to obtain a sensitivity gain of 11.8 dB compared to an MoN based on BFSK, for a fixed latency and considering potential impairments. To evaluate the power consumption overhead associated with the use of a CSS-DBB processor, we must add the contribution of the I and Q ADCs required. We can estimate it to 62.2 µW from [20], which implements a 12-bit ADC in 40 nm at 1 MS/s that could also be used in a LoRa receiver requiring oversampling. As we have compared the proposed DBB processor to the one in [11], we use their WuRX power consumption to compare our approach at the system level. They also work with a sub-GHz carrier frequency with I and Q paths. Their WuRX consumes 781 µW and with the addition of

62.2 and 116.7 µW for respectively the ADCs and the proposed DBB in SF12, this results in a total of 960 µW, leading to a 23-% overhead. To put these results into perspective, a figure of merit (FoM) used in the literature [11] linking sensitivity, power consumption and latency, is given by

$$\begin{aligned}\text{FoM } [dBm] = &- S[dBm] - 10\log(L[s]) \\ &- 10\log(P[W]).\end{aligned} \tag{8}$$

This metric can be used to compare various WuRXs. The addition of the proposed CSS-DBB processor would therefore improve the FoM by 10.9 dB compared to [11].

V. Observations and Future Perspectives

The proposed approach provides a significant FoM improvement when compared to the state of the art, which should be put into perspective. Without any countermeasure, the addressing capabilities of the WuRX are lost and synchronization is not done when a packet is detected, contrary to some other implementations. Addressing could be retrieved by using multiple CSS symbols with an ID, which could even be leveraged to further improve the sensitivity. Regarding synchronization, it is already well studied for LoRa, possibly performed in software as in [18].

Moreover, three other future perspectives should be highlighted regarding the use of CSS as a relevant detection scheme. First, with more advanced CMOS technology nodes becoming available for IoT devices, the power and area of the DBB processors will be reduced. Analog blocks of the RFFE will not scale as significantly as digital blocks. Going from 65 nm to 22 nm could lead to a power-area-product reduction by a factor of 8.5 for the DBB processor according to our estimation. This further highlights the relevance of transferring some of the WuRX's complexity to the DBB. Also, the sensitivity could be viewed as a constraint rather than an optimization metric. In this case, the improvement of the DBB processing gain could be used to relax the NF and therefore reduce the power consumption of the RFFE, which is the major contributor to the total power consumption in the evaluation from Section IV-B. Finally, the FFT used in the CSS-DBB implementation is the bottleneck of its power consumption. It could be optimized with more advanced FFT architectures, e.g., radix-2^2, which could reduce SRAM memory accesses and complex multiplications by a factor of two.

VI. Conclusion

In this work, we showed that long-range capabilities of wake-up receivers are currently limited by their DBB processor and that significant improvements can be made with acceptable power overheads by using more complex DBB processing. We proved this by leveraging a CSS modulation which can be efficiently processed at the receiver. Compared to literature with similar targets using a 4096-bit wake-up signal, CSS can achieve an improved sensitivity of 11.8 dB with an increase in power consumption from 781 to 960 µW. This is a 23-% power overhead which could be highly improved in future work and with technology scaling.

References

[1] G. M. Lovett, D. A. Burns, C. T. Driscoll, J. C. Jenkins, M. J. Mitchell, L. Rustad, J. B. Shanley, G. E. Likens, and R. Haeuber, "Who needs environmental monitoring?" *Front. Ecol. Environ.*, vol. 5, no. 5, pp. 253–260, 2007.

[2] A. Mohapatra and T. Trinh, "Early wildfire detection technologies in practice—a review," *Sustainability*, vol. 14, no. 19, 2022.

[3] D. Culler, D. Estrin, and M. Srivastava, "Guest Editors' Introduction: Overview of Sensor Networks," *Computer*, vol. 37, pp. 41–49, 2004.

[4] D. Ismail, M. Rahman, and A. Saifullah, "Low-power wide-area networks: opportunities, challenges, and directions," in *Proc. 19th Workshop Program of the Int. Conf. on Distributed Computing and Networking*. ACM, 2018, pp. 1–6.

[5] H.-C. Lee and K.-H. Ke, "Monitoring of Large-Area IoT Sensors Using a LoRa Wireless Mesh Network System: Design and Evaluation," *IEEE Trans. Instrum. Meas.*, vol. 67, pp. 2177–2187, 2018.

[6] I. Demirkol, C. Ersoy, and E. Onur, "Wake-up receivers for wireless sensor networks: benefits and challenges," *IEEE Wireless Communications*, vol. 16, pp. 88–96, 2009.

[7] J. Oller, I. Demirkol, J. Casademont, J. Paradells, G. U. Gamm, and L. Reindl, "Has Time Come to Switch From Duty-Cycled MAC Protocols to Wake-Up Radio for Wireless Sensor Networks?" *IEEE/ACM Trans. on Networking*, vol. 24, pp. 674–687, 2016.

[8] A. Kozłowski and J. Sosnowski, "Energy Efficiency Trade-Off Between Duty-Cycling and Wake-Up Radio Techniques in IoT Networks," *Wireless Personal Communications*, vol. 107, pp. 1951–1971, 2019.

[9] P. P. Mercier, B. H. Calhoun, P.-H. P. Wang, A. Dissanayake, L. Zhang, D. A. Hall, and S. M. Bowers, "Low-Power RF Wake-Up Receivers: Analysis, Tradeoffs, and Design," *IEEE Open Journal of the Solid-State Circuits Society*, vol. 2, pp. 144–164, 2022.

[10] J.-M. Kang, "LoRa Preamble Detection With Optimized Thresholds," *IEEE Internet of Things Journal*, vol. 10, pp. 6525–6526, 2023.

[11] K.-M. Kim, K.-S. Choi, H. Jung, B. Yun, J. Xu, J. Ko, and S.-G. Lee, "A -124-dBm Sensitivity Interference-Resilient Direct-Conversion Duty-Cycled Wake-Up Receiver Achieving 0.114 mW at 1.966-s Wake-Up Latency," *IEEE J. Solid-State Circuits*, pp. 1–14, 2022.

[12] P.-H. P. Wang, H. Jiang, L. Gao, P. Sen, Y.-H. Kim, G. M. Rebeiz, P. P. Mercier, and D. A. Hall, "A Near-Zero-Power Wake-Up Receiver Achieving 69-dBm Sensitivity," *IEEE J. Solid-State Circuits*, vol. 53, pp. 1640–1652, 2018.

[13] R. Piyare, A. L. Murphy, C. Kiraly, P. Tosato, and D. Brunelli, "Ultra Low Power Wake-Up Radios: A Hardware and Networking Survey," *IEEE Commun. Surv. Tutor.*, vol. 19, pp. 2117–2157, 2017.

[14] A. Dissanayake, J. Moody, H. L. Bishop, D. Truesdell, H. Muhlbauer, R. Lu, A. Gao, S. Gong, B. H. Calhoun, and S. M. Bowers, "A- 108dBm Sensitivity, -28dB SIR, 130nW to 41µW, Digitally Reconfigurable Bit-Level Duty-Cycled Wakeup and Data Receiver," in *2020 IEEE Custom Integrated Circuits Conf. (CICC)*, 2020, pp. 1–4.

[15] A. Dissanayake, H. L. Bishop, S. M. Bowers, and B. H. Calhoun, "A 2.4 GHz-91.5 dBm Sensitivity Within-Packet Duty-Cycled Wake-Up Receiver," *IEEE J. Solid-State Circuits*, vol. 57, pp. 917–931, 2022.

[16] A. Springer, W. Gugler, M. Huemer, L. Reindl, C. Ruppel, and R. Weigel, "Spread spectrum communications using chirp signals," in *IEEE/AFCEA EUROCOMM 2000. Information Systems for Enhanced Public Safety and Security*, 2000, pp. 166–170.

[17] K.-M. Kim, H.-G. Seok, O.-Y. Jung, K.-S. Choi, B. Yun, S. Kim, W. Oh, E.-R. Jeong, J. Ko, and S.-G. Lee, "A –123-dBm Sensitivity Split-Channel BFSK Reconfigurable Data/Wake-Up Receiver for Low-Power Wide-Area Networks," *IEEE J. Solid-State Circuits*, vol. 56, pp. 2656–2667, 2021.

[18] M. Xhonneux, J. Louveaux, and D. Bol, "Implementing a LoRa Software-Defined Radio on a General-Purpose ULP Microcontroller," in *2021 IEEE Workshop on Signal Processing Systems (SiPS)*, 2021, pp. 105–110.

[19] S. M. Noor, E. John, and M. Panday, "Design and implementation of an ultralow-energy fft asic for processing ecg in cardiac pacemakers," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 27, no. 4, pp. 983–987, 2019.

[20] J. Shen, A. Shikata, A. Liu, B. Chen, and F. Chalifoux, "A 12-Bit 31.1-W 1-MS/s SAR ADC With On-Chip Input-Signal-Independent Calibration Achieving 100.4-dB SFDR Using 256-fF Sampling Capacitance," *IEEE J. Solid-State Circuits*, vol. 54, pp. 937–947, 2019.