

REACH²-Mote: A Range-Extending Passive Wake-Up Wireless Sensor Node

LI CHEN, JEREMY WARNER, PAK LAM YUNG, DAWEI ZHOU,
and WENDI HEINZELMAN, University of Rochester
ILKER DEMIRKOL, Universitat Politècnica de Catalunya
UFUK MUNCUK, KAUSHIK CHOWDHURY, and STEFANO BASAGNI,
Northeastern University

A wireless sensor network that employs passive radio wake-up of the sensor nodes can reduce the energy cost for unnecessary idle listening and communication overhead, extending the network lifetime. A passive wake-up radio is powered by the electromagnetic waves transmitted by a wake-up transmitter rather than a battery on the sensor node. However, this method of powering the wake-up radio results in a short wake-up range, which limits the performance of a passive wake-up radio sensor network. In this article, we describe our design of a passive wake-up radio sensor node—REACH²-Mote—using a high-efficiency, energy-harvesting module and a very low power wake-up circuit to achieve an extended wake-up range. We implemented REACH²-Mote in hardware and performed field tests to characterize its performance. The experimental results show that REACH²-Mote can achieve a wake-up range of 44 feet. We also modeled REACH²-Mote and evaluated its performance through simulations, comparing its performance to that of another passive wake-up radio approach, an active wake-up radio approach, and a conventional duty cycling approach. The simulation results show that REACH²-Mote can significantly extend the network lifetime while achieving high packet delivery rate and low latency.

Categories and Subject Descriptors: C.2.3 [**Computer-Communication Networks**]: Network Operations

General Terms: Design, Performance

Additional Key Words and Phrases: Wireless sensor networks, passive wake-up radio, range extension

ACM Reference Format:

Li Chen, Jeremy Warner, Pak Lam Yung, Dawei Zhou, Wendi Heinzelman, Ilker Demirkol, Ufuk Muncuk, Kaushik Chowdhury, and Stefano Basagni. 2015. REACH²-Mote: A range-extending passive wake-up wireless sensor node. *ACM Trans. Sen. Netw.* 11, 4, Article 64 (December 2015), 33 pages.
DOI: <http://dx.doi.org/10.1145/2829954>

This research was funded in part by the National Science Foundation under research grant CNS-1143662 and in part by the Spanish government, MINECO, through projects TEC2012-32531, RYC-2013-13029, and FEDER.

Authors' addresses: L. Chen, J. Warner, P. L. Yung, D. Zhou, and W. Heinzelman, University of Rochester, Department of Electrical and Computer Engineering, Rochester, NY 14627 USA; emails: li.chen.83@gmail.com, jeremy.warner@rochester.edu, pyung2@u.rochester.edu, dzhou3@hse.rochester.edu, wendi.heinzelman@rochester.edu; I. Demirkol, Universitat Politècnica de Catalunya, Barcelona, Catalunya, Spain, 08034; email: ilker.demirkol@entel.upc.edu; U. Muncuk, K. Chowdhury, and S. Basagni, Northeastern University, Department of Electrical and Computer Engineering, Boston, MA 02115 USA; emails: umuncuk@coe.neu.edu, {krc,basagni}@ece.neu.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2015 ACM 1550-4859/2015/12-ART64 \$15.00

DOI: <http://dx.doi.org/10.1145/2829954>

1. INTRODUCTION

Wireless sensor networks (WSNs) are composed of several sensor nodes that can sense the physical environment (e.g., temperature, air quality, sound, pressure), process the sensed data, and send the processed data to other nodes or to the data sink(s) in the network. There are many potential applications for WSNs, including smart grid monitoring, emergency response, military surveillance, home security, and environment monitoring. As typically the sensor nodes are powered by batteries, WSNs are highly energy constrained. Additionally, in some cases, the batteries attached to the sensor nodes are difficult or even impossible to replace. Thus, minimizing the energy dissipation of a sensor node is a key problem in WSN research.

Duty cycling, where the sensor node is periodically set to the sleep mode, is one of the most commonly used methods to reduce the energy dissipation of a sensor node. As communication between two nodes can only be achieved when both the transmitter and the receiver nodes are awake, the duty cycles of all nodes must be time synchronized so the nodes all wake up at the same time; otherwise, idle listening is required until both the transmitter and receiver are awake simultaneously. However, both time synchronization and idle listening increase the complexity of the MAC protocol and waste additional energy. Furthermore, to reduce the energy dissipation of the nodes, the sensor nodes tend to be kept in the sleep mode for the majority of the time, which increases the delay for packet delivery. In the case of a mobile sink, the sensor node may be in the sleep mode when the sink comes by to collect data, and thus the sink may miss collecting that node's data. Thus, duty cycling may not be suitable for some delay-sensitive applications.

Using a wake-up radio—a low-power, secondary radio that is only used to wake up the primary radio for communication—is another solution for prolonging the lifetime of a WSN. Using a wake-up radio, the sensor node is only woken up when communication is necessary. The cost for this approach is the additional hardware needed on the devices, including a wake-up radio receiver (WuRx) and a wake-up radio transmitter (WuTx). Each sensor node with a WuRx has two working modes: sleeping mode and active mode. Most of the time, the sensors are kept in an ultra low power sleep mode, where they cannot communicate with other nodes nor perform any computation. The sensor node may wake up periodically to sense the environment and go back to sleep after the data is collected and stored in local memory. Only when a surrounding node's WuTx sends a trigger signal to start data communication and the WuRx receives this signal will the WuRx trigger the sensor node to enter the active mode, at which point it can communicate with other nodes in the network.

Two classes of wake-up radio devices have been developed: active wake-up radios and passive wake-up radios. An active WuRx requires a power supply, which commonly is the battery of the sensor node. Most active wake-up receivers provide good performance in terms of wake-up delay and wake-up distance. On the other hand, passive wake-up radio devices are powered by energy harvested from the WuTx signals (and hence do not require any energy from the sensor node's battery), which reduces the energy consumption of the sensor node but results in a shorter wake-up range than the active wake-up approach. As passive WuRxs utilize the energy harvested from the RF signals sent by the WuTx, this approach extends the lifetime of the sensor network compared to using active wake-up radios and using duty cycling.

However, there are several challenges for passive wake-up radio sensor networks. First, due to the limitations and efficiency losses in the energy harvesting process, passive wake-up radio sensor nodes operate over a shorter communication range and present longer wake-up delay than active wake-up radios. Additionally, the performance of a passive WuRx may be affected by environmental conditions, such as heavy rain, which may decrease the energy received by the WuRx, possibly making some

sensor nodes inaccessible. Furthermore, to achieve a reasonable wake-up distance, the WuTx needs to be designed to have a high energy transmission efficiency. As a result, it is difficult to build a multihop WSN where each node is equipped with both a WuTx and a passive WuRx.

To address some of these issues, in Chen et al. [2013], we introduced the REACH-Mote (Range EnhAnCing energy Harvester-Mote) and evaluated its performance through field tests. REACH-Mote is composed of a highly efficient energy harvester module [Nintanavongsa et al. 2012] and an ultra low power wake-up circuit to achieve a long-range passive wake-up. In this article, we enhance the design of the REACH-Mote to create REACH²-Mote, with an improved wake-up range achieved by applying an improved energy harvesting module and a supply voltage regulator.

We perform a thorough evaluation of the performance of REACH²-Mote through both field tests of the hardware and simulations. We compare the performance of REACH²-Mote with that of REACH-Mote and another passive wake-up radio called *WISP-Mote* [Ba et al. 2010]. The field test results show that REACH²-Mote can achieve an extended wake-up range of 44 feet, which represents a 19% increase compared to the wake-up range of REACH-Mote and a 220% increase compared to the wake-up range of WISP-Mote. Based on the physical characterization of REACH²-Mote and WISP-Mote, we developed a simulation model of the performance of REACH²-Mote and WISP-Mote. Additionally, we model a conventional duty cycling approach and an active wake-up radio approach [Pletcher et al. 2009]. Using these models, we perform simulations under several different network scenarios with a mobile sink (e.g., a data mule [Anastasi et al. 2008]) that traverses the network to collect data from the sensor nodes. The simulation results show that REACH²-Mote can significantly extend the network lifetime while achieving a high packet delivery rate (PDR) and low latency for the scenarios we tested.

The remainder of this article is organized as follows. In Section 2, we present a survey of related work. The description of the hardware design of the first-generation REACH-Mote is provided in Section 3, and the description of the hardware design of the second-generation REACH²-Mote is provided in Section 4. Section 5 presents results from field experiments using three passive wake-up radio designs (WISP-Mote, REACH-Mote, and REACH²-Mote). Simulation results under different network scenarios using REACH²-Mote, WISP-Mote, an active wake-up approach, and a duty-cycling approach are provided in Section 6. We compare the design of different passive wake-up sensor nodes in Section 7. Conclusions are drawn in Section 8, and future work is discussed in Section 9.

2. RELATED WORK

Reducing the energy dissipation of the sensor nodes is an important goal in the design of WSNs. Duty cycling is one approach to reducing energy dissipation, where the radio is periodically turned off to save energy that would be wasted on idle listening. However, as communication can only occur when the transmitter and receiver nodes are both awake, the duty cycles must be synchronized; otherwise, the nodes waste energy in idle listening, waiting until both nodes are awake.

Both synchronized protocols and asynchronous protocols have been developed for conventional WSNs to support duty cycling. Synchronized protocols such as S-MAC [Ye et al. 2002] and T-MAC [Van Dam and Langendoen 2003] negotiate a schedule between sensor nodes so that the nodes can wake up at the same time to communicate. Asynchronous protocols such as B-MAC [Polastre et al. 2004], WiseMAC [El-Hoiydi et al. 2003], and X-MAC [Buettner et al. 2006], also known as low-power listening protocols, apply preamble sampling to establish communication between the sender and the receiver. Both synchronized protocols and asynchronous protocols need to wait

until both nodes are awake before communication can begin, which wastes energy from the battery and increases the transmission delay. Increasing the wake-up/sleep ratio can improve the latency performance at the expense of wasting more energy due to unnecessary wake-ups. Thus, it is difficult for duty cycling protocols to achieve both energy efficiency and low latency.

Active wake-up radios utilize low-power wake-up circuits for WuRxS, which are powered by the batteries of sensor nodes. Thus, the energy consumption of these wake-up circuits are critical for determining the performance of the active wake-up sensor network. Van der Doorn et al. [2009] proposed a $96\mu\text{W}$ wake-up circuit and Le-Huy and Roy [2010] developed a WuRx circuit that consumes $17.8\mu\text{W}$ to achieve a low-power wake-up. The energy costs of active WuRxS are decreasing continuously. Recently, Spenza et al. [2015] presented the architecture and applications of a receiver consuming less than $1.3\mu\text{W}$ and -55dBm sensitivity. The wake-up circuits proposed in Ansari et al. [2009] and Marinkovic and Popovici [2011] only consume $2.4\mu\text{W}$ and $0.27\mu\text{W}$ by using integrated circuits, respectively. However, as all of these active wake-up receivers only achieve a wake-up sensitivity of -50dBm to -60dBm , compared to a -95dBm sensitivity for conventional sensor nodes, the wake-up range of these active wake-up circuits is much shorter than the communication range of sensor nodes. Pletcher et al. [2009] proposed an active wake-up receiver that achieves a -72dBm sensitivity with an energy cost of $52\mu\text{W}$, and Petrioli et al. [2014] proposed a discrete components wake-up receiver with -85dBm sensitivity with 1.2mW energy consumption. These two approaches provide a decent wake-up range for sensor network applications. In this work, we will compare our passive wake-up approach to Pletcher's work through simulations, as it offers a good range and low energy consumption.

Energy harvesting can be used to extend a wireless sensor node's lifetime without increasing the device's battery capacity. Energy harvesters capture energy from ambient vibration, wind, heat, light, or electromagnetic radiation and convert this into electrical energy. This energy can either be used to power an ultra low power MCU or it can be stored in a supercapacitor or battery. Supercapacitors are used when the application needs to provide large energy spikes. Batteries leak less energy and are therefore used when the device needs to provide a steady flow of energy [Basagni et al. 2013]. The generated energy is usually very small and highly dependent upon the size and efficiency of the generator; thus, a good energy harvester system must have very low internal loss of energy and good storage. For example, AmbiMax is an energy harvesting circuit and a supercapacitor-based energy storage system for wireless sensor nodes [Park and Chou 2006]. Moreover, AmbiMax is modular and enables composition of multiple energy harvesting sources, including solar, wind, thermal, and vibration.

Wireless Identification and Sensing Platform (WISP) is a research project of Intel Research Seattle assisted by the University of Washington [Sample et al. 2008]. WISP is a battery-less device that harvests power from a standard off-the-shelf RFID reader and uses this to respond to the reader. The harvested energy operates a 16-bit ultra low power MSP430 microcontroller that can perform a variety of computing tasks, such as sampling sensors and reporting this data back to the RFID reader [wisp.wikispaces.com 2010]. WISP is an open source, open architecture EPC Class 1 Generation 2 RFID tag that includes a light sensor, a temperature sensor, a strain gauge, and an accelerometer [Tapia et al. 2007]. WISP is one of many implementations of passive devices that use backscatter communication between an RFID reader and a WISP node. Liu et al. [2013] and Parks et al. [2014] proposed radio nodes for ambient backscatter communication, which transmit a signal by reflecting TV radio waves. Zhang and Ganesan [2014] implemented a bit-by-bit backscatter communication in severe energy harvesting environments. Gummesson et al. [2010] analyze the energy performance of the WISP with a hybrid energy harvester. As the backscattered signal strength is weak compared




Fig. 1. REACH-Mote main system components.

to conventional communication methods, it is very hard to build a long-range multihop backscatter network.

Passive wake-up radios, which are the focus of this article, do not rely on the nodes' battery power supplies while awaiting a wake-up signal from the wake-up transmitter. Sensor nodes that employ passive wake-up receivers tend to have longer lifetimes but shorter wake-up range compared to sensor nodes that employ active wake-up receivers. There are a few existing approaches in the literature for passive wake-up radios. Gu and Stankovic [2005] proposed a passive radio-triggered wake-up for WSNs. However, the authors did not provide an implementation of their design; they only evaluated its performance through simulation. In our previous work [Chen et al. 2013; Ba et al. 2010], we proposed three single-hop passive wake-up motes: WISP-Mote, EH-WISP-Mote, and REACH-Mote. Among these implementations, WISP-Mote is our first-version passive wake-up radio device, which is a combination of a WISP and a Tmote Sky sensor node [Ba et al. 2010]. Whenever the WISP harvests enough energy from the transmitter radio, it sends a pulse to wake up Tmote Sky from the sleep state. WISP-Mote can be awakened by an Impinj RFID reader [Impinj 2002] at a maximum distance of approximately 13 feet. Moreover, simulations show the potential advantages of WISP-Mote over duty cycling in terms of delay, collision, overhead, energy efficiency, and protocol complexity [Ba et al. 2013]. Based on the design of WISP-Mote, we developed EH-WISP-Mote, which uses a parallel harvesting circuit to extend the wake-up range. Implementation results show that EH-WISP-Mote can reach 17 feet for the wake-up range at a height of 1 foot above the ground, which is 4 feet further than WISP-Mote's maximum wake-up range, representing a 20% improvement in the maximum wake-up range performance [Chen et al. 2013]. All of these represent a promising approach for passive wake-up of the sensor nodes.

3. REACH-MOTE

A passive WuRx does not use any energy from the sensor node's battery; instead, it utilizes the energy harvested from the signal sent by the WuTx. Thus, to achieve a long-range passive wake-up, the WuRx must include a high-efficiency energy harvester. In addition, the wake-up circuit that triggers the MCU of the sensor node should operate using as little energy as possible to further extend the wake-up range. Thus, an efficient passive WuRx should be composed of a high-efficiency energy harvester, a low-power wake-up trigger generator, and a wireless sensor node. Using these components, we created a node called *REACH-Mote*, as shown in Figure 1 [Chen et al. 2013]. REACH-Mote operates as follows:

- By default, REACH-Mote is in sleep mode (i.e., the MCU on Tmote Sky, which is an MSP430 F1611, is put to LPM3 sleep mode [MSP430 2012] and the radio on Tmote Sky is turned off.

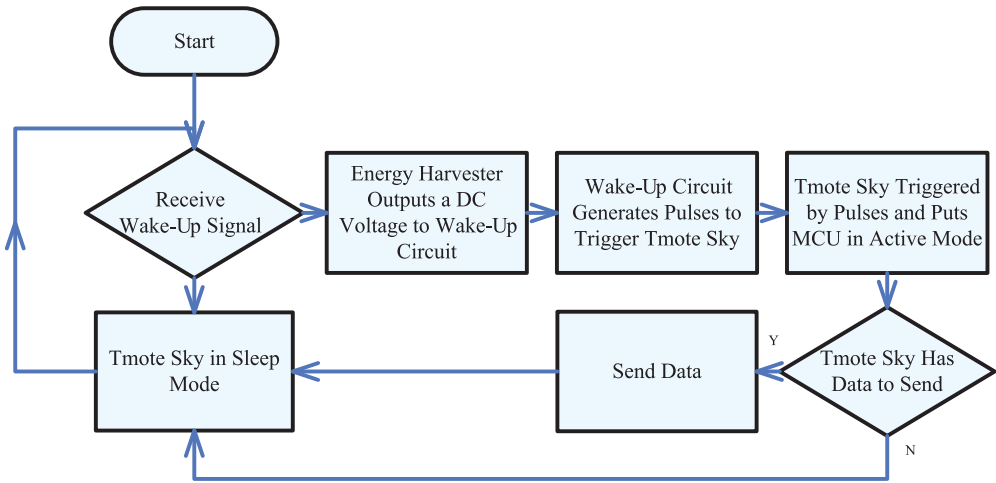


Fig. 2. REACH-Mote operation flow chart.

- When a wake-up signal is sent by the WuTx of a nearby mote or base station, the energy harvesting circuit receives the energy and outputs a DC voltage.
- The wake-up circuit generates a pulse once the DC voltage is higher than 1.5V, and this will trigger the sensor mote.
- The trigger forces the MCU on the sensor mote to be woken up, then the MCU turns on the radio (i.e., the CC2420 [CC2420 2012] on Tmote Sky).
- After waking up, if the mote has data to send, data transfer commences.
- If the mote does not have data to send, or after the data transmission is complete, the mote goes directly back to sleep mode (i.e., the MCU is set to LPM3 and the radio is turned off).

The flow chart of the REACH-Mote operation is shown in Figure 2.

3.1. Energy Harvesting Circuit Design

The RF energy harvesting circuit enhances the wake-up ability of REACH-Mote, as a more efficient energy harvester increases the wake-up distance. In this section, we describe the general design of the energy harvesting circuit and interfacing principles, and motivate the choice of specific circuit components.

3.1.1. Selection of Circuit Components. The overall aim of our design is to maximize the energy conversion from the front-end antenna to the sensor node. To achieve this, as shown in Figure 3, we carefully tune a matching circuit to balance the input impedance seen from the antenna side with the circuit load (i.e., the WuRx and Tmote Sky combination), as well as use a voltage rectifier that also functions as a multiplier. The multiplier is based on the classical Dickson's voltage multiplier circuit (Figure 4), which has several stages connected in parallel, each stage being a series combination of a diode and a capacitor. The advantage here is that because the capacitors appear in parallel with respect to each other, the effective circuit impedance is reduced. Hence, this makes the task of matching the antenna side to the load side simpler.

As the peak voltage of the AC signal obtained at the antenna is generally much smaller than the diode threshold [Yan et al. 2005], diodes with the lowest possible turn-on voltage are preferable. Moreover, since the energy harvesting circuit operates in the high megahertz range, diodes with a very fast switching time need to be used. Schottky diodes use a metal-semiconductor junction instead of a semiconductor-semiconductor

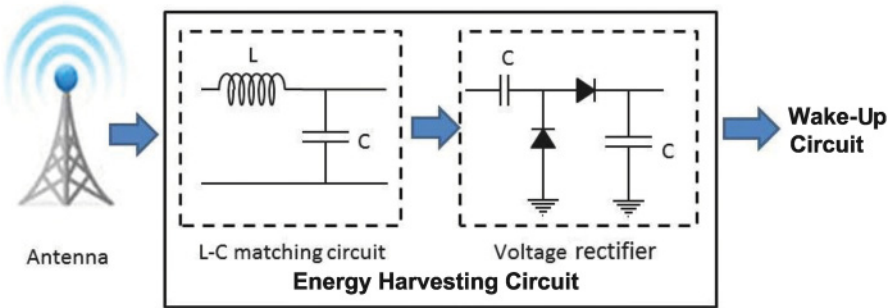


Fig. 3. Architectural view of the REACH-Mote circuit and connections.

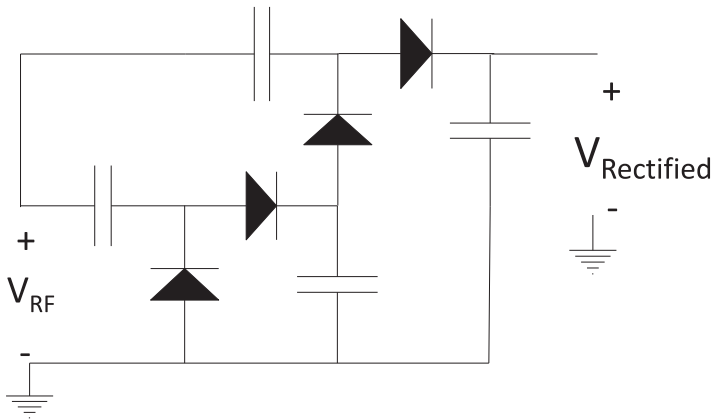


Fig. 4. Dickson diode-based multiplier.

junction. This allows the junction to operate much faster and gives a forward voltage drop of as low as 0.15V. We employ diodes from Avago Technologies' HSMS-2852, which has a turn-on voltage of 150mV, measured at 0.1mA, because this specific diode is suitable for operating in the low power region, typically considered as the range of power between -20dBm and 0dBm .

The selection of the number of multiplier stages has a major influence on the output voltage of the energy harvesting circuit [Nintanavongsa et al. 2012]. Although the output voltage is directly proportional to the number of stages used in the energy harvesting circuit, it also reduces progressively the current drawn by the load, which in turn impacts the overall charging time. We set the number of stages to 10, as this ensures sufficient output voltage of the circuit to drive REACH-Mote at 915MHz.

3.1.2. Optimization Framework and Fabrication. The selection of the precise values for the matching circuit is undertaken through an optimization framework, where a fixed input RF power is injected via the Agilent N5181 MXG RF signal generator, and the resulting changes in the output voltage values are measured through the Agilent 34401A multimeter, while sweeping the input frequency of the circuit. After we determine the frequency at which the output voltage value reaches a maxima, we add the capacitor and inductor components on the matching circuit as series and parallel, respectively, to change the frequency of the peak response and draw it closer to 915MHz, which is the RF frequency of the WuTx.




Fig. 5. Photo of the energy harvesting circuit on REACH-Mote.

Table I. Components Used to Build the Energy Harvester

Component	Value	Component	Value
Series capacitor	0.1pF	Stage capacitor	36pF
Parallel capacitor	1.0pF	Diode	HSMS-2852

Table II. Parameters Used in PCB Fabrication for Dual-Stage Circuit Design

Component	Value
Laminate thickness	62 mil FR-4
Number of layers	2-layer, one serves as a ground plane
Copper thickness	1.7 mil
Trace width	20 mil with 12 mil gap
Dielectric constant	4.6
Through-hole size	29 mil

To ensure that energy transmission from the antenna to the circuit occurs with minimal waste of energy, we use a fine granularity in the component value selection—for instance, the capacitor value is varied from 0.1pF to 10pF with 0.1pF step size. Similarly, the value of the inductor is changed from 1nH to 10nH with 1nH step size.

After selection of the series components, we repeat a similar procedure to find the proper component values for the parallel connections of the matching network. These iterations finally result in the peak voltage being attained at a frequency very close to 915MHz. Figure 5 shows the final fabricated PCB of our energy harvesting module. The PCB is fabricated with FR-4 epoxy glass substrate and has two layers, one of which serves as a ground plane. We select components with values and ratings of their performance parameter as close as possible to the ones obtained from the simulation. This data is summarized in Tables I and II.

3.2. Wake-Up Circuit

Even with the high-efficiency energy harvester circuit, the energy received from the radio is limited. Thus, the wake-up circuit of the WuRx must meet the following design requirements:

- The wake-up circuit must consume as low energy as possible to achieve a long wake-up range.

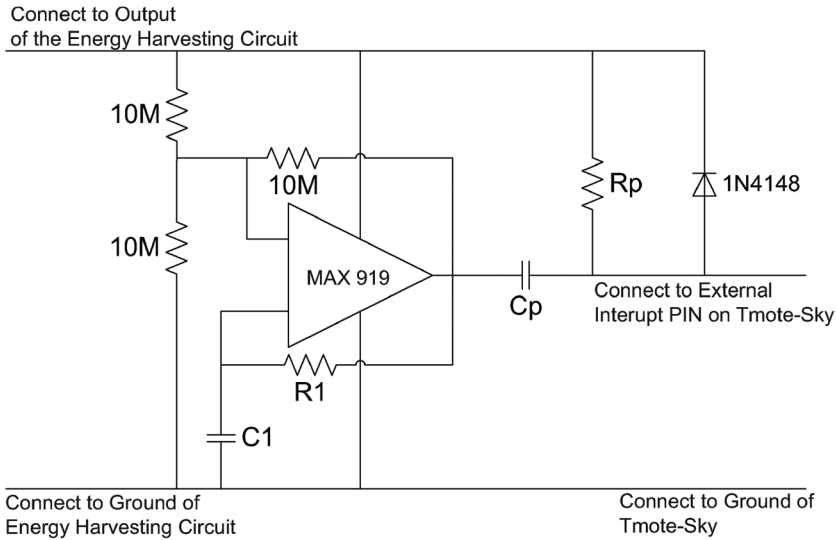


Fig. 6. Wake-up circuit of REACH-Mote.

- The wake-up circuit must generate a rising edge of 1.8V to trigger Tmote Sky to wake up from the sleep mode.
- The trigger circuit must work on a variable support voltage, as the voltage level output by the energy harvesting circuit is not stable.

Figure 6 shows the wake-up circuit of REACH-Mote. This circuit is an adaptation of a normal relaxation oscillator with a differentiator and diode clamp on the output to generate the pulse. The pulse width can be adjusted by varying the value of the capacitor C_p and the resistor R_p . The period of the pulse is determined by the value of C_1 and R_1 . In this design, we applied $C_p = 1\text{nf}$, $R_p = 270\text{k}\Omega$, $C_1 = 130\text{nF}$, and $R_1 = 8.2\text{M}\Omega$ to generate a pulse of $100\mu\text{s}$ width with a period of 1 second. Using these values, the wake-up circuit requires only $1\mu\text{A}$ with a supply voltage of 1.5V to 5V. Thus, with different input voltages from the energy harvester, the voltage output of the wake-up circuit can trigger the MCU on the sensor node. Note that this energy is drawn from the energy harvester circuit and not from the node's battery. Figure 7 shows a photo of the wake-up circuit.

3.3. Integration of REACH-Mote

We combine the RF energy harvesting circuit and the wake-up circuit, as well as Tmote Sky, to build REACH-Mote (Range EnhAnCing energy Harvester-Mote) passive wake-up radio sensor node [Chen et al. 2013]. When a wake-up signal is sent by the WuTx, the energy harvesting circuit outputs a DC voltage. The wake-up circuit starts to generate the pulse once the DC voltage is higher than 1.5V, and this will trigger the mote and put the mote's MCU into active mode in 5ms [moteiv 2006]. Note that the following steps are included in this period of time: MCU wake-up from sleep mode, wake-up of the operating system on Tmote Sky (TinyOS), and reinitialization of the radio chip (CC2420). After waking up, Tmote Sky starts the data transmission and goes back to sleep after the data transmission is complete. The energy harvesting circuit is a passive component that does not consume energy from the node's battery. The wake-up circuit is powered by the energy harvesting circuit, so the wake-up circuit also does not drain energy from the battery. Thus, all energy provided by the REACH-Mote battery




Fig. 7. Photo of the wake-up circuit on REACH-Mote.

is used for sensing, data processing, and data communication, and no energy is wasted on unnecessary communication overhead.

4. REACH²-MOTE

REACH²-Mote incorporates some design enhancements to improve the wake-up range compared to that of REACH-Mote. In particular, two approaches have been utilized to improve the efficiency of the wake-up design: improving the output of the energy harvester circuit and lowering the voltage required to trigger the MCU on Tmote Sky to wake up.

For the first approach, to improve the output of the energy harvester circuit, we note that the energy harvester circuit in REACH-Mote works as the battery supply for the wake-up circuit. Thus, increasing the number of energy harvesters and connecting them serially can increase the output voltage of the energy harvester. The serial connection between the two energy harvesters works just as a serial connection of two batteries, which can increase the voltage output from the energy harvesting circuit. As the wake-up circuit requires a minimum voltage to operate, the higher output voltage may potentially extend the wake-up range.

For the second approach, reducing the voltage required to wake up the MCU, we exploited the fact that Tmote Sky can work using different voltage values. Typically, Tmote Sky is powered by two AA batteries that provide a 3V power supply. The MCU on Tmote Sky, the TI MSP430 F1611, requires a 1.5V rising edge to be triggered with the 3V battery supply. However, a lower supply voltage can potentially decrease the requirement for the trigger signal. We designed a voltage regulator and a switch controlled by the digital I/O (DIO) of Tmote Sky to change the supply voltage of Tmote Sky between 3V and 2.5V, as 2.5V is a voltage that MSP 430 supports. We use two MCU DIO pins directly connected to the $\overline{EN1}$ and $\overline{EN2}$ pins on the TPS2042B. The $OUT1$ pin of the TPS2042B is connected to the voltage regulator AMS AS1375-BTDT-25, and the $OUT2$ pin is directly connected to the VCC of the MCU. When initializing the MCU, $\overline{EN1}$ is set to low and $\overline{EN2}$ is set to high. Thus, the voltage regulator output is connected to the VCC of the MCU. When switching the supply voltage, the MCU first sets $\overline{EN2}$ to low to enable the 3V VCC power supply, then it sets $\overline{EN1}$ to high to disable the voltage regulator output. By applying this approach, Tmote Sky can sleep at 2.5V

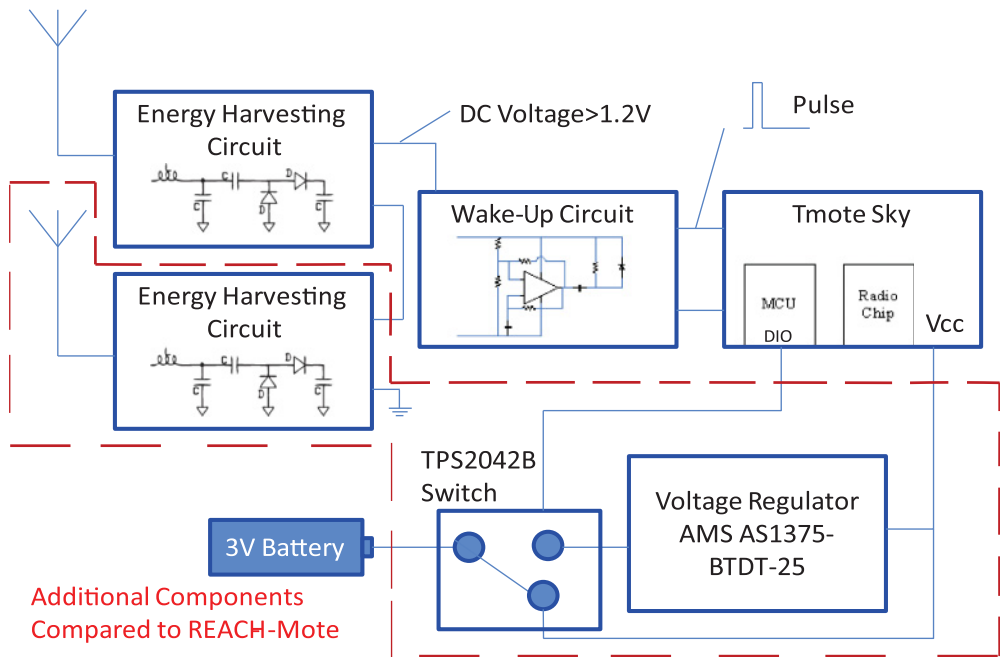


Fig. 8. Block diagram of the REACH²-Mote components.

voltage supply with a lower voltage trigger wake-up requirement. After the MCU of Tmote Sky is woken up, Tmote Sky then switches the supply voltage to 3V to obtain the best communication performance for the sensor node.

Thus, the main upgrades for REACH²-Mote compared to REACH-Mote are as follows:

- Increasing the number of energy harvesting circuits and antennas. As the antennas are separated on the mote, the additional energy harvesting circuit can provide increased energy to the wake-up circuit.
- Applying a voltage regulator to change the supply voltage of Tmote Sky. The voltage regulator reduces the amount of energy required to wake up the MCU on Tmote Sky, which thus increases the wake-up range of REACH²-Mote.

4.1. Operation of REACH²-Mote

Figure 8 shows the system diagram of REACH²-Mote. REACH²-Mote operates following the flow chart shown in Figure 9. In the following, we describe the operation principles for REACH²-Mote.

- REACH²-Mote remains in sleep mode before the WuTx transmits the wake-up signal (i.e., the MCU on Tmote Sky, which is an MSP430 F1611, is put to LPM3 sleep mode [MSP430 2012] and the radio on Tmote Sky is turned off).
- The voltage regulator maintains the battery supply voltage of REACH²-Mote at 2.5V.
- When a wake-up signal is sent by a nearby WuTx, the energy harvesting circuit receives the energy and outputs a DC voltage.
- The wake-up circuit generates a pulse once the DC voltage is higher than 1.2V, and this will trigger a wake-up of the MCU on the sensor mote. Note that the voltage requirement of wake-up has been lowered from 1.5V to 1.2V because the supply voltage of the MCU is set at 2.5V.

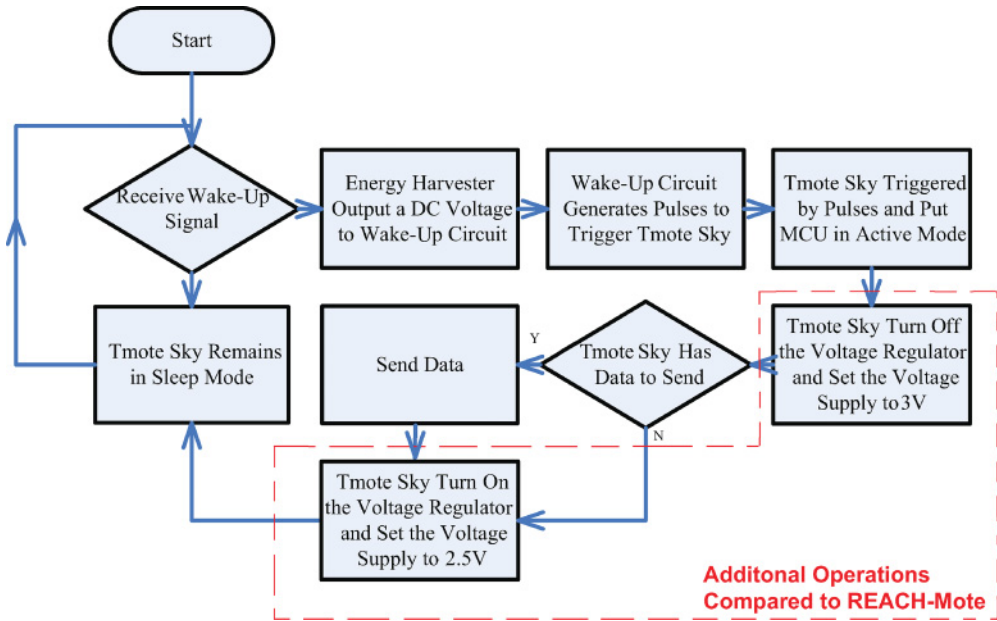


Fig. 9. Flow chart of the REACH²-Mote operation.

- The MCU changes the DIO pin on the voltage regulator and switches the power supply of the sensor node back to 3V.
- The MCU turns on the radio (i.e., the CC2420 radio on Tmote Sky). As the supply voltage is 3V at this time, the CC2420 can achieve a reasonable communication range.
- After turning on the radio, data transfer is started if the mote has data to send.
- If the mote does not have data to send, or after the data transmission is complete, the MCU switches the supply voltage back to 2.5V and the mote goes back to the sleep mode (i.e., the MCU is set to LPM3 and the radio is turned off).

The improved energy harvester circuit and the adaptation of the power supply voltage for Tmote Sky enable REACH²-Mote to extend the wake-up range compared to REACH-Mote, as shown in Section 5.1.1.

4.2. Energy Analysis of REACH²-Mote

A voltage regulator will require some energy from the node's battery. However, the lowered supply voltage also decreases the energy cost of the MCU during the sleep state. Thus, a well-selected voltage regulator is important to extend the lifetime of the sensor node. The voltage regulator used in REACH²-Mote must meet the following requirements:

- The input voltage of the voltage regulator circuit is 3V so that the input of the voltage regulator can share the same battery supply with Tmote Sky in active mode.
- The output voltage of the voltage regulator is 2.5V.
- The quiescent current of the voltage regulator should be as low as possible.

According to these criteria, we select the AMS AS1375-BTDT-25 [ams AG 2015] as the voltage regulator, as this chip only requires a quiescent current of 1 μ A. We also added a TI TPS2042B [TPS2042B 2012] to switch the supply voltage between 2.5V

and 3V. In addition, the switch consumes $1\mu\text{A}$ continuously. As the sleeping current of Tmote Sky is about $11.2\mu\text{A}$, the energy cost of the sleeping REACH²-Mote is $33\mu\text{W}$ ($28\mu\text{W}$ for the sleeping mote, $2.5\mu\text{W}$ for the switch, and $2.5\mu\text{W}$ for the voltage regulator) compared to the $33.6\mu\text{W}$ sleeping energy cost of a normal Tmote Sky powered by a 3V battery. Thus, with the new voltage regulator and switch system, the energy cost of the sensor node is lowered by 1.7%, and the wake-up voltage requirement of REACH²-Mote is decreased. Although the voltage regulator and the switch consume energy from the battery, this approach reduces the overall battery consumption of the mote. Hence, we consider this approach as a hybrid-passive WuRx approach.

5. EXPERIMENTS AND FIELD TESTS

We performed field tests to evaluate the performance of REACH-Mote and REACH²-Mote. We use the field test results for REACH²-Mote to build a simulation model to evaluate the performance of REACH²-Mote in detailed application scenarios.

5.1. Experiments and Field Tests for REACH-Mote

We evaluated the wake-up delay and wake-up distance performance of REACH-Mote through field tests and compared its performance with that of an existing passive wake-up sensor node, namely WISP-Mote [Ba et al. 2010]. The wake-up delay is mainly caused by the delay of the energy harvester, as the energy harvester circuit takes some time to accumulate enough energy to power the wake-up circuit. Thus, the efficiency of the energy harvesting circuit has a large impact on the wake-up delay. In addition, the distance between the WuTx and the WuRx impacts the wake-up delay as well, as this impacts the received energy. When the distance is short, the received energy is high and it takes less time for the energy harvesting circuit to accumulate enough energy to trigger a wake-up. We thus characterize the wake-up delay as a key metric to evaluate the performance of the wake-up sensor node.

5.1.1. Experiments and Field Test Setup. We ran several experiments in an open-space environment (an empty gymnasium). WISP-Mote is capable of both addressable wake-up and broadcast wake-up, but REACH-Mote is only capable of broadcast wake-up. Hence, we only evaluate the performance of WISP-Mote utilizing broadcast wake-up for this test for a fair comparison. In our experiments, we tested the single-hop wake-up scenario, assuming that a base station with a WuTx transmits the wake-up signal to collect data on the REACH-Mote and WISP-Mote. The base station is composed of a WuTx, Tmote Sky, and a laptop. The WuTx is composed of a Powercast wireless transmitter [Powercast 2009] and an Impinj R1000 RFID reader [Impinj 2002] controlled by the laptop. The transmit power of both the Powercast transmitter and the RFID reader is 1W. After the WuTx transmits the wake-up signal and wakes the sensor node (REACH-Mote and WISP-Mote), Tmote Sky on the sensor node transmits a short ACK packet indicating the successful wake-up to the base station. We evaluate the period between the start of the wake-up signal transmission and the reception of the ACK packet. As there are no collisions occurring in this scenario, this period represents the wake-up delay.

We placed the transmitter (WuTx) antenna 2 feet above the ground and varied the location of REACH-Mote and WISP-Mote (WuRx) in both the horizontal and vertical directions to evaluate their performance. If the mote does not respond within 100 seconds, we assume that it cannot be woken up at that particular location. Figure 10 shows the field test setup.

5.1.2. Experiments and Field Test Results. The tests are repeated with 2-foot increments in the horizontal direction (x -direction) starting 0.1 foot from the WuTx and 1-foot increments in the vertical direction (z -direction), with 0 corresponding to the ground



Fig. 10. Field test setup.




Fig. 11. Wake-up delay (in seconds) for WuTx: combination of RFID Reader and Powercast; WuRx: WISP-Mote. The delay limit of 100 seconds is used to represent the locations where wake-up is not possible.




Fig. 12. Wake-up delay (in seconds) for WuTx: combination of RFID Reader and Powercast; WuRx: WISP-Mote. The delay limit of 100 seconds is used to represent the locations where wake-up is not possible.

level. After each measurement, Tmote Sky is reset and the energy harvesting circuit is discharged. Each data point in Figures 11–14 represents the average of five tests.

As seen in Figures 11 and 12, REACH-Mote can achieve a 37-foot wake-up range, more than double the distance compared to that of WISP-Mote, which achieves a 17-foot




Fig. 13. Average results for Test1 and Test2, which are performed on a clear day. Wake-up delay (in seconds) for WuTx: combination of RFID Reader and Powercast; WuRx: REACH²-Mote. The test is performed in the x and y directions with the height set at $z = 2$ feet. The delay limit of 100 seconds is used to represent the locations where wake-up is not possible.

wake-up range. This is due to the ultra low energy consumption of the proposed wake-up circuit and an optimized energy harvesting circuit. Furthermore, the longest range is achieved at the 2-foot height, which is the same height as the wake-up transmitter.

5.2. Experiments and Field Tests for REACH²-Mote

Here, we provide the experimental results for REACH²-Mote. As we see from the previous results that the 2-foot height achieves the best results vertically (z -direction), the REACH²-Mote tests are performed only at this height. For these experiments, we vary both the x -direction and the y -direction. In addition, three sets of tests are performed during different days, with one being a rainy day to evaluate the performance of REACH²-Mote under different environmental conditions. Although these tests are performed indoors, the rainy day increases the moisture of the air, which will decrease the performance of REACH²-Mote somewhat. Each set of tests is performed three times, and the average values of the wake-up delays are calculated. The tests are repeated with 1-foot increments in the x -direction starting 0.1 foot from the WuTx and 3-foot increments in the y -direction. The other settings in these tests are the same as the tests for REACH-Mote and WISP-Mote.

Figure 13 shows the results of wake-up coverage for REACH²-Mote for Test1 and Test2, which are both performed on a clear day. We see that REACH²-Mote can achieve a wake-up distance of 44 feet, which represents a 19% improvement compared to REACH-Mote. Figure 14 shows the results of wake-up coverage for Test3, which is performed on a rainy day. As shown in Figure 14, during the rainy day, REACH²-Mote achieves a 43-foot wake-up distance, which shows that the high moisture in the air does little to degrade the performance of REACH²-Mote.

In the y -direction, as the WuTx on the base station is composed of a directional antenna, the results show that REACH²-Mote can be woken up at ± 19 feet. These results are used in the modeling for the simulation to further evaluate the performance of REACH²-Mote in different network scenarios.




Fig. 14. Results for Test3, which is performed on a rainy day. Wake-up delay (in seconds) for WuTx: combination of RFID Reader and Powercast; WuRx: REACH²-Mote. The test is performed in the x and y directions with the height set at $z = 2$ feet. The delay limit of 100 seconds is used to represent the locations where wake-up is not possible.

6. SIMULATION RESULTS

Due to the prototype phase of the hardware, we cannot build many REACH²-Motes to perform a full-scale test in a large network. Hence, to evaluate the performance of REACH²-Mote in a network scenario with multiple REACH²-Motes, we build an energy harvesting model of REACH²-Mote based on the field test results. In addition, we build a communication model for REACH²-Mote and WISP-Mote, as well as for an active wake-up scenario and for a duty cycling approach. In this way, we can compare the performance of these different approaches for a range of network scenarios. Additionally, we build a simulation scenario for a particular application—air pollution monitoring—and evaluate the performance of these approaches for this application.

6.1. Models Created for the Simulation

To perform the simulations, we modeled the energy harvesting process of REACH²-Mote by measuring the wake-up delay. We assume that the sensor node will be woken up when the energy harvester receives enough energy to trigger the MCU. After that, we build a communication model for the communication between the sensor nodes and the base station(s).

6.1.1. Energy Harvesting Model. An energy harvesting model is developed to indicate the amount of energy harvested for the wake-up based on the locations of the WuTx and the WuRx. For the energy harvesting model, we make the following assumptions. First, we assume that the amount of energy harvested from the transmitter at a fixed location (x, y) in a unit time is constant. We denote this location-specific constant value with $E_h(x, y)$. We assume that a wake-up circuit consumes E_c amount of energy when it wakes up the MCU on the sensor node. In addition, the capacitor leaks E_l amount of energy per unit time when the wake-up circuit is not active. Thus, the amount of energy in a REACH²-Mote capacitor at time t when it is not sending a wake-up trigger to the MCU is

$$E_t = E_{t-1} + E_h(x, y) - E_l, \quad (1)$$

and the energy in the capacitor at time t when REACH²-Mote is woken up is

$$E_t = E_{t-1} + E_h(x, y) - E_c. \quad (2)$$

Note that the leakage when the wake-up circuit is active is negligible because $E_c \gg E_l$. The values E_c and E_l are measured through field tests. To do this, we charged the capacitor and turned on the wake-up circuit, and then we measured the voltage change on the capacitor to calculate E_c . Then we turned off the wake-up circuit and measured the leakage E_l .

Assuming that there is no energy stored at the beginning of the simulation, we can calculate the energy stored in the capacitor of the WuRx. We measured the voltage value on the capacitor (C_w) when it is just sufficient to trigger a wake-up. Then we calculate the energy based on the following equation:

$$E'_t = \frac{1}{2} C_w V_t'^2. \quad (3)$$

Let $T_d(x, y)$ define the wake-up delay when REACH²-Mote is deployed at location (x, y) relative to the base station. With the assumption of constant energy harvesting at one location,

$$E_h(x, y) = E'_t / T_d(x, y). \quad (4)$$

Note that as E'_t is the energy that is barely sufficient to trigger a wake-up, this represents the threshold energy to turn on the wake-up circuit.

Figure 15 shows the energy harvesting model that we are using in the simulation framework.

6.1.2. Communication Model. To compare the performance of REACH²-Mote, WISP-Mote, an active wake-up approach, and duty cycling approach, we build communication models for these approaches. Note that the approach of the active wake-up is based on the work described in Pletcher et al. [2009], as it is the only active wake-up with -72 dBm sensitivity (i.e., long wake-up range). The communication is modeled based on time slots, where each time slot is 10ms.

For REACH²-Mote, we build the communication model based on the energy harvesting model. When a sensor node is woken up, it performs carrier sensing using its communication radio. The node will sense the channel immediately after it wakes up. If the channel is clear, the sensor node will transmit its data to the base station. The base station will provide an ACK once it successfully receives the data. If the channel is busy, the sensor node will back off for a random number of time slots. If the transmission is not successful (i.e., an ACK is not received from the base station), the sensor node will back off for another random number of time slots and retransmit the data.

For WISP-Mote, we build the wake-up model based on the wake-up probability model given in Ba et al. [2010]. When the node is located in the wake-up range of the WuTx, the node has a given probability to wake up. After the node is woken up, it acts the same as REACH²-Mote.

For active wake-up, we assume that the sensor node is woken up as soon as the base station moves into the wake-up range of the sensor node. After that, the sensor node performs carrier sensing in the same way as for REACH²-Mote and WISP-Mote.

For the duty cycling approach, the base station transmits a beacon packet once every eight time slots and waits for a response for the remaining seven slots. If there is no response from a sensor in these seven slots, the base station transmits the beacon packet again. The sensor node remains in the sleeping mode until a preset timer wakes it up. The timer is set based on the ratio of active/sleep mode, which represents different duty cycle values. After the sensor node is woken up by the timer, it starts to listen for the channel for eight time slots to guarantee not missing the beacon signal if a base station is nearby. If the sensor node receives the beacon packet, it will randomly select one of the next seven slots to transmit data to the base station. Otherwise, it will reset

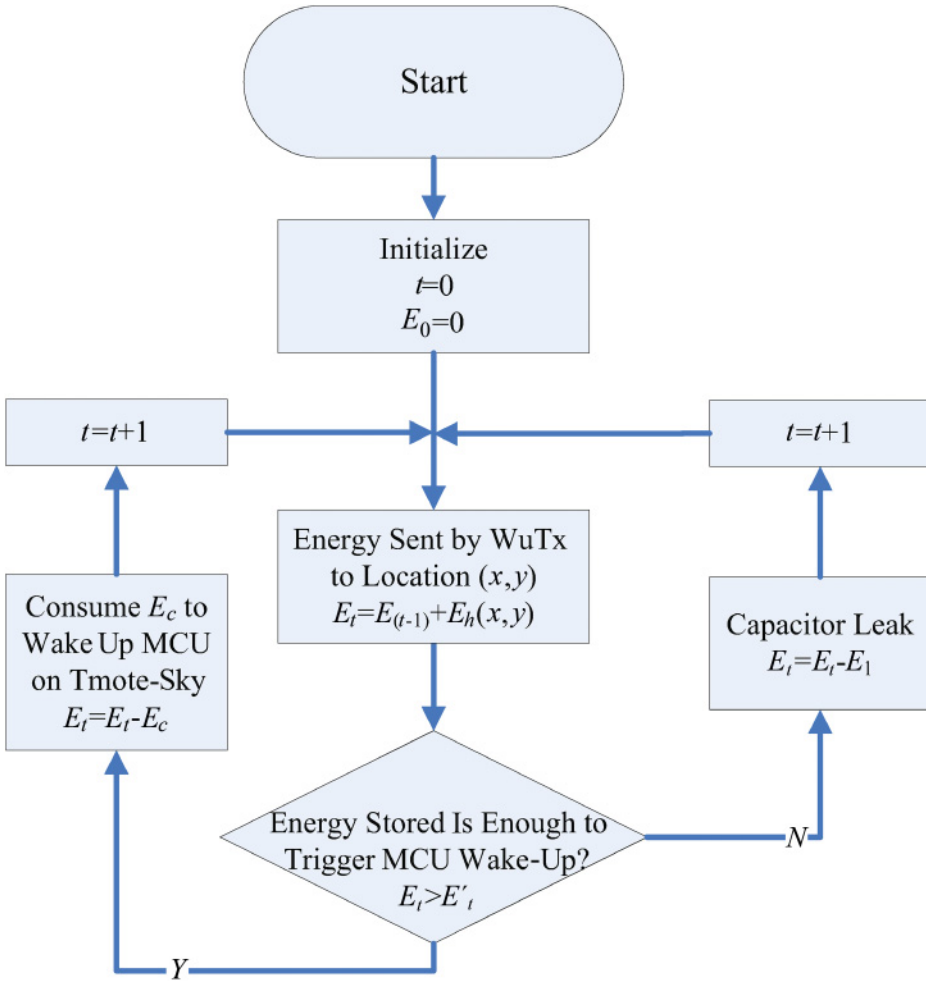


Fig. 15. Energy harvesting model for the simulations.

the wake-up timer and return to the sleep mode. If the transmission to the base station is not successful due to collisions, the sensor node will back off for a random number of time slots and pick another random slot in the seven slots to retransmit the data.

For all four approaches, the sensor node will receive an ACK packet after a successful transmission. The ACK packet notifies the sensor node that the base station is still within its communication range and that no collisions occurred during the data transmission. Thus, the sensor node can continue to transmit other packets stored in its buffer. After emptying its buffer, or if the base station goes out of communication range and no longer sends ACK packets, the sensor node will not receive the ACK for a period of time and it will return to the sleep mode.

6.2. Simulation Setup

To evaluate the performance of the investigated approaches, we consider two categories of application scenarios: one with a low data rate requirement and one with a high data rate requirement. In the low data rate requirement scenarios, the sensor

nodes generate packets with a relatively long interval. This category simulates the sensing tasks that do not require continuous monitoring, such as air pollution control, temperature, and moisture monitoring, where a measurement/reading might be taken only once an hour or even once a day. On the other hand, a high data rate requirement sensing task generates packets much more frequently and performs continuous sensing observations, such as for hazard monitoring.

The simulations are performed in Matlab and utilize the following simulation setup:

- The sensor nodes are deployed randomly in an area of $200\text{m} \times 200\text{m}$.
- There are one or multiple mobile base stations that move with a random direction mobility model with a speed of 10 m/s [Nain et al. 2005].
- The nodes generate packets according to the designated packet generation rate periodically and store these packets in their buffers. The sensor nodes can have finite buffer size or infinite buffer size depending on the scenario. For finite buffer size, the oldest packet is dropped when the buffer is full.
- For the wake-up scenarios, once the base station is within the wake-up range of the sensor nodes, they wake up according to the model described in Section 6.1.1.
- For the duty cycling approach, the sensor node wakes up according to its internal timer.
- After the sensor nodes wake up, they apply the communication model described in Section 6.1.2.
- Each simulation run lasts for 6 hours with a timestep of 10ms.

In each category—both low data rate and high data rate—three sets of simulations are performed as detailed next:

- (1) *Set 1*: 100 sensor nodes in the $200\text{m} \times 200\text{m}$ area. There is one mobile base station collecting data. The sensor nodes have infinite buffer size. The packet generation rate changes from 0.02pkt/min to 0.2pkt/min for category 1 and 0.2pkt/min to 2pkt/min for category 2.
- (2) *Set 2*: The same as Set 1 except that the buffer size is 10pkt instead of unlimited.
- (3) *Set 3*: Varying the number of base stations from 1 to 10. The packet generation rate is 0.02pkt/min for category 1 and 0.2pkt/min for category 2 with unlimited buffer. The number of sensor nodes is 100.

We also implemented an air pollution monitoring scenario in the simulations to evaluate the performance of these approaches in a real application. In this scenario, 100 sensor nodes are deployed along the road. Each sensor node is equipped with the following air quality sensors: CO gas sensor, CO₂ gas sensor, CH₄ gas sensor, NH₃ gas sensor, NO₂ gas sensor, and volatile organic components sensor. Each node will collect air pollution information once every hour. The base station moves along the designed route to collect air pollution data once a day. When the base station establishes communication with a sensor node, it downloads the stored sensed data and updates the timer on the sensor node. Thus, all sensor nodes will have approximately synchronized timers so that all sensor nodes will sense the air pollution information roughly at the same time. The route is 10km long, and the simulation runs for 2 days.

6.3. Simulation Results

In all of the simulations, we collect data for five performance metrics to evaluate the performance of the different approaches:

- Average buffer size represents the memory requirement needed to store the packets that have not been sent. The lower the average buffer size, the less memory required on the sensor node.

- Average collisions per packet represents the collisions that occur during the communication with the base station. The higher the number of collisions, the higher the retransmission rate, which will cost additional energy.
- Average packet delay measures the delay between when a packet is generated and when the packet is received by the base station. A high packet delay is caused by missed wake-ups, short wake-up range, or high collisions in data transmission.
- Energy consumption per packet represents the energy efficiency in data transmission. Packet retransmission, unnecessary wake-up for the wake-up approaches, and unnecessary idle listening for the duty cycling approach will increase this value. A lower energy consumption per packet represents a better energy efficiency.
- PDR calculates the ratio between the number of packets generated by the sensor node and the number of packets delivered to the base station.

6.3.1. Set 1 Simulation Results. Figure 16 shows the performance of each approach with varying packet generation rates from 0.02pkt/min to 0.2pkt/min (category 1). In this set of simulations, there are 100 nodes deployed in the area and one base station moving within the target area to collect the data. The buffer size is assumed to be unlimited for sensor nodes in this set of simulations. We can see that none of the approaches requires much buffer space, as the packet generation rate is relatively low. The buffer requirements for REACH²-Mote are lower than for WISP-Mote, as the longer wake-up range increases the possibility of packet delivery. The 0.1% duty cycling, WISP-Mote, and REACH²-Mote achieve a low collision rate. Among these, WISP-Mote is a little less than the others, as WISP-Mote provides a low wake-up range, which decreases the probability of waking up multiple sensor nodes at the same time to transmit data. The 10% duty cycling provides the best delay performance and REACH²-Mote and active wake-up perform almost the same as the 10% duty cycling approach. REACH²-Mote and WISP-Mote result in the best energy consumption performances, as both approaches are passive wake-up sensor nodes. The active wake-up approach doubles the energy consumption compared to the passive wake-up approaches. The 10% duty cycling results in the worst energy efficiency, as expected, since it wastes a lot of energy on unnecessary idle listening. Although WISP-Mote performs well in terms of energy efficiency, it results in the worst buffer requirement and delay result, as the wake-up range of WISP-Mote is short.

Figure 17 shows the simulation results when the packet generation rate is varied from 0.2pkt/min to 2pkt/min (category 2). This simulation aims to evaluate the performance of each approach when the sensor nodes require a high data transmission rate. Results show that all approaches, except REACH²-Mote and the active wake-up approach, require higher buffer occupancies, as increasing the packet generation rate leads to a lower PDR, and more packets are stored in the buffer for these approaches. Referring to the average packet delay and packet delivery ratio results, we find that REACH²-Mote and the active wake-up approach can deliver most of their packets, so REACH²-Mote and active wake-up approaches increase little when the packet generation rate increases. As we do not implement addressable wake-up for the active wake-up approach, the active wake-up leads to a high collision rate due to the large wake-up range (i.e., more nodes being woken up simultaneously). Note that for these results, when the packet generation rate is 2pkt/min, the results show the performance for each approach in a heavy data rate scenario. Compared to duty cycling and the active wake-up approach, the passive wake-up approaches result in a huge advantage in energy cost (50% less than the active wake-up approach and 90% less than the 0.1% duty cycling approach) with high PDR and low packet delay. In addition, passive wake-up requires less memory for the buffer compared to the other approaches.

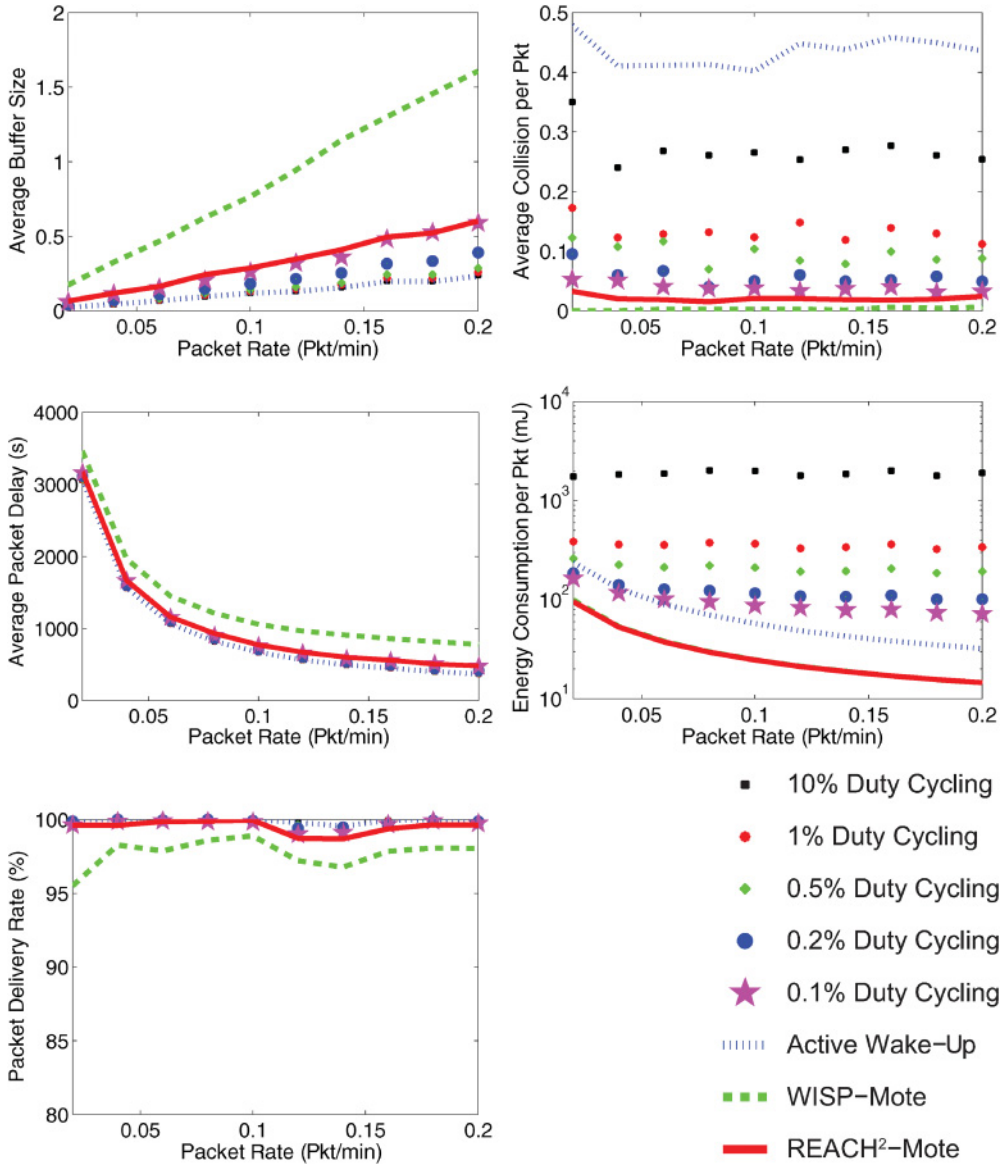


Fig. 16. Simulation results for different packet generation rates from 0.02pkt/min to 0.2pkt/min (100 sensor nodes, one base station, unlimited buffer).

6.3.2. Set 2 Simulation Results. Figure 18 shows the simulation results for the limited buffer case for low packet generation rate scenarios, and Figure 19 shows that of high packet generation rate scenarios. The packet generation rate varies from 0.02pkt/min to 0.2pkt/min (category 1) and from 0.2pkt/min to 2pkt/min (category 2). For the packet generation rate from 0.02pkt/min to 0.2pkt/min, the results are similar to the unlimited buffer results, as the low packet generation rate does not require much storage in memory. The effects of the limited buffer size are more visible as the packet generation rate increases. All approaches, except the active wake-up approach and 10% duty cycling, achieve lower PDR performance with a limited buffer in this scenario.

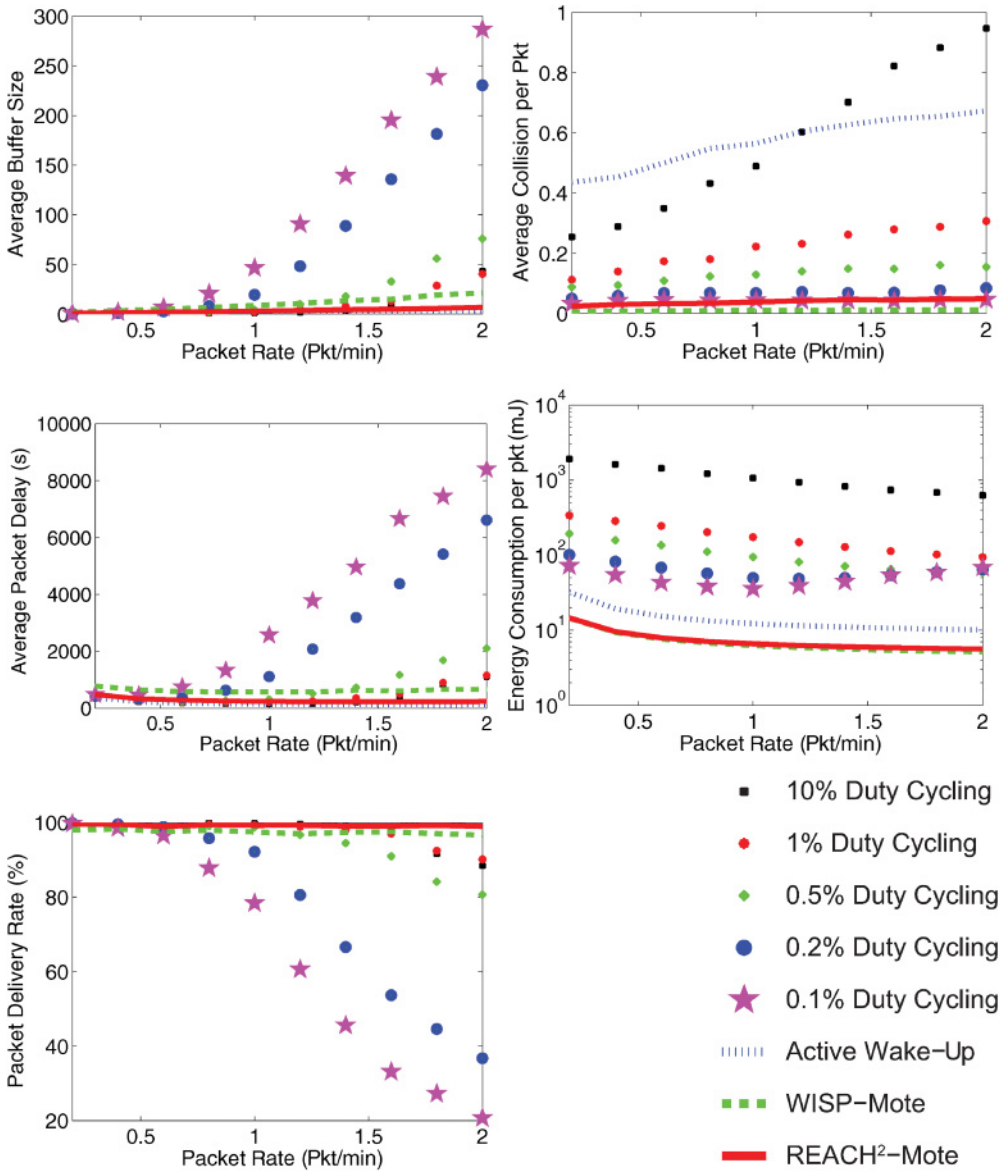


Fig. 17. Simulation results for different packet generation rates from 0.2pkt/min to 2pkt/min (100 sensor nodes, one base station, unlimited buffer).

REACH²-Mote can still provide a decent performance in terms of PDR while requiring only 40% of the energy necessary for the active wake-up approach and 0.7% of the energy necessary for the 10% duty cycling case.

For the simulation results when the packet generation rate is 0.02pkt/min and 2pkt/min for the limited buffer scenario, REACH²-Mote outperforms all other approaches in terms of energy efficiency. Active wake-up performs the best in terms of packet delivery ratio and latency with about double the energy consumption compared to REACH²-Mote. A high duty cycling approach performs well in terms of packet

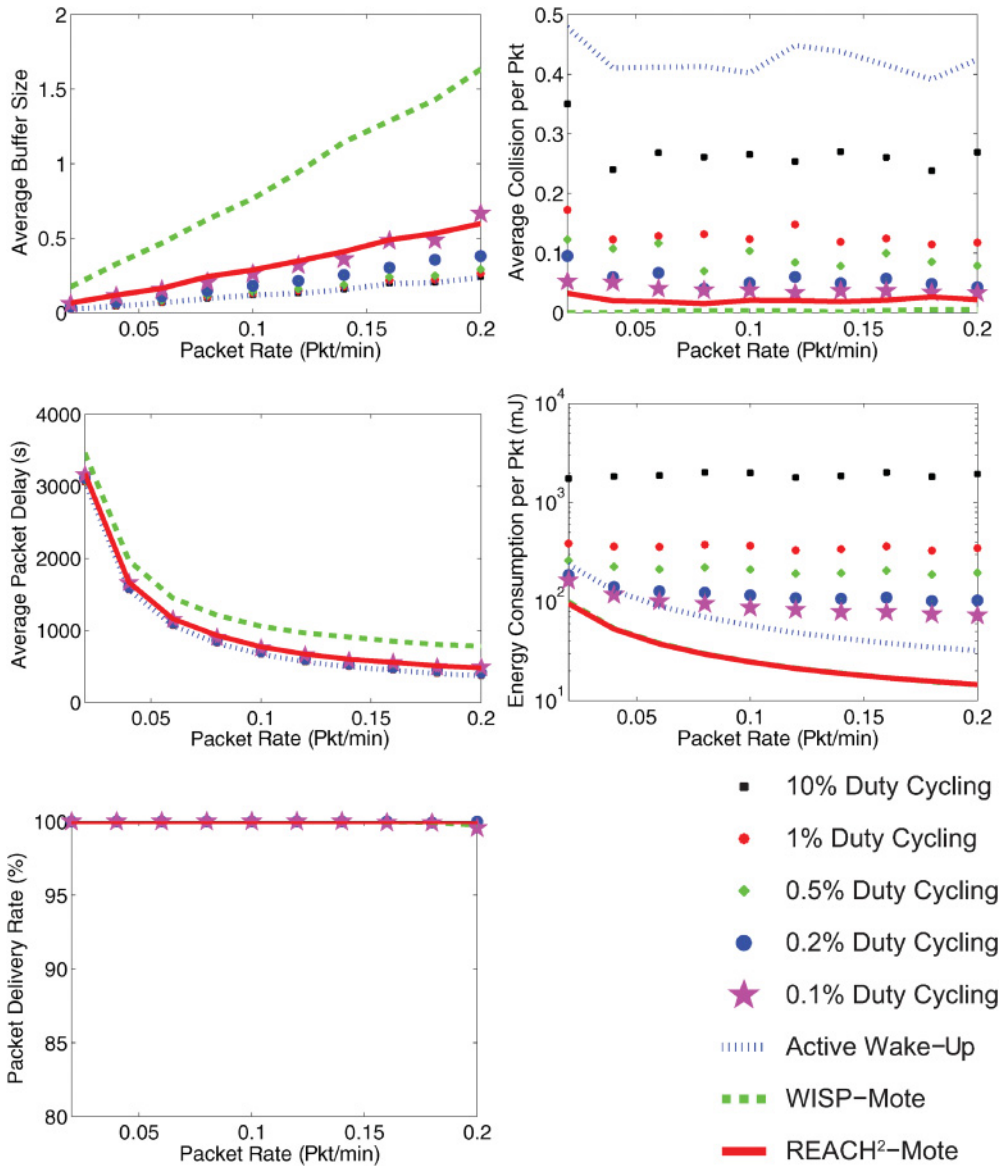


Fig. 18. Simulation results for different packet generation rates from 0.02pkt/min to 0.2pkt/min (100 sensor nodes, one base station, limited buffer).

delivery ratio and latency. However, duty cycling requires much more energy than the different wake-up approaches.

6.3.3. Set 3 Simulation Results. Figures 20 and 21 show the results of the performance of each approach with increasing the number of base stations. The results show that increasing the number of base stations can increase the performance for each approach. Even with a high packet generation rate, all approaches can result in a good PDR. REACH²-Mote and WISP-Mote result in the best energy efficiency performance compared to the other approaches.

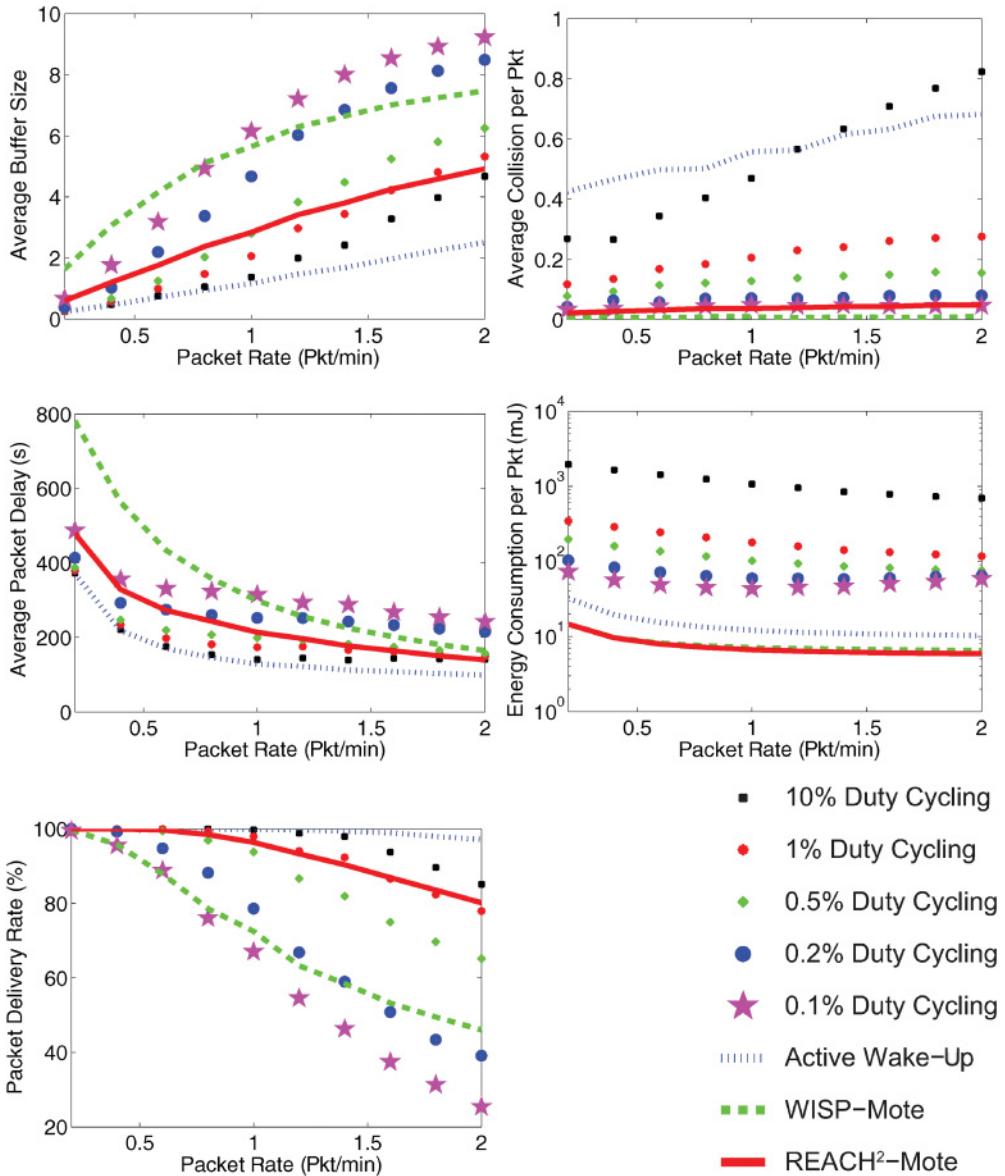


Fig. 19. Simulation results for different packet generation rates from 0.2pkt/min to 2pkt/min (100 sensor nodes, one base station, limited buffer).

6.3.4. Air Pollution Monitoring Scenario. Figure 22 shows the simulation results for the air pollution monitoring scenario, in which the base station moves along the designed route to collect air pollution data from 100 sensor nodes once a day. The results show that all approaches require a limited buffer, as the packet generation rate is low. In addition, the average collision rate is very low for all approaches, as this scenario represents a sparse network. The packet delay is mainly caused by the interval between the visits of the base station so that all approaches lead to high packet delays. The low duty cycling approach leads to higher delay compared to the other approaches, as some

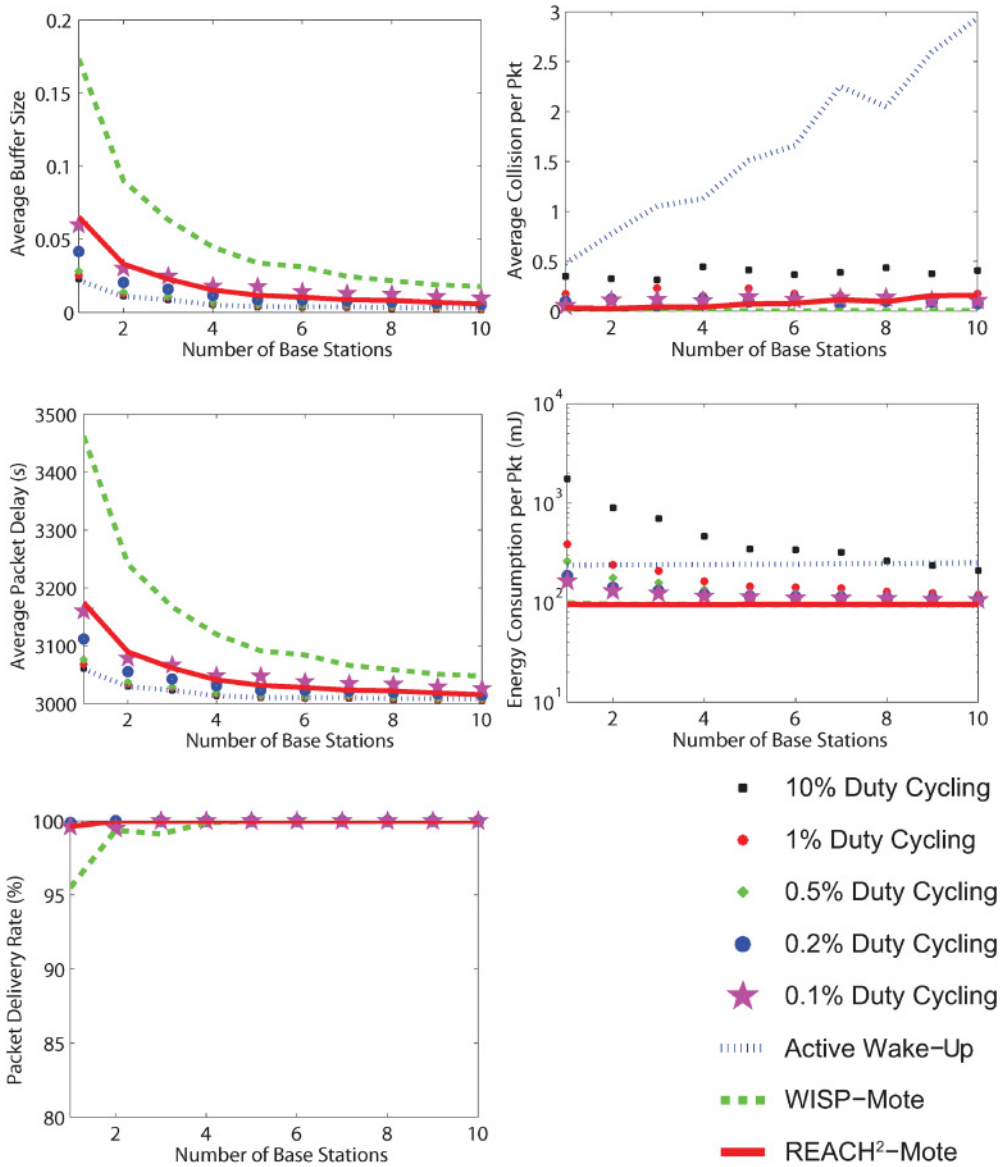


Fig. 20. Simulation results as the number of base stations varies from 1 to 10 (0.02pkt/min, 100 sensor nodes, unlimited buffer).

nodes miss the base station when it comes by. The results show that REACH²-Mote, WISP-Mote, and the active wake-up require much less energy compared to the duty cycling approach. As the data rate of this scenario is relatively low, a duty cycling approach wastes much of its energy on idle listening, especially for the 10% duty cycling. The energy cost of REACH²-Mote (108mJ) is only 41% of that required for active wake-up (263mJ). As well, all wake-up approaches perform well in terms of PDR. The 10% duty cycling is the only approach that results in good PDR among all

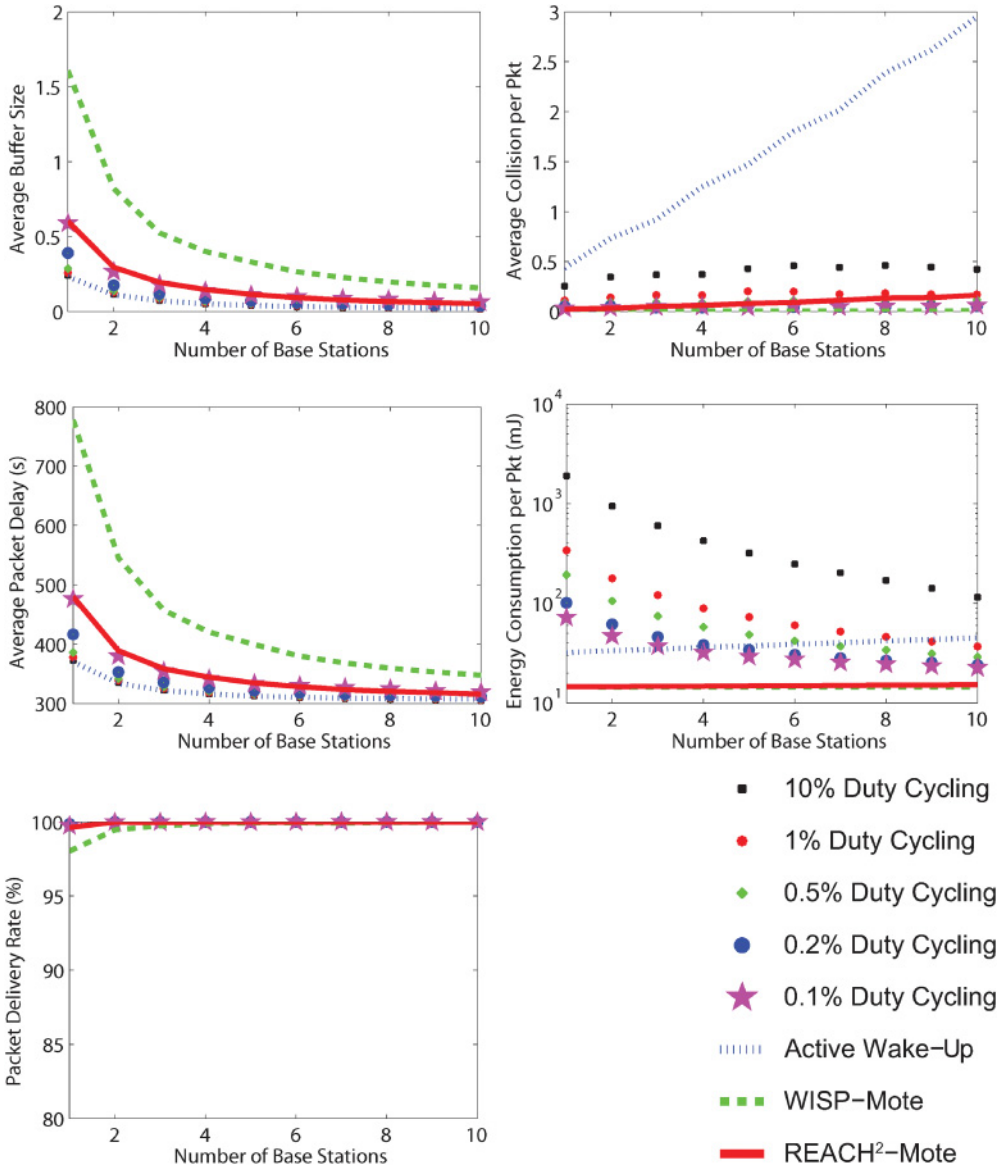


Fig. 21. Simulation results as the number of base stations varies from 1 to 10 (0.2pkt/min, 100 sensor nodes, unlimited buffer).

duty cycling approaches, as a lower duty cycle leads to a higher probability of missing an opportunity to communicate with the base station.

6.3.5. Conclusions on Simulation Results. These four sets of simulations show that REACH²-Mote and WISP-Mote provide the best energy performance compared to all other approaches. These two approaches can save quite a bit of energy compared to the 0.1% duty cycling approach. Considering that the 0.1% duty cycling performs worst among all duty cycling approaches in terms of buffer size, latency, and PDR, REACH²-Mote and WISP-Mote outperform duty cycling in most metrics evaluated.

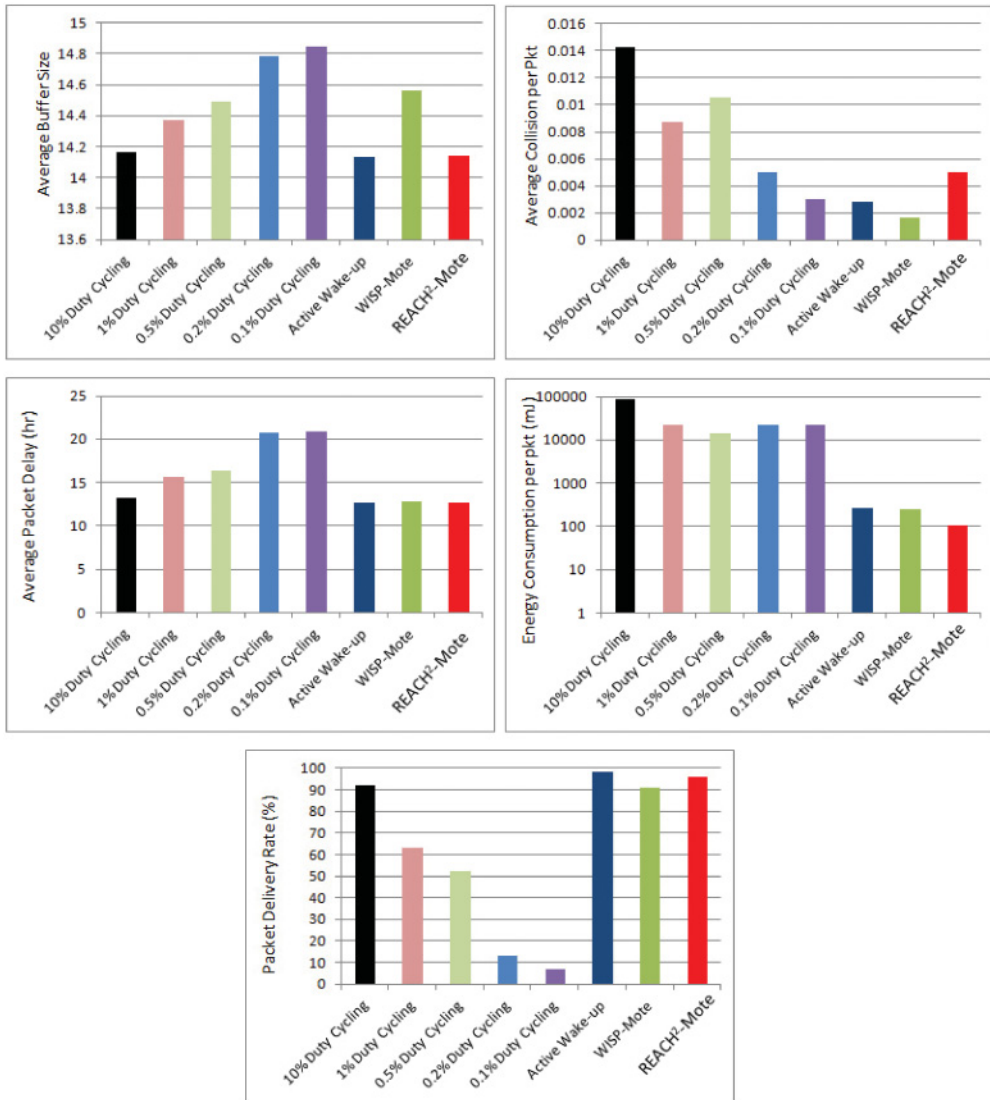


Fig. 22. Simulation results for the air pollution monitoring scenario.

Compared to the active wake-up approach, REACH²-Mote and WISP-Mote result in huge energy savings. REACH²-Mote can also provide better collision performance with similar performance in terms of buffer size and PDR compared to active wake-up. As REACH²-Mote and WISP-Mote are both passive wake-up sensor nodes, they result in very close energy consumption performance. However, as WISP-Mote provides a shorter wake-up range, REACH²-Mote outperforms WISP-Mote in terms of buffer size requirement, latency, and PDR.

The pollution monitoring scenario analysis shows us that the duty cycling approach is not suitable for a low collection rate scenario. All wake-up approaches perform well in this scenario, but REACH²-Mote results in the highest energy efficiency.

7. A COMPARISON OF DIFFERENT RADIO WAKE-UP DESIGNS

Both field test and simulation results show that different WuRx designs lead to different wake-up performance. To design a wake-up sensor node for WSNs, some design options must be considered. In this section, we compare and contrast the following wake-up radio designs: (1) battery-powered WISP, (2) WISP-Mote, and (3) REACH-Mote and REACH²-Mote. These three different wake-up radios utilize different approaches to achieve passive wake-up.

7.1. Battery-Powered WISP

Prior work has shown that a battery-powered WISP can be woken up at a distance of 45 feet by an RFID reader. Although the WISP is powered by the battery, the wake-up circuit, which is composed of an s-1000c20-n4t1x voltage detector and a NLSV1T244 translator, is powered by the energy harvested by the WISP. Thus, the battery-powered WISP is a simple and effective passive wake-up sensor node. The WISP uses an MSP430 F2132 ultra low power MCU that works at 1.8V as the microcontroller. It is clear that this approach has the following advantages:

- The design of the wake-up circuit on the WISP is simple and efficient.
- As the WISP uses a backscatter approach to communicate with the RFID reader, the energy cost for data transmission on the node is very small.
- As the CPU is an ultra low power MCU, the amount of harvested energy required to wake up the MCU is low. Thus, the wake-up distance is improved.

However, as the WISP is powered by a low-power MCU and uses backscatter communication, there are some disadvantages as well:

- The ultra low power MCU does not provide powerful data processing capabilities. Thus, it may be difficult to implement data fusion and postprocessing on the sensor node.
- As the WISP uses a backscatter approach to communicate with the RFID reader, it is very hard to build a long-range multihop WSN.
- As the wake-up signal triggers the MCU when it receives enough energy from the energy harvesting circuit, a battery-powered WISP can only perform broadcast wake-up. Thus, some false alarms may occur for this approach.

Thus, we can see that the battery-powered WISP approach is very suitable to build a single hop wake-up, single hop, shorter-range communication sensor node.

7.2. WISP-Mote

WISP-Mote is built by combining a WISP passive RFID tag (without battery support) and Tmote Sky. The energy harvested by the WISP first wakes up the MCU on the WISP. Then, the WISP transmits a wake-up trigger to Tmote Sky to wake up the MCU on Tmote Sky. Then, Tmote Sky starts its CC2420 radio chip to communicate with the base station. Due to this dual MCU design, the WISP-Mote design has the following advantages:

- As the MCU on the WISP wakes up before waking up Tmote-Sky MCU, the WISP's MCU can decode the information transmitted by the WuTx and detect an address transmitted by the WuTx. Thus, WISP-Mote can achieve ID-based wake-up.
- As the MCU on the WISP can decode information transmitted by the WuTx, there are no false wake-ups for the WISP-Mote approach.
- As Tmote Sky uses the CC2420 as the radio chip, the communication range of WISP-Mote is much higher than that of the battery-powered WISP. Considering a mobile

data mule scenario, if the data mule wakes WISP-Mote and moves out of the wake-up range, it is still possible for WISP-Mote to upload all of its data to the data mule.

However, the benefit of the dual MCU design of WISP-Mote also leads to some disadvantages:

- As the sensor node is composed of two MCUs, the cost of this approach is much higher than the other approaches.
- As the MCU on the WISP is powered by the energy harvested by the WISP, the energy cost of the wake-up circuit is much higher than the other approaches. Thus, the WISP-Mote approach results in a comparatively short wake-up range.

Using an even lower-power MCU on the WISP may improve the performance of WISP-Mote. However, considering that the wake-up circuit of both the battery-powered WISP and REACH²-Mote cost around $2\mu\text{W}$ power, it is very hard for an MCU to achieve such a low energy cost. Thus, building a long-range wake-up sensor node using dual MCUs is difficult.

7.3. REACH-Mote and REACH²-Mote

REACH-Mote and REACH²-Mote use an energy harvesting circuit and a wake-up circuit to achieve long-range wake-up. Thus, REACH-Mote and REACH²-Mote cost much less than WISP-Mote. In addition, REACH-Mote and REACH²-Mote have the following advantages:

- The cost in terms of hardware is much lower than for WISP-Mote.
- The wake-up range is much longer than WISP-Mote and is comparable to the battery-powered WISP.
- As REACH-Mote and REACH²-Mote use MSP430 F1611 MCUs, the additional computation ability can support complex data fusion and processing operations that cannot be supported by the battery-powered WISP.
- As Tmote Sky uses the CC2420 radio chip, the communication ranges of REACH-Mote and REACH²-Mote are much higher than that of the battery-powered WISP. Considering a mobile data mule scenario, if the data mule wakes REACH-Mote or REACH²-Mote and moves out of the wake-up range, it is still possible for REACH-Mote and REACH²-Mote to upload all of their data to the data mule.
- In a sensor network, it is possible to use REACH-Motes and REACH²-Motes to build a hybrid network. REACH-Mote and REACH²-Mote can work as a cluster head to collect data from other sensor nodes through a standard 802.15.4 network. Then, when they are woken up by a data mule, they can upload all of the information from their cluster to the data mule.

However, REACH-Mote and REACH²-Mote have the following disadvantages:

- As the wake-up circuit of REACH-Mote and REACH²-Mote does not contain an MCU, it is impossible for this approach to achieve ID-based wake-up.
- As REACH-Mote and REACH²-Mote can only perform broadcast wake-up, some false wake-ups will occur.

Comparing REACH-Mote and REACH²-Mote with WISP-Mote and battery-powered WISP, we can see that REACH-Mote and REACH²-Mote provide a good wake-up range while supporting a sensor node with a conventional MSP430 MCU and a CC2420 802.15.4 radio, which makes this approach suitable for a wide range of existing sensor networks.

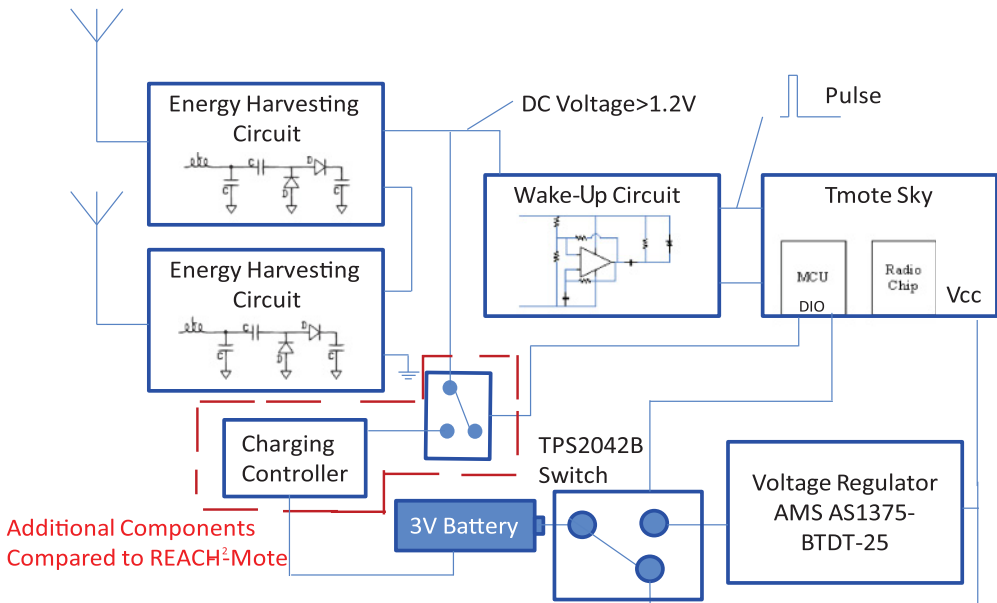


Fig. 23. Block diagram of REACH²-Mote with charging module.

8. CONCLUSIONS

In this article, we presented the design and evaluation of the REACH²-Mote passive wake-up radio sensor node, which utilizes energy harvesting and an efficient wake-up circuit for extended wake-up range. We evaluated our implementation of REACH²-Mote through field tests and compared its performance with that of the first-generation REACH-Mote and WISP-Mote, an existing passive wake-up sensor node. The field test results show that REACH²-Mote can extend the wake-up range to 44 feet compared to a 37-foot wake-up range for REACH-Mote and a 17-foot wake-up range for WISP-Mote. As the communication range of Tmote Sky working indoors is around 160 feet, REACH²-Mote can achieve a wake-up range that is almost one third of the communication range of Tmote Sky. Thus, REACH²-Mote is a passive wake-up sensor node that can be deployed in real WSNs. In addition, as the WuRx of REACH²-Mote requires less battery energy while waiting for a wake-up signal from the WuTx compared to Tmote Sky, more battery energy on REACH²-Mote can be used for either sensing the data or transmitting the data to a base station, eliminating most of the overhead in communication.

To evaluate the performance of REACH²-Mote in a network, we modeled the hardware of REACH²-Mote and evaluated its performance through simulations. We compared the results with that of a network employing WISP-Motes, an implemented active wake-up approach, and a duty cycling approach. The results show that REACH²-Mote outperforms the other approaches in energy efficiency while performing comparable to the other approaches in terms of packet latency and packet delivery performance with higher scalability. We also qualitatively compared the performance of REACH²-Mote with a battery-powered WISP and WISP-Mote to determine the advantages and disadvantages of the different approaches to passive radio wake-up for WSNs.

9. FUTURE WORK

REACH²-Mote operates using broadcast-based wake-up, which means that false wake-ups may be an issue. As the energy harvesting circuit works at 915MHz rather than

2.4GHz, the potential for false wake-ups is less severe, as this is outside the range of common transmissions such as WiFi. To determine the severity of false wake-ups in a real implementation, we ran experiments and found that there were 10 false wake-ups in 1 day. Although this is a relatively low number, a more noisy radio environment may increase the number of the false wake-ups. Therefore, dealing with false wake-ups is a topic for future research. Adding a low-power MCU chip to check an ID sent by the WuTx can reduce most of the false wake-ups. However, the MCU must consume energy from either the battery or the energy harvester; using battery energy to operate the MCU will decrease the node lifetime, whereas using the harvested energy to operate the MCU increases the load to the energy harvester, which will decrease the wake-up range. Thus, further research is needed to determine the best way to enable ID-based wake-up within the context of REACH²-Mote.

As the harvested energy is not used after the sensor node is woken up, this energy can be used to charge the sensor node, which can potentially increase the lifetime of the sensor node. Figure 23 shows a block diagram of this potential sensor node. After the REACH²-Mote node wakes up, the MCU will set a DIO pin to close the switch of the charging circuit. After that, the energy harvested for the wake-up circuit goes to the charging controller. The charging controller works as a voltage regulator, which limits the output voltage to protect the battery. Thus, the energy harvested by REACH²-Mote can be used to charge the battery after the node is woken up. Some future work is needed to further develop this idea.

REFERENCES

- ams AG. 2015. Ams—Designer and Manufacturer of Sensors, Sensor Interfaces, Power Management, and Wireless ICs. Retrieved October, 25, 2015, from <https://www.ams.com/eng>.
- Giuseppe Anastasi, Marco Conti, and Mario Di Francesco. 2008. Data collection in sensor networks with data mules: An integrated simulation analysis. In *Proceedings of the IEEE Symposium on Computers and Communications (ISCC'08)*. IEEE, Los Alamitos, CA, 1096–1102.
- Junaid Ansari, Dmitry Pankin, and Petri Mähönen. 2009. Radio-triggered wake-ups with addressing capabilities for extremely low power sensor network applications. *International Journal of Wireless Information Networks* 16, 3, 118–130.
- He Ba, Ilker Demirkol, and Wendi Heinzelman. 2010. Feasibility and benefits of passive RFID wake-up radios for sensor networks. In *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM'10)*. IEEE, Los Alamitos, CA, 1–5.
- He Ba, Ilker Demirkol, and Wendi Heinzelman. 2013. Passive wake-up radios: From devices to applications. *Ad Hoc Networks* 11, 8, 2605–2621.
- Stefano Basagni, M. Yousof Naderi, Chiara Petrioli, and Dora Spenza. 2013. Wireless sensor networks with energy harvesting. In *Mobile Ad Hoc Networking: The Cutting Edge Directions*, S. Basagni, M. Conti, S. Giordano, and I. Stofmenovic (Eds.). Wiley, 701–736.
- Michael Buettner, Gary V. Yee, Eric Anderson, and Richard Han. 2006. X-MAC: A short preamble MAC protocol for duty-cycled wireless sensor networks. In *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems*. ACM, New York, NY, 307–320.
- CC2420. 2012. CC2420—Proprietary 2.4 GHz—Wireless Connectivity—Description & Parametrics. Retrieved October, 25, 2015, from <http://www.ti.com/product/cc2420>.
- Li Chen, Stephen Cool, He Ba, Wendi Heinzelman, Ilker Demirkol, Ufuk Muncuk, Kaushik Chowdhury, and Stefano Basagni. 2013. Range extension of passive wake-up radio systems through energy harvesting. In *Proceedings of the 2013 IEEE International Conference on Communications (ICC'13)*. IEEE, Los Alamitos, CA, 1549–1554.
- Amre El-Hoiydi, J.-D. Decotignie, Christian Enz, and Erwan Le Roux. 2003. WiseMAC, an ultra low power MAC protocol for the wiseNET wireless sensor network. In *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems*. ACM, New York, NY, 302–303.
- Lin Gu and John A. Stankovic. 2005. Radio-triggered wake-up for wireless sensor networks. *Real-Time Systems* 29, 2–3, 157–182.

- Jeremy Gummesson, Shane S. Clark, Kevin Fu, and Deepak Ganesan. 2010. On the limits of effective hybrid micro-energy harvesting on mobile CRFID sensors. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services*. ACM, New York, NY, 195–208.
- Impinj. 2002. Home—Impinj. Retrieved October, 25, 2015, from <http://www.impinj.com/>.
- Philippe Le-Huy and Sébastien Roy. 2010. Low-power wake-up radio for wireless sensor networks. *Mobile Networks and Applications* 15, 2, 226–236.
- Vincent Liu, Aaron Parks, Vamsi Talla, Shyamnath Gollakota, David Wetherall, and Joshua R. Smith. 2013. Ambient backscatter: Wireless communication out of thin air. In *ACM SIGCOMM Computer Communication Review* 43, 4, 39–50.
- Stevan J. Marinkovic and Emanuel M. Popovici. 2011. Nano-power wireless wake-up receiver with serial peripheral interface. *IEEE Journal on Selected Areas in Communications* 29, 8, 1641–1647.
- moteiv. 2006. Tmote Sky Low Power Wireless Sensor Module. Retrieved October, 25, 2015, from <http://www.eecs.harvard.edu/~konrad/projects/shimmer/references/tmote-sky-datasheet.pdf>.
- MSP430. 2012. MSP430F1611—MSP430F1x—Ultra-Low Power—Description & Parametrics. Retrieved October, 25, 2015, from <http://www.ti.com/product/msp430f1611>.
- Philippe Nain, Don Towsley, Benyuan Liu, and Zhen Liu. 2005. Properties of random direction models. In *Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'05)*, Vol. 3. IEEE, Los Alamitos, CA, 1897–1907.
- Prusayon Nintanavongsa, Ufuk Muncuk, David Richard Lewis, and Kaushik Roy Chowdhury. 2012. Design optimization and implementation for RF energy harvesting circuits. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 2, 1, 24–33.
- Chulsung Park and Pai H. Chou. 2006. Ambimax: Autonomous energy harvesting platform for multi-supply wireless sensor nodes. In *Proceedings of the 3rd Annual IEEE Communications Society on Sensor and Ad Hoc Communications and Networks (SECON'06)*, Vol. 1. IEEE, Los Alamitos, CA, 168–177.
- Aaron N. Parks, Angli Liu, Shyamnath Gollakota, and Joshua R. Smith. 2014. Turbocharging ambient backscatter communication. In *Proceedings of the 2014 ACM Conference on SIGCOMM (SIGCOMM'14)*. ACM, New York, NY, 619–630.
- Chiara Petrioli, Dora Spenza, Pasquale Tommasino, and Alessandro Trifiletti. 2014. A novel wake-up receiver with addressing capability for wireless sensor nodes. In *Proceedings of the 2014 IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS'14)*. IEEE, Los Alamitos, CA, 18–25.
- Nathan M. Pletcher, Simone Gambini, and Jan Rabaey. 2009. A 52 W wake-up receiver with 72 dBm sensitivity using an uncertain-IF architecture. *IEEE Journal of Solid-State Circuits* 44, 1, 269–280.
- Joseph Polastre, Jason Hill, and David Culler. 2004. Versatile low power media access for wireless sensor networks. In *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems*. ACM, New York, NY, 95–107.
- Powercast. 2009. Wireless Power Solutions—Powercast Corp. Retrieved October, 25, 2015, from <http://www.powercastco.com/>.
- Alanson P. Sample, Daniel J. Yeager, Pauline S. Powladge, Alexander V. Mamishev, and Joshua R. Smith. 2008. Design of an RFID-based battery-free programmable sensing platform. *IEEE Transactions on Instrumentation and Measurement* 57, 11, 2608–2615.
- Dora Spenza, Michele Magno, Stefano Basagni, Luca Benini, Mario Paoli, and Chiara Petrioli. 2015. Beyond Duty Cycling: Wake-Up Radio with Selective Awakenings for Long-Lived Wireless Sensing Systems. Retrieved November 12, 2015, from <http://wwwusers.di.uniroma1.it/~spenza/pubs/WRxINFOCOM15.pdf>.
- Emmanuel Munguia Tapia, Stephen S. Intille, and Kent Larson. 2007. Portable wireless sensors for object usage sensing in the home: Challenges and practicalities. In *Ambient Intelligence*. Springer, 19–37.
- TPS2042B. 2012. TPS2042B—Fixed Current Limited Switch—USB Power and Charging Port Controllers—Description & Parametrics. Retrieved October, 25, 2015, from <http://www.ti.com/product/tps2042b>.
- Tijs Van Dam and Koen Langendoen. 2003. An adaptive energy-efficient MAC protocol for wireless sensor networks. In *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems*. ACM, New York, NY, 171–180.
- Bas Van der Doorn, Winelis Kavelaars, and Koen Langendoen. 2009. A prototype low-cost wakeup radio for the 868 MHz band. *International Journal of Sensor Networks* 5, 1, 22–32.
- wisp.wikispaces.com. 2010. WISP Challenge. Retrieved October, 25, 2015, from <http://wisp.wikispaces.com/WISP+Challenge>.
- Han Yan, Jose G. Macias Montero, Atef Akhnounk, Leo C. N. De Vreede, and Joachim Burghartz. 2005. An integration scheme for RF power harvesting. In *Proceedings of the STW Annual Workshop on Semiconductor Advances for Future Electronics and Sensors*. 64–66.

Wei Ye, John Heidemann, and Deborah Estrin. 2002. An energy-efficient MAC protocol for wireless sensor networks. In *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies*, Vol. 3. IEEE, Los Alamitos, CA, 1567–1576.

Pengyu Zhang and Deepak Ganesan. 2014. Enabling bit-by-bit backscatter communication in severe energy harvesting environments. In *Proceedings of the 11th USENIX Conference on Networked Systems Design and Implementation (NSDI'14)*. 345–357.

Received September 2014; revised July 2015; accepted September 2015