A Hybrid Approach to Enhance Cognitive Wireless Sensor Networks with Energy-Efficient Software-Defined Radio Capabilities

Ramiro Utrilla, Alba Rozas, Javier Blesa, Alvaro Araujo
Department of Electronic Engineering
Universidad Politécnica de Madrid
{rutrilla, albarc, jblesa, araujo}@b105.upm.es

Abstract

The loss of efficiency and dependability in wireless communications, caused by the increasing spectrum scarcity problem, is particularly critical for battery-powered devices such as sensor nodes. Cognitive Wireless Sensor Networks (CWSNs) arise in order to mitigate this situation by adding cognitive radio capabilities to them. However, the closed architecture and limited resources of traditional nodes represent a major constraint to perform certain required cognitive tasks. On the other hand, Software-Defined Radios (SDRs), which have the flexibility and performance needed to overcome the aforementioned limitations, currently have a power consumption too high for these networks. In this work, we propose a hybrid methodology of operation that consists of exploiting SDR technology only for those actions that strictly require its high flexibility, using traditional fixed hardware transceivers, which demand less node's resources, for the remaining tasks. Then, we present the architecture and the main electronic components of a platform able to operate according to this methodology. This solution constitutes a significant reduction in power consumption compared with existing low-power SDRs, while maintaining the functionality needed for research in CWSNs.

Categories and Subject Descriptors

B.m [**Hardware**]: Miscellaneous—*emerging tools and methodologies*

General Terms

Design, Performance

Keywords

Spectrum scarcity, cognitive wireless sensor networks, software-defined radio, low-power platform

International Conference on Embedded Wireless Systems and Networks (EWSN) 2017
20–22 February, Uppsala, Sweden

ISBN: 978-0-9949886-1-4

1 Introduction

The global adoption and use of wireless communications is growing rapidly in all areas of society and industry. According to forecasts [3], there will be over 11.6 billion of mobile-connected devices by 2020. As a result, global mobile data traffic is expected to grow to 30.6 exabytes per month by then, nearly an eightfold increase over 2015. Most of these devices are aimed to operate in the same license-free Industrial, Scientific and Medical (ISM) bands, causing these to reach a very high level of occupancy. This fact poses an important challenge to overcome: the spectral coexistence.

To address this challenge and achieve a reliable communication, the Cognitive Radio (CR) paradigm was defined [7]. Basically, it consists of a fully reconfigurable wireless transceiver with learning capabilities that is aware of its own resources, its surrounding environment and the user communication needs, and automatically adapts its internal parameters to the network and user demands.

Wireless Sensor Networks (WSNs) are one of the areas with the highest demand for cognitive capabilities because most of them operate in those overcrowded bands and they have to make a very efficient use of energy. For this last reason, traditional sensor nodes generally just consist of a Commercial Off-The-Shelf (COTS) radio chip and some sensors connected to a basic Microcontroller Unit (MCU). COTS radio chips have a high power efficiency and performance because they implement the lower layers of the communications stack as an Application-Specific Integrated Circuit (ASIC), only providing a packet-level interface with the MCU. Nevertheless, this closed architecture, along with the limited node's resources, impose significant constraints to perform some required CR tasks, such as: spectrum sensing or Physical (PHY) and Medium Access Control (MAC) layers adaptation, which have achieved relevant improvements in wireless communications efficiency [8].

In contrast, SDRs, which have the entire communication protocol stack implemented in software, reconfigurable hardware or a combination of both, have demonstrated an extraordinary flexibility and performance to overcome the aforementioned limitations [9], being the most appropriate technology for communication protocol researching, prototyping and verification. Nevertheless, most of them were designed from the CR perspective, i.e., without considering the unique challenges posed by the limited

resources of CWSNs. Therefore, most existing SDRs rely on chips with better performance than the typical ones in sensor nodes, so the consumption and price are far higher.

Thus, one of the major challenges in CWSNs [1, 2] is to find a compromise between a traditional sensor node architecture and a more flexible one, with enough computing performance to implement CR strategies, but maintaining a reasonable power consumption to guarantee the feasibility of operating in a realistic long-term testbeds.

In this work, we address this challenge in two ways. As a first contribution, we propose a hybrid methodology of operation that consists of exploiting SDR capabilities only for those actions that strictly require them, using traditional fixed hardware transceivers, that demand less node's resources, for the remaining tasks. As a second contribution, we present the architecture and the main electronic components of a platform able to operate according to this methodology.

The rest of the paper is organized as follows. In Section 2, some low-power SDR approaches are reviewed. In Section 3, the power consumption of these reference platforms is analyzed in detail. Based on this analysis, a new hybrid methodology to operate CWSN nodes along with the main aspects of their hardware design are presented in Sections 4 and 5 respectively. Finally, conclusions and future lines of work are offered in Section 6.

2 Related Work

During the last years some low-power SDR approaches oriented to battery-powered devices have appeared. The aim of these works is to promote experimental research of the PHY/MAC layers in low-power wireless communications.

Dutta et al. [4] propose to build a small, low-cost, and low-power SDR platform around the Microsemi's SmartFusion System-on-Chip (SoC) FPGA, which would act as the mixed-signal processing back-end. This device integrates a flash-based FPGA fabric, a 100 MHz ARM Cortex-M3 processor, and programmable analog circuitry. Thus, it gives more flexibility than traditional fixed-function MCUs without the excessive cost of soft processor cores on traditional FPGAs. Moreover, compared to SRAM-based FPGAs, flash-based FPGAs use non-volatile flash memory cells to load the active configuration that defines the function of the programmable logic elements and their connections. This sort of cells greatly reduces the static (leakage) consumption, which is crucial for low-power operation.

Along with the SmartFusion, they propose to use the MAX2831 highly-integrated transceiver as a 2.4 GHz front-end. This device allows direct access to the In-phase and Quadrature (I/Q) components of the baseband signals and it integrates almost all the circuitry required to implement the radio front-end functions into a single chip that draws very low power in sleep mode and can wake up rapidly when needed. This high level of integration and flexibility enables a considerable reduction in size, time to market and bill of materials, as most elements of an SDR are found in just two chips. Finally, these two components could be connected directly through the SmartFusion on-chip analog front-end for low baseband modulation rates, or by using a high-speed ADC/DAC stage for higher data rates.

Based on this previous work, Kuo et al. [5] implement the μSDR platform and demonstrate that their approach is the first one capable of running for a full day on a pack of AA batteries thanks to its sub-Watt power consumption in sleep mode. Moreover, they also achieve a significant reduction in cost and size compared with the existing SDR platforms.

Along the same lines, Szilvási et al. [11] develop another low-power SDR platform, called MarmotE SDR, based again on a Microsemi's SmartFusion flash-FPGA and a Maxim Integrated transceiver, this time the 2.4 GHz MAX2830. In their work, they evaluate the FPGA logic resource utilization and the platform power consumption. Its results show that the MarmotE SDR has sufficient computational resources to approach the WSN research from the PHY/MAC layers perspective and its power consumption allows a several hours long battery-based operation, which can be further extended by duty cycling.

However, it should be noted that when the flexibility of SDR operation is only required to perform some sporadic tasks, as is the case of CWSNs, these systems introduce penalties in several aspects, as explained throughout this work.

In the next section, we analyze these approaches in depth to subsequently propose a series of additional solutions aimed at achieving a greater consumption reduction.

3 Power Consumption Analysis

As shown in Section 2, there are already a couple of low-power SDR approaches that have achieved a substantial reduction of the energy consumption compared to previously existing SDR systems. In fact, both are very similar, relying on the SmartFusion SoC as a mixed-signal processing back-end and MAX283X transceivers as 2.4 GHz front-ends.

Tests conducted on these platforms constitute a good proof of concept of the range of possibilities that this technology gives. However, their power consumption is still too high to permit CWSNs long-term deployments, and it is also higher than the one of WSN solutions based on COTS radio chips by orders of magnitude. Moreover, the appearance of new IC solutions make existing approaches obsolete. For example, both platforms are based on the original SmartFusion, which does not support any ultra-low power mode while its new version, SmartFusion2, does. Therefore, these platforms have been studied in detail to propose a new methodology of operation, and a series of significant hardware design improvements aimed at achieving a greater consumption reduction.

Although some enhancements are straightforward, like the upgrade to the SmartFusion2, it is interesting to observe the consumption data of these reference systems in order to understand where and under which power mode the largest current draw occurs.

Figure 1 depicts the energy consumption of the MarmotE SDR platform in transmit, receive, and sleep modes. It represents the measurements performed by Szilvási et al. [11, 10]. Specifically, they implemented a Binary Frequency-Shift Keying (BFSK) PHY layer, that was chosen to be able to compare the measurements with the CC1000

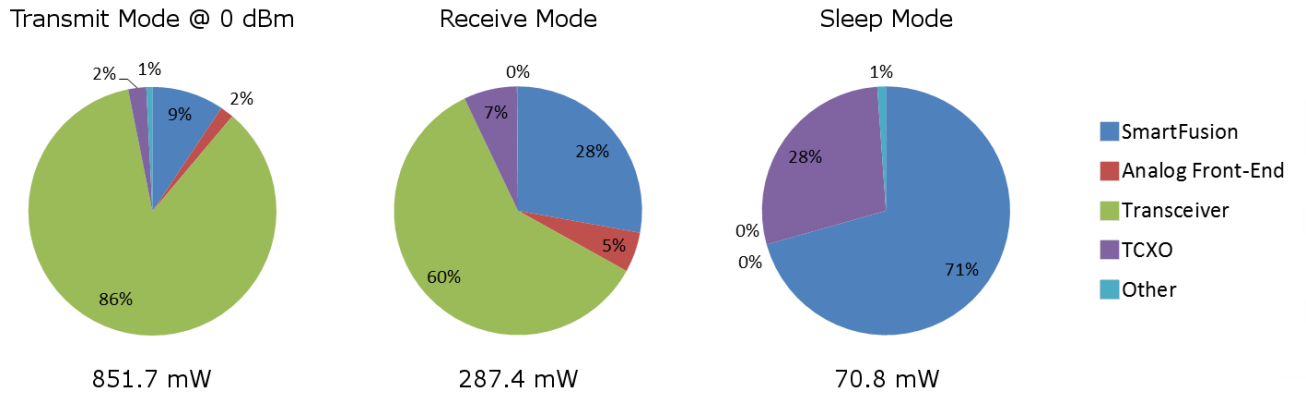

Figure 1. Power consumption of the MarmotE SDR platform in various operation modes.

transceiver, used by reference WSN nodes, MICAz and TelosB. It is assumed that the power consumption of the μSDR platform will be very similar to the MarmotE SDR, considering that the main parts, which are responsible for the largest energy expenditure, are identical or analogous in both platforms.

The first thing that stands out regarding Figure 1 is that the radio transceiver used is responsible for more than half of the power consumption in those modes in which it is active, and it reaches 86% of the total consumption in the transmit mode. Moreover, its I/Q interface is analog, so a bidirectional analog-to-digital conversion stage, also called Analog Front-End (AFE), is needed in between the transceiver and the processing back-end. Since the AFE included in the SmartFusion has a maximum sample rate of 600 ksps, which seriously limits the baseband modulation rates, an external 22 Msps ultra-low power AFE is used in the platform, specifically the Maxim Integrated MAX19706. This device has a power consumption of 15 mW when operating, which must be added to the one of the transceiver. Therefore, these two components constitute one of the most critical parts of the design, with a high impact in the final energy utilization. Thus, it would be very desirable to find a radio front-end solution with a lower energy consumption.

On the other hand, as CWSNs are mainly orientated to low duty cycle applications, the sleep mode is where most time is spent, thus its consumption must be as lower as possible in order to extend the network lifetime. This is where the ultra-low power modes provided by the Microsemi's SmartFusion2 will play a huge role. While the SmartFusion is responsible for a 50 mW consumption in sleep mode, as shown in Figure 1, the SmartFusion2 includes important low power features that greatly lower it. Specifically, in Flash*Freeze mode, the FPGA fabric consumes only 1 mW, while during normal operation, the static power of a 50K gate model is 10 mW. Moreover, thanks to the flash-based technology, it takes no more than about 100 μs to enter or exit Flash*Freeze mode, and the device is able to retain SRAM content, I/O state, and register data. Thus, after exiting this mode, the SmartFusion2 continues to operate where it left off, drastically reducing power consumption. Apart from these features, the ARM Cortex-M3 processor and the SoC peripherals can also be put into low power states if required [6].

Finally, it is noteworthy that the Temperature Compensated Crystal Oscillator (TCXO), which is used to drive the radio transceiver stably and accurately, has a fixed energy consumption of 20 mW. This is because it is always enabled by design, even in sleep mode. Such kind of details are critical and should be avoided in low-power systems, whose sleep modes should be able to turn off all non-essential parts.

Having reviewed the MarmotE SDR in depth, if we compare the above information with the power consumption of the CC1000, we can conclude that even with all the proposed improvements there is still a big difference between them. Specifically, based on the specifications and assuming a 3.3 V supply voltage, the crystal oscillator turned-on, a BFSK modulation in the 433 MHz band and an output power of 0 dBm, the consumption data of the CC1000 in transmit, receive and sleep modes are just 34.3 mW, 24.4 mW and 0.66 μW respectively. That difference is the cost of having a more flexible node architecture with enough resources to acquire cognition, and it must be compensated by the efficiency gains achieved at a network level. That is, what matters is not only the consumption of a node in each operating mode, but also the total energy spent by the network to provide a certain service, which will depend on the efficiency of the communication between the nodes within it. Therefore, cognitive capabilities should be used to reduce the time spent by nodes in transmit and receive modes. At the same time, it should be noted that as the problem of spectral saturation gradually increases, nodes without cognitive capabilities will have more problems with their communication and, therefore, will spend more time and energy on this task.

4 Hybrid Methodology to Operate CWSN Nodes

As reviewed in the previous section, even though the consumption of SDR platforms can be significantly lowered, it is still much higher than the one of traditional fixed

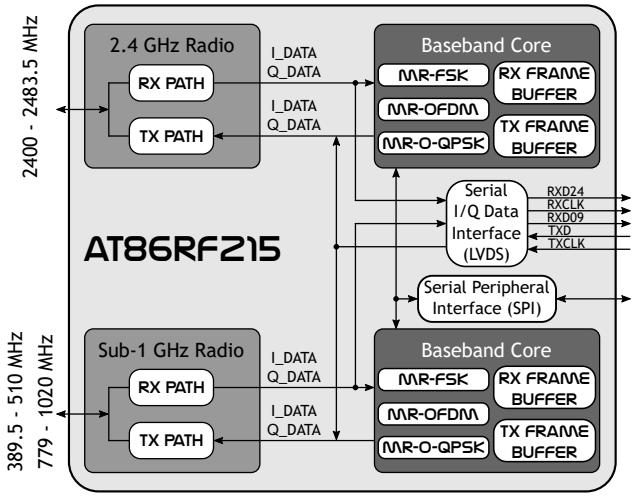

Figure 2. AT86RF215 block diagram.

hardware transceivers. Thus, to achieve a more efficient solution, suitable to permit the deployment of CWSNs long-term testbeds, it is essential to consider additional strategies to operate it in a more efficient manner, especially the radio stage, which usually dominates the system power budget.

Undoubtedly, CWSNs greatly benefit from the SDR technology in regard to certain tasks, such as the spectrum sensing or the customization of the lower layers of the communication stack. However, considering real possible applications, there is a substantial percentage of the time in which it is not necessary to use any of those power-hungry cognitive features. For example, after sensing the radio environment and adapting some communication parameters, nodes can operate normally until they start to have communication problems again, and they need to repeat the sensing-adaptation process. Another example would be if a CWSN employs a custom modulation only for the control channel, using a standard one for the rest of the data. In those moments, when no special task or customization is required, the use of an SDR system results in a considerable penalty in terms of power consumption and resource usage. Conversely, the use of a fixed function transceiver in these situations, such as the CC1000, would allow us to take advantage of its higher efficiency in those terms. This would mean a great reduction in the nodes' consumption, while the network could still benefit from having cognitive features.

In addition, the Hardware Description Language (HDL) implementation and evaluation of the PHY/MAC layers for SDR systems implies a major commitment of time and money because of its complexity. This is a price that must be paid when working with SDRs, even when these layers do not include any cognitive feature.

Thus, we propose that the optimal solution in all the aforementioned aspects is to be able to change between an SDR and a traditional transceiver, and only use the first one for those specific tasks that strictly require it. This hybrid methodology to operate nodes, conceived as a compromise between hardware flexibility and energy efficiency, supposes a new approach for the CWSN field.

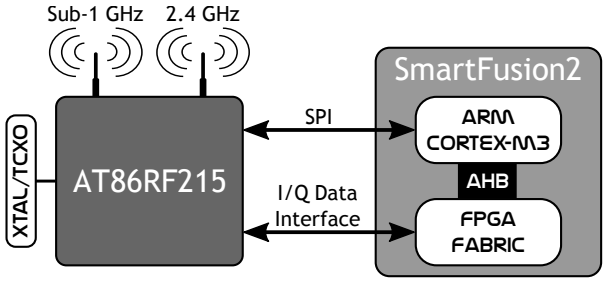

Figure 3. MIGOU platform architecture.

5 Platform Architecture

As a second contribution of this work, we present MIGOU, a platform able to operate according to our proposed hybrid methodology. In this section, we discuss its architecture and main electronic components, which have been carefully selected attending to the limited resources of sensor nodes.

The Atmel's AT86RF215 (Figure 2) is a device comprised of two highly-integrated radio transceivers that is specifically designed to minimize the number of external components required on the Printed Circuit Board (PCB). It allows simultaneous independent reception in both transceivers, which operate in the sub-1 GHz and 2.4 GHz bands respectively. Moreover, each one is paired with an optimized baseband core that supports a wide variety of data rates with three modulation schemes. This allows a highly efficient implementation of PHY layers compliant with the IEEE 802.15.4g-2012 and the ETSI TS 102 887-1 standards, and the use of other proprietary ones.

Alternatively, users can route the I/Q data stream directly to/from an external processor. This is done through a serial Low Voltage Differential Signal (LVDS) interface. Furthermore, since this is a digital interface, the need for an intermediate AFE stage between both devices is eliminated.

With the mentioned features, each radio transceiver of the AT86RF215 supports two operating modes, which do not need to be the same. In baseband mode, the internal highly-optimized baseband core controls the radio and processes the data encoding/decoding for transmitting and receiving from the internal frame buffers. However, this core is limited to the supported modulations and configurations options. On the other hand, in the I/Q radio mode, the radio is controlled from an external microcontroller via a Serial Peripheral Interface (SPI) and the data are exchanged via the LVDS interface, allowing the use of an external baseband processor for implementing custom SDR features to perform cognitive strategies. The ability to switch between these two modes of operation allows a node to carry out the hybrid methodology proposed in the previous section.

Therefore, most of the required elements of our low-power platform with SDR capabilities are found mainly in two chips that can be connected directly: the AT86RF215 as a radio front-end, and the SmartFusion2 SoC as a control unit and mixed-signal processing back-end. Specifically, the ARM Cortex-M3 contained in the SoC is responsible for setting and controlling the transceiver via SPI. When

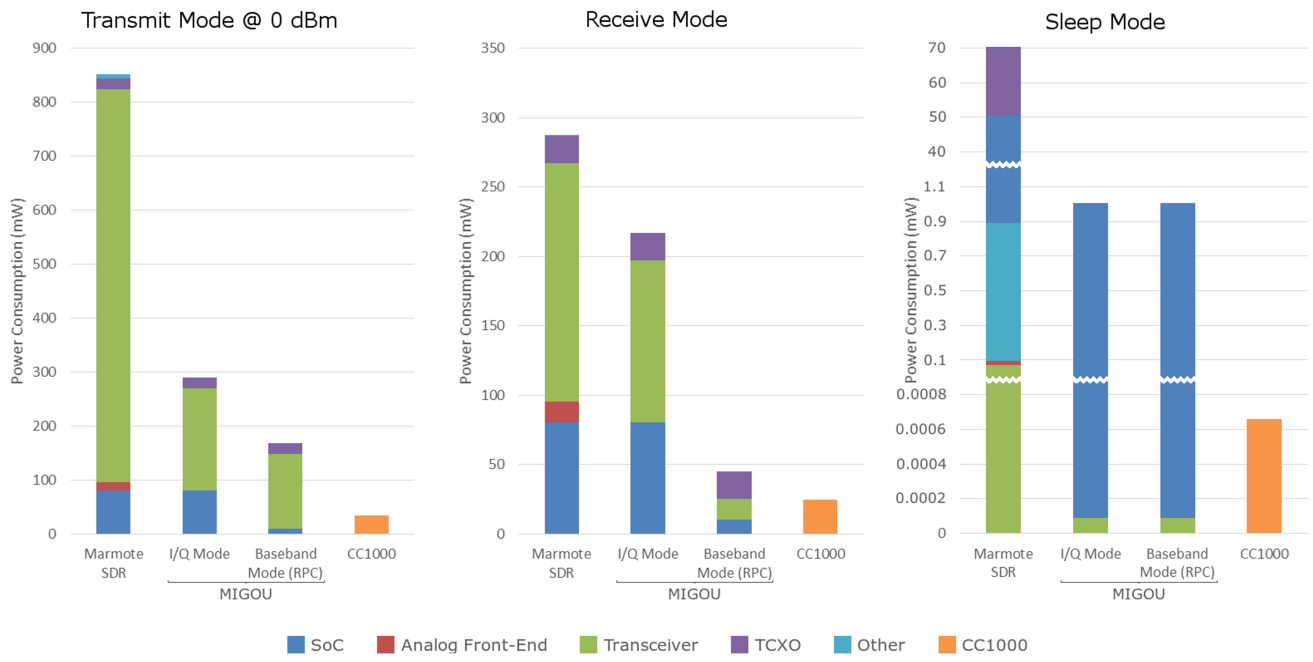

Figure 4. Power consumption comparison in milliwatts (mW).

the platform is in baseband mode, these are the only two parts involved. On the other hand, when any custom processing must be performed, it will take place in the FPGA fabric, which is connected with the AT86RF215 via the I/Q data interface, as shown in Figure 3. The internal communication between the FPGA fabric and the ARM Cortex-M3 is based on the AMBA High-performance Bus (AHB), which provides enough bandwidth and avoids data bottleneck problems. Thus, the boundary between software and hardware has a great flexibility, which is very desirable for SDRs.

Moreover, this proposed platform not only allows to implement our hybrid methodology aimed at reducing power consumption, but it also represents a reduction of current draw even operating in a single mode, because of the new components selected. This can be seen in Figure 4, where a comparison between MIGOU, the MarmotE SDR and the CC1000 is shown, all of them using a binary FSK modulation. The power consumption data of the MarmotE SDR and the CC1000 were extracted from [11], while the ones of the proposed platform were calculated based on datasheets' specifications. For the AT86RF215, we assumed that the 2.4 GHz radio front-end is employed while the sub-1 GHz transceiver is in sleep mode, the power supply is 3.0 V and, when receiving in baseband mode, the operation with Reduced Power Consumption (RPC), with on-off time 1:16, is enabled. For the SmartFusion2, we employed the information provided in [6], which is 1 mW of consumption in Flash*Freeze mode and 10 mW of static power in baseband mode, when any processing task has to be performed by the SoC. When the platform is in I/Q mode, the modulation/demodulation takes place in the FPGA, so we assumed the same consumption as the MarmotE SDR, which employs the SmartFusion. Finally, the same criterion was used for the consumption of the TCXO except for the sleep mode, where it will be turned off as explained in Section 3.

Regarding this comparison, it can be stated that our proposed system leads to a great reduction in power consumption compared with the MarmotE SDR reference platform. In fact, it is closer to the CC1000 integrated radio chip in those terms. That reduction is especially large in sleep mode, which is where most time is spent and, therefore, where the impact on the final application consumption will be higher. Moreover, as expected, there is also a remarkable difference in energy consumption between the I/Q and baseband modes, especially in reception. Thus, exploiting our hybrid methodology to find a good balance between the use of the I/Q mode for implementing cognitive features, and the use of the baseband mode for saving energy, will have significant benefits in terms of power consumption and it will be an important line of research in the future.

On the other hand, the platform is equipped with dedicated hardware to monitor its own power consumption while it is operating, facilitating its characterization. This feature is essential for experimental purposes, but it is also highly recommended for real applications as such information can be used as an input in cognitive algorithms.

Finally, the size and price of the SmartFusion2 and the AT86RF215, as the major components of MIGOU, confirm the feasibility of implementing a platform for the CWSN field with a similar size as traditional nodes, and a cost around 100€, depending on the resources of the selected device models.

6 Conclusions

Given the spectrum scarcity problem, mostly in unlicensed ISM bands, and the forecasts regarding the increasing pervasion of wireless communications, it is essential to optimize the use of the spectrum to ensure the proper functioning of services and devices in the near future.

To achieve this, SDR-capable platforms seem to be the best option as they give full access to the entire communication stack and offer a high flexibility to customize it, greatly improving the ability to sense and adapt to the radio environment. However, the consumption of the current solutions is too high for real battery-powered devices. This prevents application areas of great interest, like WSNs, IoT or smartphones, to benefit from their valuable features.

In this work in progress, we face this challenge with a double contribution. On the one hand, we propose a hybrid methodology of nodes operation that consists of exploiting SDR capabilities only for those actions that strictly require them, using traditional fixed hardware transceivers, that demand less node's resources, for the remaining tasks. This approach represents a promising paradigm in the field of CWSN, and opens a wide range of possibilities as it allows to add any specific functionality to the different layers of the communication stack, while keeping the option of using highly optimized cores for the standard operation.

On the other hand, we present the architecture and the main electronic components needed for the development of a platform able to operate according to this methodology. These components have been carefully selected based on the analysis of the consumption of the reference platforms. Thus, the use of the SmartFusion2 SoC as the processing back-end, achieves a reduction in consumption of an order of magnitude in sleep mode, which for low duty-cycled applications, such as CWSNs, has a huge impact in terms of energy efficiency. Moreover, the AT86RF215, as a radio front-end, not only enables the implementation of the aforementioned methodology with a single radio chip, but it can also operate in multiple ISM bands, adding an extra flexibility to the platform.

Finally, the reduced number of components required, as well as their prices and sizes, confirm the feasibility and suitability of the platform for applications based on autonomous nodes with limited resources. Then, when finished, MIGOU will allow to test and evaluate existing works in real scenarios, and to develop new cognitive solutions aimed at improving the efficiency of wireless communications.

7 Acknowledgments

This work was supported by Mundo Reader S.L., under the BQ brand as part of its BQ Chair program.

8 References

[1] O. Akan, O. Karli, and O. Ergul. Cognitive radio sensor networks. *IEEE Network*, 23(4):34–40, July/August 2009.

[2] V. Blaschke, S. Nagel, and F. Jondral. Mechanisms for the adaptation of the physical layer in a cognitive radio. In *Proceedings of the 9th European Conference on Wireless Technology (ECWT)*, pages 43–46, September 2006.

[3] Cisco Systems Inc. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2015-2020. Accessed: 2016-04-18.

[4] P. Dutta, Y.-S. Kuo, A. Lédeczi, T. Schmid, and P. Völgyesi. Putting the Software Radio on a low-calorie diet. In *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks (HOTNETS)*, pages 1–6, October 2010.

[5] Y.-S. Kuo, P. Pannuto, T. Schmid, and P. Dutta. Reconfiguring the software radio to improve power, price, and portability. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems (SenSys)*, pages 267–280, November 2012.

[6] Microsemi Corporation. SmartFusion2 Lowest Power FPGAs White Paper, 2012. Accessed: 2015-12-10.

[7] J. Mitola III. *Cognitive Radio. An integrated agent architecture for software-defined radio*. PhD thesis, Royal Intitute of Technology (KTH), Stockholm, Sweden, May 2000.

[8] A. Mitra. A PHY-MAC cross layer design for low PER with adaptive modulation and coding. In *International Workshop on Cognitive Radio (IWCR)*, pages 1–5, December 2010.

[9] Y. Qin, Z. He, and T. Voigt. Towards accurate and agile link quality estimation in wireless sensor networks. In *Proceedings of the 10th IFIP Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net)*, pages 179–185, June 2011.

[10] S. Szilvási. *Advanced RF techniques for Wireless Sensor Networks: the Software-Defined Radio approach*. PhD thesis, Institute for Software Integrated Systems (ISIS), Vanderbilt University, Nashville, TN, USA, May 2014.

[11] S. Szilvási, B. Babják, P. Völgyesi, and Á. Lédeczi. MarmotE SDR: Experimental platform for low-power wireless protocol stack research. *Journal of Sensor and Actuator Networks (JSAN)*, 2(3):631–652, September 2013.