A Novel wake-up Receiver with Addressing Capability for Wireless Sensor Nodes

Chiara Petrioli*, Dora Spenza*, Pasquale Tommasino† and Alessandro Trifiletti†
*Computer Science Deparment, University of Rome "La Sapienza"
Email: {petrioli, spenza}@di.uniroma1.it
†Department of Information, Electronical and Telecommunications Engineering, University of Rome "La Sapienza"
Email: {tommasino, trifiletti}@die.uniroma1.it

***Abstract*—Emerging low-power radio triggering techniques for wireless motes are a promising approach to prolong the lifetime of Wireless Sensor Networks (WSNs). By allowing nodes to activate their main transceiver only when data need to be transmitted or received, wake-up-enabled solutions virtually eliminate the need for idle listening, thus drastically reducing the energy toll of communication. In this paper we describe the design of a novel wake-up receiver architecture based on an innovative pass-band filter bank with high selectivity capability. The proposed concept, demonstrated by a prototype implementation, combines both frequency-domain and time-domain addressing space to allow selective addressing of nodes. To take advantage of the functionalities of the proposed receiver, as well as of energy-harvesting capabilities modern sensor nodes are equipped with, we present a novel wake-up-enabled harvesting-aware communication stack that supports both interest dissemination and convergecasting primitives. This stack builds on the ability of the proposed WuR to support dynamic address assignment, which is exploited to optimize system performance. Comparison against traditional WSN protocols shows that the proposed concept allows to optimize performance tradeoffs with respect to existing low-power communication stacks.**

I. INTRODUCTION

In the last decade, considerable research efforts have been devoted to address the energy bottleneck problem that severely limits the lifetime of battery-power wireless sensor motes. As communication is widely regarded as the major source of power consumption, the leading type of solution to enable long-lasting Wireless Sensor Network (WSNs) is currently to adopt protocols that operate at very low duty cycles. However, even if such approach allows to extended the typical duration of a node's battery, it suffers from reduced performance due to the inherent tradeoff between energy efficiency (i.e., low duty cycling) and data latency.

The need to lower the energy consumption of communication without increasing data latency has recently driven research on WSNs towards architectures based on wake-up radios. The use of a wake-up receiver enables on-demand communication, allowing motes of a sensor network to keep their main transceiver off unless data transmission or reception is needed. This approach virtually eliminates the need for idle listening, which is a critical source of energy consumption in traditional WSNs [1]. In addition, if the wake-up radio supports selective addressing of nodes, energy wastage due to overhearing can also be significantly reduced.

In this paper we describe the design and prototyping of a novel wake-up receiver (WuR) architecture operating in the 2.4 GHz ISM band. The proposed concept, based on an innovative pass-band filter bank with high selectivity capability, allows selective addressing of nodes by combining both frequency-domain and time-domain addressing space. Multiple wake-up addresses can be dynamically assigned to nodes. Moreover, as we used for the WuR the same ISM band employed for WSN communication, there is no need for additional hardware to transmit wake-up addresses. Rather, off-the-shelf motes such as TelosB can send wake-up sequences by using their main transceiver. The developed wake-up radio allows to optimize the latency vs. energy consumption performance tradeoffs with respect to existing low-power WSN communication stacks, which achieve energy saving by requiring nodes to follow low duty cycles. The additional feature of the developed concept, i.e., the ability to assign multiple wake-up addresses to nodes and to dynamically change them, brings in an additional feature: The possibility to exploit dynamic address assignment to optimize system performance. We show how to take advantage of this functionality by proposing a novel wake-up-enabled communication stack that provides both interest dissemination and convergecasting primitives. This communication stack also includes support for energy harvesting capabilities embedded in modern sensor nodes, which allows to use energy readily available from the environment (e.g., from solar light, wind, body movements, etc.) to supplement or even completely replace traditional batteries [2].

The contribution of this paper is as follows:

- We propose a novel topology of wake-up radio with addressing capability, and describe its prototyping and validation through lab and in-field experiments.
- We prove power consumption scaling of the proposed WuR by carrying out simulations in microelectronic technology of the most critical sections.
- We present a wake-up-radio-enabled low-power low-latency communication stack that outperforms traditional duty-cycle approaches.

The remainder of this paper is organized as follows: After discussing related work in Section II, we present the architecture of the designed wake-up receiver in Section III. In Section IV we describe a novel communication stack that exploits the features of the proposed design to support both interest dissemination and convergecasting primitives. Section V reports the result of our performance evaluation, including experimental validation of the WuR design, proof of consumption scaling, and performance comparison of the proposed communication stack against traditional WSN protocols. Finally, Section VI concludes the paper.

II. Related Works

Several technological solutions have been recently under investigation to design wireless sensor networks based on wake-up radios. A critical aspect is the best tradeoff between receiver sensitivity (and therefore wake-up range), and power consumption. Passive wake-up receivers are a very attractive approach, because they operate making solely use of the power available in the radio spectrum, without requiring any external energy supply. However, they generally have limited sensitivity, so the wake-up range they can reach is significantly shorter than the typical communication range of a sensor node. Gu and Stankovic presented in [3] one of the earliest passive radio-triggered circuit, able to operate at 10 feet (about 3 meters) from the source of the radio signal. Based on SPICE simulation results, enhanced designs and the use of ultra-low-power active components enable the operating range to be extended to up to 100 ft with 55 ms wake-up latency. More recently, Chen et al. proposed an architecture for passive wake-up that combines an energy-harvester circuit with an ultra-low-power pulse generator to trigger mote wake-ups within a range of up to 37 ft [4]. Active WuRs require external power supply, but have higher sensitivity than passive wake-up receivers, and can thus reach significantly longer operating ranges. For example, the wake-up receiver presented by Drago et al. in [5] achieves -82 dBm sensitivity leveraging a crystal-less architecture that combines non-coherent energy detection with a broadband-IF super-heterodyne scheme. A receiver sensitivity of -83 dBm is achieved by the 868 MHz OOK receiver proposed in [6], which comprises an analogue superheterodyne front-end and two digital 31 bit correlating decoders. Other recent approaches include the use of hardware to be attached to the node, based on discrete components, subsystems, or RFID [7] to get a cost-effective off-the-shelf solution. In [8] a board implementing a wake-up radio operating at 868 MHz is proposed, which allows to reach a sensitivity of -52 dBm that results in a wake-up distance of up to 40 meters at an output power level of +10 dBm in the transmitter. The node presented in [9] by Gamm and Reindl, implemented on a PCB, makes use of a commercially available WUR with addressing capability at 125KHz, which allows a maximum working range of 90m in open air field with a sending power of +20 dBm and a packet error rate of 5%. Other recent solutions, with varying performance in terms of sensitivity, maximum reachable wake-up distance, latency and power consumption, are surveyed in [10], [11]. Some of the existing approaches are range-based, meaning that every node within the operating range is waken up whenever a radio trigger signal is transmitted, while others also support ID-based wake-ups. Despite the increasing choice of wake-up architectures nowadays available, however, so far comparable little attention has been devoted to the design of communication protocols based on radio-triggering features. In particular, to the best of our knowledge, no work to date investigates communication stacks based on the use of a wake-up radio to selectively wake up individual nodes and/or group of nodes depending on their *state*, thus associating a semantic meaning to wake-up addresses.

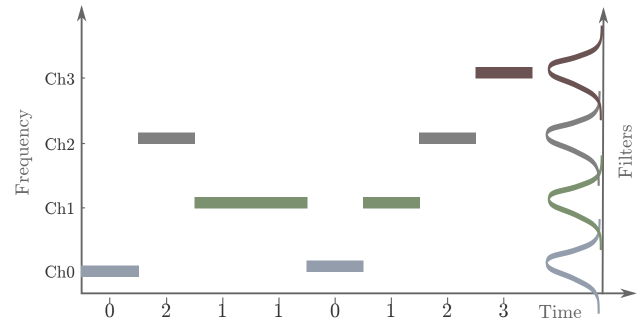

Fig. 1. Example of a wake-up sequence transmitted over four IEEE 802.15.4 channels: The address '02110123' is sent.

III. Architecture of the proposed wake-up receiver

Our proposed design exploits incoherent detection. As the receiver is not synchronized in phase with the transmitter, the use of a power-hungry Phase-Locked Loop is not required, thus reducing power consumption. As we take as a reference the IEEE 802.15.4 standard, the overall RF receiver has to approach a sensitivity of -85dBm. Such sensitivity has to be fulfilled with a Signal to Noise Ratio (SNR) greater than 12 dB. An instantaneous dynamic range (DR) greater than 30dB has been also specified, by considering a minimum spurious rejection of 42dB. Other requirements include support for dynamic address reprogramming and for assigning multiple addresses to the same receiver. In our WuR design, a node address consists of a sequence of continuous-wave or narrow-band pulses belonging to some of the 16 channels of IEEE 802.15.4 standard operating in the 2.4 GHz ISM band. Wake-up signals are OOK (On/Off Key) modulated in the same band by selecting a certain number c of carriers. Figure 1 shows an example of a wake-up sequence transmitted over four IEEE 802.15.4 channels.

The novel architecture proposed for the wake-up receiver is shown in Figure 2. The critical portion of the receiver consists of the chain of the Low Noise Amplifier (LNA) and the filter bank. Both of them have to be very-low-power blocks, but the LNA has to guarantee gain and sensitivity, whereas the filters have to show high selectivity to reject noise and interference, together with center frequency stability. The bank of band-pass filters and peak detectors allow detection of waveforms. The output of the receiver is a set of c DC instantaneous voltages proportional to the incoming signal at frequency $f_i, i = 1 \dots c$. Then, an additional digital board provides analog to digital conversion of the c DC voltages, decoding of the received address and triggering of wake-up interrupts to the node.

The choice of parameter c is an important requirement for the WuR design. Increasing values of c results in larger addressing space without increasing the wake-up time, but they also require a higher selectivity of each filter. The filter bank is thus a critical component of the overall wake-up receiver. At early design stage, we considered both passive surface acoustic wave (SAW) filters and active-inductor filters for filter bank implementation. However, performance stability requirements and preliminary measurements have driven the choice towards active filters, even if larger power consumption is expected. Our chosen filter topology is based on the classical

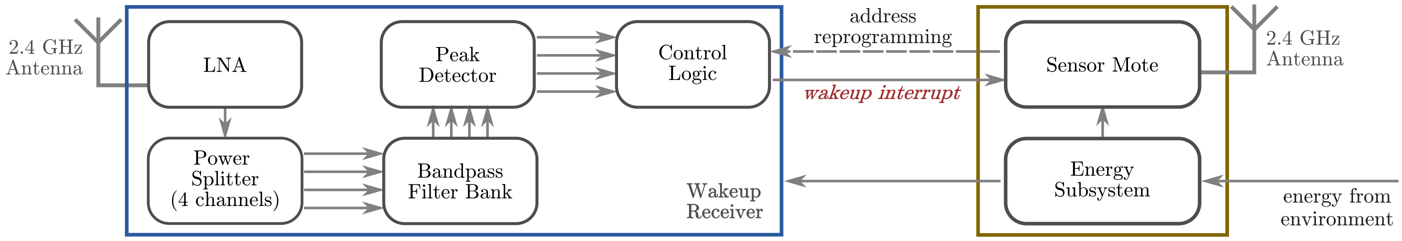

Fig. 2. Block diagram of the sensor mote with wake-up receiver and energy subsystem with harvesting capabilities. The depicted wake-up receiver uses 4 channels for node addressing.

L-C tank, where the inductor is not a passive component. In fact, in integrated technology the required filter selectivity cannot be reached by using spiral inductors due to inductor loss, corresponding to a Q not greater than 15.

The filter bank board has been designed by exploiting a filter topology making use of active inductors [12], [13]. Several first order LC filters based on active inductors with very narrow 3dB bandwidth have been presented in previous works. This is usually obtained by the use of negative series resistance with high risk of practical instability, and therefore of oscillation. On the contrary, we think that active inductor should always be unconditionally stable; as a consequence high Q in filters should be obtained by using higher order filters composed by stable low-loss active inductors and passive capacitances. We remark that this aspect is usually underestimated in the literature and, to our best knowledge, only few works give information about AI stability, power handling and noise. In most previous works the inductance value, as well as its quality factor, are controlled by acting on the bias current of the active devices that sets the transconductance. This approach limits the linearity of the active inductor for relatively high signal power, because it requires a bias-dependent small-signal gain, and therefore a strongly nonlinear trans-characteristic curve of the active device. Because the active inductor is an intrinsically noisy circuit, the resulting dynamic range is usually small, then often useless in RF front-end circuitry. For these reasons our approach in the presented Active Inductor (AI) topology is to control both the inductance value and the series resistance without affecting the linearity of the active component of the gyrator, that can therefore be high. This approach is based on the insertion of a suitable variable passive compensation network in the gyrator. Both the unit-gain voltage amplifier and the inverting transconductance amplifier are highly linear, with fixed gain. The compensating network provides variable attenuation and delay.

IV. Software design: wake-up-enabled harvesting-aware communication stack

In the following the describe a wake-up-enabled communication stack that exploits the features of the proposed WuR design to support both interest dissemination and convergecasting primitives.

A. Interest Dissemination

The interest dissemination primitive refers to transmission of commands from the sink to the sensor nodes, which is typically implemented by means of either low-power broadcasting or multicasting primitives. In the following we focus on broadcasting, as the same approach can be followed to enable communication on any overlay created by existing multicasting protocols.

In the literature different schemes have been proposed to implement broadcasting. Flooding is the basic implementation of broadcasting, which however suffers from broadcast storm and from collisions in case a reliable (acked) version of the flooding is implemented. In case nodes do not follow a duty cycle, each node transmits only one copy of the broadcast packet. However, in this case the energy consumption is high and mechanisms to avoid congestion in case of dense networks are particularly critical. When nodes follow a low duty cycle either all neighbors ON times are synchronized (the toll to pay being the control message exchange and complexity typical of synchronous approaches) or each node must transmit the broadcast packet multiple times to be able to reach all neighbors, which significantly increases energy consumption.

A first contribution is therefore a new flooding protocol, called FLOOD-WUP, that solves these problems achieving high reliability and better latency vs. energy consumption performance. In FLOOD-WUP nodes are assigned shared broadcast addresses[1]. Without loss of generality, we assume that each node is assigned two broadcasting wake-up addresses, named w_a and w_b. This concept can be generalized to N addresses, N being set as the maximum number of different broadcast packets which can be generated in the network during the interest dissemination. The sink sends the first broadcast packet to its neighbors, preceding it with the wake-up sequence w_a. The following broadcast packet will be sent using the wake-up sequence w_b, the third using the wake-up sequence w_a, and so on. To each broadcast packet a unique sequence number is associated. Nodes in the network are initially in sleep, i.e., with their MCU and radio set to power down mode, with the wake-up radio active with address set to w_a. When a node receives a wake-up signal with sequence w_a, it wakes up and sets its radio to RX mode to receive the broadcast packet B. Upon reception of B, each node changes its broadcast wake-up address from w_a to w_b. This means that nodes that have already received the broadcast packet will not wake up when duplicated packets are transmitted, as the the sequence w_a that precedes each duplicated packets is no longer active for those nodes. After receiving the broadcast packet B, a node waits a random time and retransmits the broadcast

[1] In case the same WSN infrastructure is shared by multiple systems, which can be seen as different overlays exploiting the WSN, different orthogonal broadcast addresses can be assigned to each overlay. The same applies in case overlays enabling multicast communications have been determined by means of multicast protocols.

packet preceding it with the same wake-up sequence w_a.

For reliability purposes, each node can retransmit the same packet multiple times (a number n_{broad} of times) to compensate for collisions (kept low by the jitter), PER in the broadcast packet transmission, and possible wake-up sequence false negatives. Each of those packets is retransmitted by using the same wake-up sequence it was received with. However, it could happen that a node, due to faults or to wake-up sequence false negatives[2] that are not solved by multiple packet transmissions, loses connection to the proper sequence of broadcast wake-up signal transmissions (i.e., that it expects w_a to be transmitted over the network when w_b is transmitted, or vice versa). In this case, a node detects it missed a broadcast packet based on the sequence number of the last packet it received, and requests the lost packet to be retransmitted by sending a dedicated control message to its neighbors.

In FLOOD-WUP, nodes consume energy only when they receive or transmit packets. Energy consumption due to reception is limited to the reception of the first copy of the broadcast packet, as nodes do not wake-up when duplicate broadcast packets are transmitted. Energy consumption for transmission is limited to the transmission of the broadcast packet and the wake-up sequence once or, if re-transmissions are envisioned to add robustness, n_{broad} times.

B. Converge Casting

In this section, we propose a hop-count-based converge-casting protocol, named GREEN-WUP, which is energy-harvesting and wake-up-radio aware. In GREEN-WUP packet forwarding is realized by opportunistically selecting the next-hop relay among the *whole* set of neighbors, rather than just among the awake ones. To this end, we assume that nodes activate a subset of possible unicast wake-up addresses based on their state. In other words, the network has agreed on wake-up sequences that have a semantic meaning. In particular, nodes activate a given wake-up sequence w_i (i.e., they wake-up if such sequence is transmitted) in case they are in the corresponding state i. Such a state expresses how good a node is to serve as a relay, and can be computed based on factors such as its current and expected energy intake from harvesting, its residual energy, its hop count, its queue size, the quality of its channel, and so on.

In the specific case of GREEN-WUP, we propose that each node is associated with a wake-up address w composed by two subsequences, i.e., $w = w_l w_e$. The first subsequence, w_l, is set based on the node hop count, which is obtained by hop count information piggybacked in broadcast packets during the initial interest dissemination phase. The second subsequence, w_e, is set by each node based on its residual energy. Energy levels are discretized into a given number k of classes. By computing its fractional residual energy (normalized with respect to a nominal maximum battery size), each node determines the class index i it belongs to, and it sets the subsequence w_e of its address accordingly. In GREEN-WUP, to higher energy class indexes correspond higher normalized fractions of residual energy. The class with index k is used to indicate nodes that, due to their energy state, are preferred relays. In particular, environmentally-powered nodes whose battery is at capacity may experience periods of *energy peak*, during which they harvest energy at a rate that exceeds their current power consumption. Since this excess energy would be wasted if not used immediately, a node experiencing an *energy peak* sets $w_e = k$ to reflect the fact that it is ideally able to perform forwarding at no cost. Nodes having their battery level greater than a given threshold are also included in this class. The sequence w is periodically recomputed by each node to reflect changes in its hop count and energy state, thus resulting in modifications of the activated wake-up address when the node local state changes.

[2] False positives are not a problem, since nodes that erroneously wake up will receive a duplicated packet, which can be detected based on its sequence number.

Whenever a node N with hop count h has a packet, either self-generated or previously received, to transmit, it sends a wake-up sequence w addressing, among the nodes with hop count $h-1$, those with $w_e = k$. Sequence w is followed by a short packet that includes a temporary wake-up address w_N, which potential relays can use to contact back the sender. Node N then goes to sleep. Receiving nodes that wake up and correctly receive the packet pick a random jitter, and then answer with a CTS packet. The CTS, preceeded by the wake-up sequence w_N, includes the unicast wake-up sequence that node N can use to contact the answering potential relay. After transmitting the CTS packet, potential relay nodes go back to sleep. Node N picks the best relay and sends the unicast wake-up sequence to wake it up, followed by the actual data packet transmission. If no node answers the initial wake-up sequence, as there are no neighbor belonging to the addressed class, the sender node iterates the procedure, addressing nodes with hop count $h-1$ and energy state $k-1$. Then, if no CTS packet is received, node N tries addressing nodes with hop count $h-1$ and energy state $k-2$, and so on, until a relay is found[3]. Once a relay is selected, the sender transmits its buffered packets in a burst, performing a back-to-back transmissions in which data packets are individually acknowledged.

V. Performance Evaluation

A. WuR prototyping and design validation

To assess the performance of the proposed WuR architecture, a prototype has been designed and implemented by means of COTS components, except for the filter bank that has been fabricated on RF-60A substrate using standard microwave SMD passive components and BFP420 bipolar transistors (Fig. 3). As the topology of the filter is novel and its performance were not fully known at the beginning of the design, $c = 4$ channels were considered for the prototype hardware. The filter bank is designed to provide four channels at nominal frequencies f_1 = 2410 MHz, f_2 = 2435 MHz, f_3 = 2455 MHz and f_4 = 2480 MHz. The central frequencies f_i are tunable to allow for receiver re-configurations.

To make the evaluation of each block of the prototype easier, the RF portion of the receiver is subdivided into three separate connectorized boards (Fig. 4). The first board contains

[3] In case no node is found with hop count $h-1$ it could be the case that the estimated hop count is no longer valid, e.g. due to faults or link dynamics. In this case the node downgrades its hop count, sending a broadcast wake-up sequence to its neighbors and a control packet through which it asks for their hop count. It then sets its hop count to the minimum value received plus 1.

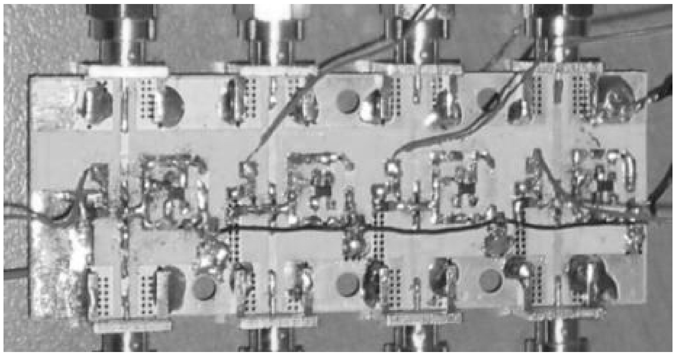

Fig. 3. Prototype board of the filter bank providing four channels at nominal frequencies 2410, 2435, 2455 and 2480 MHz.

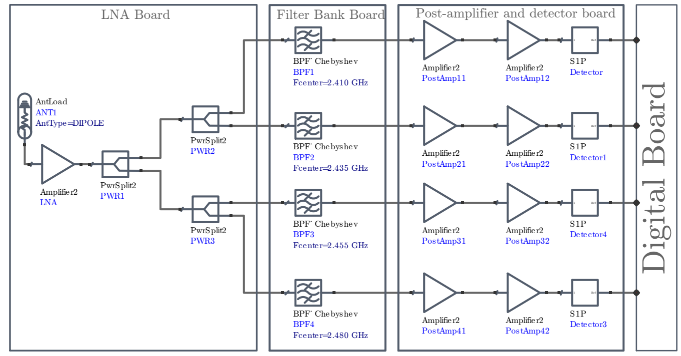

Fig. 4. Block scheme of the overall WuR prototype, showing the three separate connectorized boards and the digital board.

the antenna, the low-noise amplifier and three 1:2 power splitters. The second board hosts the filter bank, designed accordingly to the AI technique discussed in Section III. The architecture is the same for all the c filters. The bandpass filter is a shunt resonator, with the capacitance C and the proposed grounded AI. The variable capacitance C has been realized with the Infineon BBY58 varactor. This allows a flexible use of the filter bank, as each filter can be tuned at the desired frequency minimally affecting attenuation and bandwidth. The third board contains two post-amplifier stages and a RF power detector. The MAX2015 logarithmic detector allows DC conversion of the RF waveform, with an 18 mV/dB gain and sensitivity equal to -45 dBm. Therefore, a 38dB of gain in the linear receiver portion is sufficient to guarantee an overall sensitivity of -83 dBm.

Validation of the described prototype included experimental evaluation of the performance of both the individual boards and the overall RF receiver. As a first step, accurate characterization and fine tuning of the filter bank board has been performed by removing the antenna and by feeding the input of the receiver with a signal source. Three different filter bank prototypes have been designed, fabricated, and measured. The first one is composed of passive L-C filters, the second one of active filters with fixed central frequency and the third one of active filters with tunable central frequency by voltage control on varactor diodes. Experimental measurements on the RF portion of the receiver have shown that passive filters show limited selectivity, while active filters with varactors provide more robustness in terms of performance variation versus component technological spread with respect to the ones with fixed central frequency. Therefore, active filters with varactors has been chosen to compose the overall receiver. The obtained dynamic range of the filter from noise floor to -1dB compression point is 70dB with a power consumption of about 1mW. During the second characterization step, performance of the overall RF section has been evaluated by driving the LNA with a CW RF source and by reading the DC voltages at the output of the 4 power detectors. As the output impedance of the LNA board and the input impedance of the post-amplifier and detector board differ from 50Ω, both the insertion loss and the central frequency of the filter needed fine tuning. Moreover, the termination impedances could affect system stability. However, re-optimization of the filter bank allowed to overcome this problem without losing filter selectivity.

To perform a full evaluation of the overall wake-up receiver we run several sets of in-field experiments. In these tests we determined the wake-up characteristics of the WuR prototype by considering two off-the-shelf motes acting as transmitters and receivers of wake-up signals. The receiver node is a TelosB mote interfaced with the WuR prototype. We designed and implemented a firmware for the wake-up receiver, including an UART control protocol for communication between the mote and the WuR, which is required for dynamic address reprogramming. A dedicated TinyOS application was developed to keep statistics about the wake-up signals received, and to test dynamic address reprogramming. The transmitter node is a MTM-CM3300 mote equipped with a TI CC2420 with external SMA antenna and additional power amplifier. WuR sequences are sent by a sender node by sequentially transmitting an unmodulated carrier for a given period of time, called the symbol duration, on each of the IEEE 802.15.4 channels that compose the target wake-up address. Note that no additional hardware is necessary to send WuR addresses. Rather, each symbol of the wake-up sequence is transmitted by the node by putting its TI CC2420 transceiver into unmodulated carrier test mode [14] for the required period of time. We modified the CC2420 driver of TinyOS to add support for sending wake-up sequences with customizable length and symbol duration. In the reported set of experiments, different wake-up sequences are sent by the transmitter node over two distinct IEEE 802.15.4 channels: channels 12 (2410 MHz) and 26 (2480 MHz). Each wake-up address is defined as a sequence of 8 symbols. For brevity, wake-up addresses are encoded so that 0 represents a signal sent on channel 12, and 1 represents a symbol sent on channel 26. Thus, the address `00001111` represents the sequence: CH12 CH12 CH12 CH12 CH26 CH26 CH26 CH26. To test address recognition, the transmitter nodes sends four different wake-up sequences: `00000000`, `11111111`, `00001111` and `01010101`. We deployed the mote outdoors, in Villa Borghese gardens in Rome, and varied the distance between the transmitter and the receiver between 10 and 120 meters For each distance, each wake-up sequence of the four considered is sent 100 times, while wake-up statistics are recorded at the receiver node. Figure 5(a) shows the results of these experiments, reporting the fraction of wake-up addresses successfully received (shown in green with crossed fill pattern) with respect to the total number of wake-up addresses sent, for each distance between the transmitter and the receiver. The fraction of false negatives

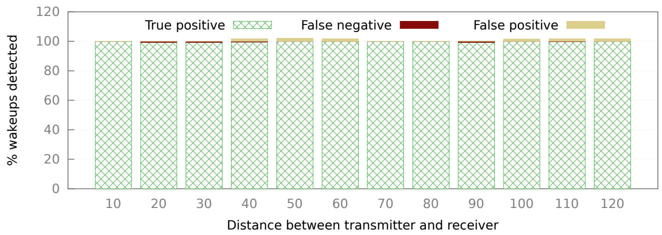

(a)

Fig. 5. WuR design validation: summary of representative results of in-field experiments.

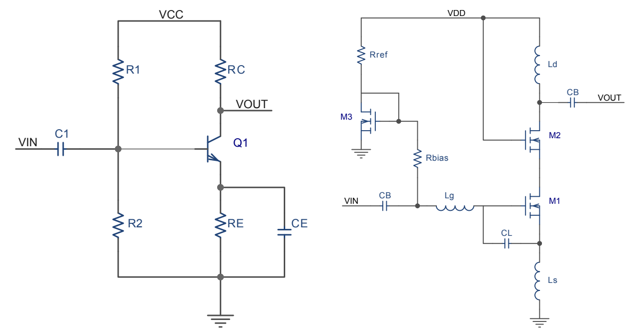

(a) (b)

Fig. 6. Design and characterization of the most critical blocks of the WuR in integrated technology: (a) emitter follower biasing configuration, and (b) circuit schematic of the low noise amplifier.

and false positives is represented by dark red and light gold bar, respectively. On average, both false negative and false positives are lower than 1%.

B. Power consumption of the wake-up receiver

The design and characterization of the radio receiver has been validated in terms of obtainable power consumption, by designing the most critical blocks, i.e. the RF LNA and the filter bank, in integrated technology.

1) Filter design: A filter prototype has been realized at schematic level with IHP SGB25V technology. We investigated its performance in order to define the best configuration for the filters bank application. The reference goals in the design are: power consumption, bandwidth at -3 dB, P_{-1dB} compression point, and stability versus power and temperature. Active Inductor (AI) compensating network and polarization circuit topologies of the active device have been deeply investigated in order to find the best results. To verify the feasibility of an active inductor designed with the proposed approach with an IHP process, we realized a first solution using the same polarization network of the discrete filter and without taking into account power consumption as a requirement. The obtained equivalent inductance value is 136 nH with a real part of less than 0.1 Ohm. This result gives the possibility to design a filter with high quality factor and 3dB bandwidth less than 10MHz. The dynamic range of the circuit is less than 50dB, but it can be significantly improved by using high impedance values at the filter terminations. An important problem related to this configuration is the stability of the filter response versus temperature variations, which significantly affects its performance. The proposed solution for temperature compensation uses a simple polarization network realized with an emitter follower configuration. A simplified schematic of the proposed structure is shown in Figure 6(a). This solution tends to be relatively immune to temperature instability, because as the current in the collector increases, the voltage across the emitter resistor also increases. This acts as negative feedback, since it reduces the voltage difference between base and emitter (VBE). This counteracts the transistor's tendency to conduct more current, thus maintaining thermal stability. Due to the capacitance CE, RE here only remedies the temperature instability and does not enter into the AC gain equation. At AC, the signal passes through CE and not through RE as CE will have much lower impedance at signal frequencies than RE (Fig. 6(a)). This configuration has been implemented with a DC voltage of 1.2 V and current of 140 μA, for a power consumption of 168 μW.

2) Low-noise amplifier (LNA) design: The proposed LNA has been designed following the topology shown in Figure 6(b). It is an inductor-source-degenerated low noise amplifier suitable for wireless applications. In this circuit the transistor M3 is a biasing transistor that forms a current mirror with M1. The transistor M2 isolates tuned input from output to increase reverse isolation, also reducing the effect of Miller capacitance C_{gd}. The inductance L_g sets the resonant frequency f_0=2.4 GHz together with C_L, while L_s matches input impedance. L_d increases the gain and also works as band pass filter with the output capacitance. The resistor R_{BIAS} must be dimensioned so as that its equivalent current noise would be small enough to be ignored, while C_B are blocking capacitors chosen to have negligible reactance at f_0=2.4 GHz. In a power-driven design approach, the LNA noise figure is higher than the global minimum NF, and minimized versus power constraints and power dissipation. The proposed circuit topology is a tradeoff between performance and power consumption, working with a bias current of 1mA with a voltage supply of 1.2V, and having the following specifications: Power gain at f_0 = 16.3 dB; Noise figure = 3.2 dB; -1dB Compression Point = -26.8dBm; OIP3 = 0dB.

C. Communication protocol evaluation

In the following we show the results of a comparative performance evaluation of our proposed wake-up-based protocols with traditional interest dissemination and converge casting protocols. Simulations have been performed using GreenCastalia [15], an open-source extension we have developed for the popular Castalia simulator [16] that allows to model networks of embedded devices with energy-harvesting capabilities. We implemented both FLOOD-WUP and GREEN-WUP in GreenCastalia, together with a novel dedicate WuR module that captures the details of the proposed wake-up receiver. The channel data rate was set to 250 Kbps. Communication hardware of each node includes a TI CC2420 transceiver and, when running WUP protocol, a wake-up receiver. Unless otherwise specified, the default Castalia settings were used for channel and radio models. Packet collisions are determined by using the additive interference model, according to which simultaneous transmissions from multiple nodes are calculated as interference by linearly adding their effect at the receiver. Packet reception probability for each link is computed based on SINR, packet size, and modulation type.

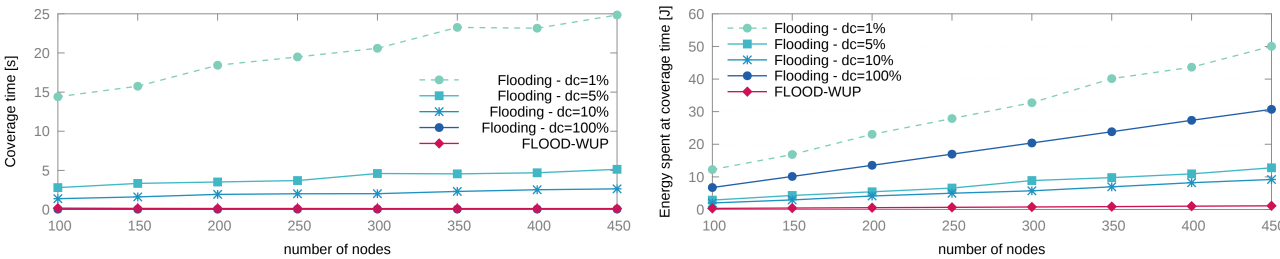

(a) (b)

Fig. 7. Performance comparison of flooding and FLOOD-WUP during interest dissemination: (a) time needed to reach all nodes; and (b) total energy consumed by the network.

The energy model we used is that of TelosB. Each node operates in one of three power modes, depending on the state of its MCU and transceiver: transmit (TX), receive (RX), and sleep (SLEEP). Mimicking the TinyOS implementation, the MCU is put into a low-power sleep state (LPM3) when the radio is off and the node is not busy with other activities (e.g., sampling). Assuming the node is powered at 3V, and considering nominal datasheet values [17], we set the power consumption of the node in TX, RX and SLEEP state to 58.5 mW, 65.4 mW, and 15.30 μW, respectively. Transition delays between states are set according to CC2420 and MSP430 datasheets [14], [18]. When simulating WUP protocols, we assume that each node in the network is equipped with a WuR receiver using 2 channels, which is kept active when the main transceiver of the node is down. According to experimental and power scaling results, the sensitivity of the wake-up receiver is set to -83 dBm, while its power consumption is of 1.62 mW. All results reported in the following have been obtained by averaging results over 50 simulation runs.

1) Interest dissemination: In the following we show the results of a comparative performance evaluation of flooding and FLOOD-WUP. Our implementation of the flooding protocol is based on the Castalia TunableMAC module, an highly-configurable MAC with duty cycling support. In this set of experiments, N sensors are distributed uniformly at random in a square area of side 100 meters, where N varies between 100 and 450 nodes. The sink, which is deployed at the center of the field, sends commands to the rest of the network. Nodes using the flooding protocol follow asynchronous wake-up schedules according to a given duty cycle d. We considered nominal duty cycles $d \in \{0.01, 0.05, 0.1, 1.0\}$.

Figure 7 shows the results of our performance evaluation in terms of time needed to reach all nodes and total energy consumed by the network during interest dissemination. In all tests, the percentage of nodes reached during the interest dissemination is 100%. As shown in the figure, the performance of FLOOD-WUP in terms of coverage time are very close to those of Flooding with 100% duty cycle: on average, the coverage time achived by FLOOD-WUP is approximately 80ms higher than that of Flooding with 100% duty cycle. However, the energy spent by FLOOD-WUP for interest dissemination is, on average, less than 4% of that consumed by nodes running Flooding with 100% duty cycle. Even in case of lower duty cycles, e.g., d=10%, FLOOD-WUP requires on average 8 time less energy than Flooding, while achieving a coverage time that is up to 24 time less than that of Flooding.

2) Converge casting: In this section we compare the performance of GREEN-WUP against that of the Collection Tree Protocol [19], the de-facto routing protocol for WSNs. We implemented CTP in GreenCastalia based on the CTP implementation publicly available at [20], which we extended to support low-power listening. LPL parameters are set accordingly to the TinyOS 2.1 implementation. Nodes using CTP follow asynchronous wake-up schedules, performing periodic receive checks every l ms, $l \in \{0, 50, 100, 250, 500\}$ ($l = 0$ representing the case in which LPL is not used). In this set of experiments, the network is composed by nodes powered by an harvesting subsystems consisting of a photovoltaic cell and a rechargeable battery. Power harvested by the nodes over time is simulated based on traces of solar availability we collected by interfacing TelosB motes with photovoltaic cells [21]. 100 nodes are distributed uniformly at random in a square area of side 100 meters. The sink, which is deployed at the center of the field, is always active. The network runs a periodic data collection application. Every 60 seconds, a given percentage p of source nodes, $p \in \{1, 3, 5, 10, 15\}$, chosen randomly, perform a sensor reading and generate a packet that must be delivered to the sink. The rest of the nodes in the networks are relays. Each simulation run lasts 700 seconds. The default size of data packets is 128 bytes.

Figure 8 shows the results of our performance evaluation in terms of packet delivery ratio, average latency, and total energy consumed by the network (excluding the sink node). CTP with no LPL and GREEN-WUP achieve very similar performance in terms of packet delivery ratio, as the PDR of GREEN-WUP is, on average, less than 0.2% smaller than that of CTP with no LPL. Overall, GREEN-WUP shows very good performance in terms of reliability, obtaining a packet delivery ratio that is always greater than 98.5%, and around 99.5% on average. As for latency, CTP with no LPL predictably achieves the best performance, having an average end-to-end latency that is approximately 100ms smaller than that of GREEN-WUP. However, its energy consumption is up to 33 time greater than that of GREEN-WUP. With respect to CTP with LPL=50ms, GREEN-WUP obtains an end-to-end latency that, on average, is 50ms higher, but its energy consumption is up to a factor of 4.26 smaller. For $l \in \{100, 250, 500\}$, GREEN-WUP always outperforms CTP+LPL in terms of both packet delivery ratio and packet latency, while significantly reducing energy consumption. For example, GREEN-WUP obtains an end-to-end data latency that is up to 16 times smaller than that of CTP+LPL=500ms, consuming at the same time approximately 45% less energy. Despite the use of LPL with

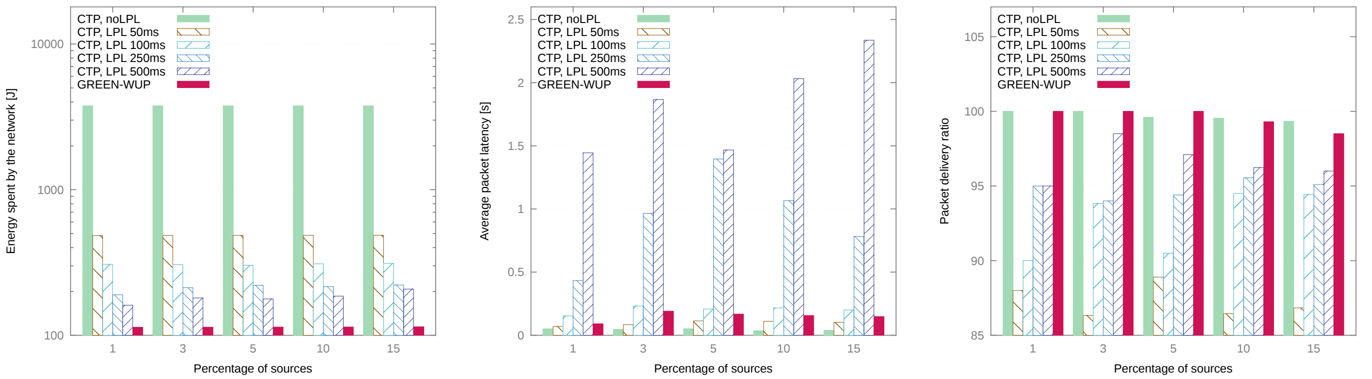

(a) (b) (c)

Fig. 8. Performance comparison of CTP and GREEN-WUP for increasing percentage of source nodes: (a) total energy consumed by the network [note the y axis is in logscale]; (b) average packet latency; and (c) packet delivery ratio.

higher value of l allows to reduce energy consumption, it has detrimental effects on PDR and end-to-end latency.

VI. Conclusion

In this paper a novel topology of wake-up radio with addressing capability has been proposed. The considered architecture, based on an innovative topology of band-pass filter, has been evaluated with both an implementation with COTS components and simulations in microelectronic technology of the most critical sections to prove power consumption scaling. Simulation of a single filter in BiCMOS integrated technology has shown a power consumption of 168μW. Taking advantage of the proposed received architecture, we have presented two novel protocols, named FLOOD-WUP and GREEN-WUP, that exploits selective wake-ups and dynamic address assignment to optimize system performance. Comparison against traditional WSN protocols, such as Flooding and CTP+LPL, have shown that FLOOD-WUP and GREEN-WUP outperform existing approaches in addressing the latency vs. energy consumption tradeoffs, achieving superior performance and significant energy saving with respect to duty-cycle-based solutions.

Acknowledgment This research has received funding from the FP7 project GENESI, grant agreement n. 257916. D. Spenza also gratefully acknowledges support from a Google Europe Doctoral Fellowship.

References

[1] W. Ye, J. Heidemann, and D. Estrin, "An energy-efficient MAC protocol for wireless sensor networks," in *Proceedings of IEEE INFOCOM 2002*, vol. 3, 2002, pp. 1567–1576 vol.3.

[2] S. Basagni, M. Y. Naderi, C. Petrioli, and D. Spenza, "Wireless Sensor Networks with Energy Harvesting," in *Mobile Ad Hoc Networking: The Cutting Edge Directions*. John Wiley and Sons, Inc., March 2013, ch. 20, pp. 701–736.

[3] L. Gu and J. Stankovic, "Radio-triggered wake-up for wireless sensor networks," *Real-Time Systems*, vol. 29, no. 2-3, pp. 157–182, 2005.

[4] L. Chen, S. Cool, H. Ba, W. Heinzelman, I. Demirkol, U. Muncuk, K. Chowdhury, and S. Basagni, "Range extension of passive wake-up radio systems through energy harvesting," in *Proceedings of IEEE ICC 2013*, 2013, pp. 1549–1554.

[5] S. Drago, D. Leenaerts, F. Sebastiano, L. Breems, K. Makinwa, and B. Nauta, "A 2.4ghz 830pj/bit duty-cycled wake-up receiver with -82dbm sensitivity for crystal-less wireless sensor nodes," in *Proceedings of IEEE ISSCC 2010*, Feb 2010, pp. 224–225.

[6] H. Milosiu, F. Oehler, M. Eppel, D. Fruhsorger, S. Lensing, G. Popken, and T. Thones, "A 3-μW 868-MHz wake-up receiver with -83 dBm sensitivity and scalable data rate," in *Proceedings of ESSCIRC 2013*, Sept 2013, pp. 387–390.

[7] J. Ansari, D. Pankin, and P. Mähönen, "Radio-triggered Wake-ups with Addressing Capabilities for Extremely Low Power Sensor Network Applications," *IJWIN*, vol. 16, no. 3, pp. 118–130, 2009.

[8] G. Lu, D. De, M. Xu, W.-Z. Song, and J. Cao, "Telosw: Enabling ultra-low power wake-on sensor network," in *Proceedings of INSS 2010*, June 2010, pp. 211–218.

[9] G. Gamm and L. Reindl, "Range extension for wireless wake-up receivers," in *9th International Multi-Conference on Systems, Signals and Devices, IEEE SSD 2012*, March 2012, pp. 1–4.

[10] V. Jelicic, M. Magno, D. Brunelli, V. Bilas, and L. Benini, "Analytic comparison of wake-up receivers for wsns and benefits over the wake-on radio scheme," in *Proceedings of ACM PM2HW2N 2012*, 2012, pp. 99–106.

[11] J. Oller, I. Demirkol, J. Casademont, J. Paradells, G. U. Gamm, and L. Reindl, "Performance evaluation and comparative analysis of sub-carrier modulation wake-up radio systems for energy-efficient wireless sensor networks," *Sensors*, vol. 14, no. 1, pp. 22–51, 2013.

[12] V. Stornelli, G. Ferri, G. Leuzzi, and A. De Marcellis, "A tunable 0.5-1.3 ghz cmos 2nd order bandpass filter with 50ω input-output impedance matching," in *Proceedings of IEEE ISCAS 2006*, May 2006, pp. 4 pp.–.

[13] S. Del Re, G. Leuzzi, and V. Stornelli, "A new approach to the design of high dynamic range tunable active inductors," in *Proceedings of INMMIC 2008*, Nov 2008, pp. 25–28.

[14] Texas Instruments, "CC2420 Datasheet, SWRS041c," Available from the TI website, p. 85.

[15] D. Benedetti, C. Petrioli, and D. Spenza, "GreenCastalia: An Energy-harvesting-enabled Framework for the Castalia Simulator," in *Proceedings of ACM ENSSys 2013*, 2013, pp. 7:1–7:6. [Online]. Available: http://senseslab.di.uniroma1.it/greencastalia

[16] A. Boulis, "Castalia: Revealing Pitfalls in Designing Distributed Algorithms in WSN," in *Proceedings of ACM SenSys 2007*, 2007, pp. 407–408.

[17] Crossbow Technology, "TelosB mote platform datasheet," 2004, document Part Number: 6020-0094-01 Rev B.

[18] Texas Instruments, "MSP430F161x Datasheet," Available from the TI website, p. 77.

[19] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis, "Collection tree protocol," in *Proceedings of ACM SenSys 2009*, 2009, pp. 1–14.

[20] U. Colesanti and S. Santini, "The collection tree protocol for the castalia wireless sensor networks simulator," Department of Computer Science, ETH Zurich, Tech. Rep. 729, June 2011.

[21] T. La Porta, C. Petrioli, C. Phillips, and D. Spenza, "Sensor-mission assignment in rechargeable wireless sensor networks," *ACM Transactions on Sensor Networks*, To appear.