Radio-Triggered Wake-Up Capability for Sensor Networks

Lin Gu and John A. Stankovic

Department of Computer Science, University of Virginia
{lingu, stankovic}@cs.virginia.edu

Abstract

Power management is an important technique to prolong the lifespan of sensor networks. Many power-management protocols employ wake-up/sleep schedules, which are often complicated and inefficient. We present power management schemes that eliminate such wake-up periods unless the node indeed needs to wake up. This type of wake-up capability is enabled by a new radio-triggered hardware component. We evaluate the potential power saving in terms of the lifespan of a sensor network application, using experiment data and SPICE circuit simulations. Comparing the result with always-on and rotation-based power management schemes, we find the radio-triggered scheme saves 98% of the energy used in the always-on scheme, and saves over 70% of the energy used in the rotation-based scheme. Consequently, the lifespan increases from 3.3 days (always-on) or 49.5 days (rotation-based) to 178 days (radio-triggered). Furthermore, a store-energy technique can extend operating distance from 10 feet to 22 feet, or even longer if longer latency is acceptable. We also present amplification and radio-triggered IDs which can further enhance performance.

1. Introduction

A sensor network is comprised of a number of low-power devices with sensing and computing capability. In many sensor network systems, the power supply for the network nodes is usually a depletable power source, such as batteries. To increase the lifespan of sensor networks, researchers have designed a number of power management schemes.

Many power management schemes take advantage of the energy saving features of sensor network hardware. For example, the ATmega128 processor, which is designed for embedded systems and is used in the Berkeley Mica2 mote, has six working modes with different energy saving features [10] [1]. In one of the working modes, the processor shuts down all the hardware components except for the memory, a timer, and the interrupt handler, hence energy consumption reduces to less than 1% of the active working mode.

Power management schemes need to control when a network node should enter a high-power running mode and when to enter a low-power sleep mode. The high-power to low-power transition can usually be done with a set of instructions that shuts down hardware components, and the power management scheme may perform this action when certain conditions hold, e.g., there are no events in the system for a long time. The low-power to high-power transition is, however, a tricky problem because the network node has its CPU halted and is unaware of the external events. In many applications, it is desirable to have the network awakened when some events of interest happen. But the network node cannot easily know exactly when events happen.

To solve this problem, many power management schemes require that each network node wake up periodically to listen to the radio channel [12]. When an event of interest happens, some nodes (possibly some sentry nodes) detect the event and send power management messages to the network. All the nodes that were in their listening mode and hear the power management messages stay awake – they do not enter sleep mode. By choosing a good wake-up/sleep schedule, the network may save much energy without compromising the system functionality. The implementation of the wake-up/sleep scheduling often involves a timer that wakes up the CPU via an interrupt.

The wake-up/sleep scheduling approach has some disadvantages. First, the design of a good wake-up/sleep schedule is often application dependent and complicated. Hence, it is hard to design a general power management service based on wake-up/sleep scheduling. For each application, the designer needs to carefully analyze the timing of the system events and tune the scheduling parameters, otherwise some nodes in the network may miss wake-up calls. Second, a good wake-up/schedule often involves collaboration among a group of nodes, or even all the nodes in the network. This often implies that the network needs a time synchronization service. With low-speed processors and radio communication links, to perform high-quality time synchronization in sensor networks is an even more challenging task than in traditional distributed systems. Finally, a common phenomenon is that, in most of the wake-up periods, no event happens and the nodes enter sleep mode again. This means that nodes wake up too often, and it is a waste of energy.

Another approach is to use a low-power stand-by hardware component to watch the environment when the node enters sleep mode. For example, a node can use a stand-by radio-transceiver subsystem to listen to the radio channel when the node sleeps. When the stand-by radio transceiver receives radio signals, it wakes the node up. Otherwise, the node keeps sleeping. PicoRadio has such functionality [16]. It separates the radio hardware for data communication and channel monitoring. A separate low-power radio, called the "wake-up radio", monitors the radio channel and wakes up the node when a power management beacon is received [18].

The disadvantage of a stand-by component is that it uses extra energy. For example, the most important power management related event is a power management message. To keep listening to the power management messages, the node needs to use a radio transceiver in the listening mode. In current sensor network systems, a listening transceiver uses a non-negligible amount of energy.

We propose another approach to power management. The events of interest in a system often contains energy, and the moment when an event happens is exactly the moment when we want the system to enter wake-up mode. So we can potentially use the energy in the event to trigger the transition of the system from sleep mode to wake-up mode.

In our study, the events of interest are wireless power management messages themselves. We use the energy in the radio communication signals to trigger system transitions and, based on this technique, explore designs of novel power management schemes. Hence, we name our power management scheme radio-triggered power management. The reason for focusing on power management messages is that this event also triggers most of the transitions in real applications. Besides, the components to collect and manipulate energy in the form of electromagnetic waves is easily available from the market.

In our radio-triggered power management, a special hardware component – a radio-triggered circuit – is connected to one of the interrupt inputs of the processor. The circuit itself does not have any power supply. The node can enter sleep mode without periodic wake-up. When a power management message is sent by another node (possibly a sentry node) within a certain distance, the radio-triggered circuit collects enough energy to trigger the interrupt to wake up the network node. This is significantly different than using the radio transceiver to listen to messages because a listening radio transceiver requires help from the processor (or a radio subcontroller) to conduct channel monitoring and message parsing and the listening process consumes energy of the node. The radio-triggered circuit, on the contrary, is powered by the radio signals themselves. It is powered off when there is no suitable radio signals and naturally starts working when suitable radio signals arrive. Except for activating the wake-up interrupt, the radio-triggered circuit is independent of any other components on the node.

Our wake-up mechanism is simple, timely, and energy-efficient. Based on it, flexible power management protocols can be constructed to ensure network nodes neither wake up too often nor miss wake-up calls.

The rest of the paper is organized as follows. Section 2 describes the related background. Section 3 presents basic radio-triggered hardware and a basic radio-triggered power management scheme, followed by performance evaluation in Section 4. After Section 5 introduces a more sophisticated radio-triggered hardware component, Section 6 presents two extensions to the radio-triggered hardware – the more powerful amplified radio-triggered circuit and the more selective radio-triggered IDs. Finally, Section 7 concludes the paper.

2. Background and related work

Knowing that energy efficiency is a vital part of performance in sensor networks, people studied the energy consumption of different operations in sensor nodes and tried to find suitable sensor node architectures to reduce power consumption [17].

Though hardware design emphasizes low-power features to reduce the energy consumption in the working state, it is obvious that more energy saving can only come from putting nodes into a deep sleep mode. In one of the earliest deployed sensor networks, researchers at Berkeley and Intel Berkeley Research Lab put the network in sleep mode and used wake-up messages to inform the whole network (800 nodes) to enter working mode when the event of interest happened.

To implement the wake-up, timer-based protocols appeared. In [11], the authors suggest that a sensor node keep awake for 100ms in every cycle of 4 seconds. This results in a duty cycle of 2.5%. Lower duty cycles are sometimes possible. [11] suggested a very low duty cycle of 0.00125%. But it needs special support from the radio transceiver hardware.

Recent power aware protocols employ carefully designed scheduling techniques to save energy. In [20], an energy efficient MAC protocol, S-MAC, is proposed. In S-MAC, nodes listen to the channel and synchronize their schedules with neighbors. In one of the experiments, S-MAC can keep the node in sleep mode for 60-70% of the time. This approximately means the duty cycle is around 30-40%. As another example, the wireless sensor network deployed on the Great Duck Island carefully scheduled all the events in the network to acquire as much sleep time as possible without missing events of interest [14]. This scheduling needs a priori knowledge about the system.

In some sensor networks, a fraction of the network nodes operate in full power mode or have a much higher duty cycle than other nodes in the network. These nodes, called sentries, provide a sensing coverage or communication backbone as required by the application. The system service that controls the selection and rotation of sentries is called a sentry service, which is relevant to power management and energy efficiency. If the application requires a full-time monitoring of possible events of interest, the sentry service selects a number of nodes as sentries to form a monitoring network, which provides the required sensing coverage. The sentry service saves energy by selecting as few sentries as possible. Also, sentries send out power management packets to wake up non-sentries when the events of interest happen [12]. The work presented in this paper can significantly improve sentry services.

Current power management schemes use power management packets or power management radio signals extensively. But they use them only as signals, not as energy sources. On the other hand, the research on passive RFID has been utilizing the energy in radio signals to drive RFID circuitry [19]. But RFIDs usually employ powerful readers to send strong radio signals to excite the transponders. This is not a good solution in wireless sensor networks.

3. Basic radio-triggered power management and basic radio-triggered circuits

In this section, we present the radio-triggered power management in its simplest form, and describe the related hardware support. We first describe a basic radio-triggered power management scheme in Subsection 3.1. To support basic radio-triggered power management, a design of a circuit powered by electromagnetic waves and generating a certain voltage output is described in Subsection 3.2. The output voltage can be used as an interrupt signal. Subsection 3.3 evaluates the effectiveness of the basic circuit design, and subsection 3.4 discusses possible engineering improvements that can make the basic design more efficient.

3.1. Basic radio-triggered power management

Most of the current power management schemes use a carefully designed wake-up/sleep schedule. In the wake-up period, the node listens to incoming radio packets and chooses to go back to sleep or stays awake based on the presence and content of the packets. We call the packets that wake up nodes "wake-up packets".

To keep the network responsive, the wake-up period should happen frequently enough so that the nodes wake up more or less "immediately" when needed. This results in a large portion of the wake-up periods to be in vain – in most cases, the nodes wake up and, after waiting for an interval without hearing a wake-up packet, decide to go back to sleep again.

The radio-triggered power management scheme aims to avoid the useless wake-up periods. In this scheme, nodes stay asleep after they decide to go to sleep. They do not wake up until a special radio signal is sent by another node. The special radio signal wakes up the sleeping node almost instantly. Consequently, we save all the energy spent in previous wake-up and listen intervals.

In the basic radio-triggered power management, the special radio signal is just a wake-up packet, as used in current power management protocols. If supported by hardware, the wake-up packet is sent at a special radio frequency. Other types of radio communication, at a different radio frequency, do not wake up the nodes even if the nodes are within the radio communication range. Note that hardware cost for adding multiple-frequency support is usually fairly low. Many recent low-end radio transceivers support multiple-frequency operations. [4] In some embedded systems where a single frequency is used, e.g., Berkeley Mica2 motes, the radio transceiver can actually support multiple frequencies and some modification to the peripheral hardware setting can enable multiple-frequency operations[2]. The separation of power management radio communication from other types of radio communication by frequencies avoids confusion and facilitates advanced power management features. For example, if the application requires that part of the network be in sleep mode when the rest is exchanging packets, the designer can control the awakened portion of the network by only sending special wake-up signals within the portion to be awakened. In the rest of the network, where no special wake-up signals are present, the nodes can stay in sleep mode even if other types of radio communication, e.g., routing to a base-station via a backbone routing path, exist.

There are four main requirements for radio-triggered power management.

- A node should wake up almost instantly when it receives a wake-up packet. Here "instantly" means that it should be ready to work after the wake-up packet has finished arriving.
- The node should use approximately the same amount of energy in sleep mode as in power management protocols without radio-triggered support.
- The node should not wake up when the event of interest does not happen.
- The node should not miss wake-up calls – when the event of interest happens, the node should not keep sleeping and thus miss the event.

3.2. Design of the basic radio-triggered circuit

In this subsection we present the design of the basic radio-triggered circuit that collects radio energy to trigger an interrupt. Using this circuit, we can implement basic radio-triggered power management. The circuit is called "basic" because the design is simple and, consequently, it can provide limited power. This means it can only operate within a limited distance (10 feet in our example) from the source of the radio signal. While the basic circuit is effective, some improvements can increase the operating distance. In later sections, we present more advanced designs that can operate over a longer distance.

All circuits consume energy. The ideal situation, however, requires that the stand-by hardware use no power supply. Hence stand-by hardware must power its circuit by collecting energy from the environment. For the radio-triggered hardware, its function is triggered when and only when there is a suitable radio signal. Therefore, a natural way of powering the circuit is to collect the energy in the radio signals, which must be present when the function is needed. The energy in the radio signals is not much, but still enough to drive a very small "start" signal that turns on the node which then can be powered on its own power supply.

So the radio-triggered circuit has two essential tasks. First, it needs to collect energy from radio signals, which are electromagnetic waves. Second, it needs to distinguish trigger signals from other radio signals, thus meeting the requirement of not waking up for false positives. In the design of the basic radio-triggered circuit, both tasks are fulfilled by the antenna.

An antenna transforms electromagnetic wave into current. By choosing the length, size, or shape of an antenna, we can make it selective for radio signals – it is strongly activated when an electromagnetic wave of a certain frequency is present, but is weakly activated otherwise. [1]

If a certain radio frequency is present and the antenna is activated, the electrical energy on the antenna can be used to

1 For simplicity, we assume that the radio-triggered circuit has its own antenna, which is tuned to the frequency the radio-triggered circuit uses. Implementations may choose to share a single antenna with normal data communication while maintaining frequency selectivity by techniques such as tunable and switchable passive filters. [15]

power other parts of the circuit, and generate an output voltage signal.

Fig. 1 shows a schematic of a design of the basic radio-triggered circuit. The antenna in the circuit can be a patch, monopole, dipole, or any other antenna, given it provides suitable selectivity and efficiency. In Fig. 1, the antenna reacts to the electromagnetic wave and generates an input voltage, which is rectified by a diode, D_{out}, to produce the output voltage, V_{out}. D_{out} should have a low bias threshold.

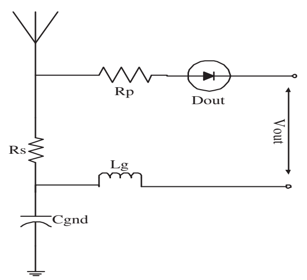

Figure 1. Basic radio-triggered circuit

In a real system, engineers may choose to add more components to the circuit to enhance performance and reliability. In this paper, we keep the circuits simple so that their functionality is easier to understand.

3.3. Effectiveness of the basic radio-triggered circuit

We evaluate the effectiveness of the basic radio-triggered circuit in the context of sensor network platforms and applications.

The radio communication in sensor networks is usually low-power. This is a critical difference from the RFID technology. A typical radio transceiver for a sensor network node transmits radio signals from -10dBm to 10dBm. For example, the Chipcon CC1000 radio transceiver transmits with 10dBm power when the potentiometer setting is 255. [4]

In many scenarios, a sensor network node is able to transmit a small number of control radio signals with high strength. Suppose the node transmits the control signal at 10dBm(10mW), and suppose the distance between the sending node and the basic radio-triggered circuit is 10 feet (about 3 meters). The power received on the basic radio-triggered circuit's antenna, P_r, is given by the following formula.

$$P_r = \frac{P_s G_s G_r \lambda^2}{(4\pi \cdot D)^2} \quad (1)$$

Here P_s is the transmission power. G_s and G_r are both antenna parameters. G_s is the transmission antenna gain, and G_r is the receiving antenna gain. λ is the wavelength of the electromagnetic wave. D is the distance between the transmitting node and the receiving antenna on the basic radio-triggered circuit.

Using Berkeley Mica2 motes as an example, we find the transmission power is 10dBm (10mW). The antenna gain varies depending on the length, size and shape of the antenna. We choose a typical 5/8-wavelength monopole antenna as an example [3]. According to the manufacturer, the gain of the antenna is 8.2dB. Considering environmental issues, we choose 8dB as the transmitting antenna's gain and 6dB as the receiving antenna's gain. They corresponds to numerical values of 6.31 and 3.98 for G_s and G_r, respectively, in the formula above. The wavelength, λ, depends on the radio frequency used in the system. The Berkeley Mica2 mote uses 433MHz radio which has a wavelength of 0.69 meter. With the distance, D, assigned a value of 3 meters, the received power is given by the following equation.

$$P_r = \frac{10 \cdot 6.31 \cdot 3.98 \cdot 0.69^2}{(4 \cdot 3.14 \cdot 3)^2} \mathrm{mW} = 0.084\mathrm{mW}$$

With the received power known, we can calculate the voltage on the receiving antenna, V_{in} by the equation

$$V_{in} = \sqrt{P_r R_s}$$

Here R_s is the internal Resistance of the receiving antenna. Suppose R_s is $7k\Omega$, we have

$$V_{in} = \sqrt{0.084 \cdot 10^{-3} \cdot 7 \cdot 10^3 V} = 0.77V$$

When there is no radio signal at 433MHz, V_{out} is zero. When the radio signal is present, the antenna generates an input voltage source of 0.77V. The input voltage source drives the circuit and produces an output voltage.

Note that we assume a relatively large resistance for R_s. This implies that the load circuit that V_{out} connects to should have a sufficiently large impedance. This is a limit to the load circuit design. But we believe it is not difficult to design a circuit with sufficient impedance in most cases.

We are most interested in knowing the strength of V_{out}, i.e., how high V_{out} is. This depends on the property of the diode, D_{out}. Section 4 presents a quantitative analysis of V_{out} with the SPICE circuit simulator. Suffice for now assume that V_{out} is more than $0.6V$.

We argue that, for sensor network systems, an electrical signal of $0.6V$ is sufficient to trigger an interrupt, which can then trigger various functionalities as desired. There is no theoretical limit on how "high" the high voltage must be, given that the interrupt input circuitry does not confuse it with a low voltage. Traditionally, computer designers choose to use a relatively high voltage level, e.g., 5V, to represent a logic "high" in order to maintain a wide "dead band" between the "high" and "low" to resist noise. In our system, when V_{out} grows beyond 0.6V, we consider it as a valid situation to trigger the interrupt. Put in another way, any noise that the system is designed to handle should not drive V_{out} up to a fraction of 0.6V. Hence, we only need a gap wide enough to avoid unstable behavior of the interrupt circuit. This is a lower requirement than in traditional computers. Furthermore, many recent low-power embedded systems choose to operate on very low voltage to reduce power consumption [5] [7]. Some of them can operate on 1.2V power supply and some accept as "high" a 1V voltage. Using comparators, the interrupt voltage threshold can be reduced to even much lower than 0.6V, as long as the wake-up signal can still be distinguished from noise. The next subsection discusses using comparators to accomplish a very low voltage threshold. With such low-power technology, interrupt devices that handle 0.6V input are definitely feasible in the near future.

Therefore, the radio-triggered circuit is activated when radio signals are present and triggers an interrupt. We call it a radio-triggered interrupt. In the Berkeley Mica2 mote, the wake-up logic is implemented as an interrupt caused by a timer. With slight modification on the wake-up circuitry, the wake-up logic can work with the radio-triggered interrupt.

Therefore, we have proven the effectiveness of the basic radio-triggered circuit. When a suitable radio frequency exists, the circuit generates a signal strong enough to drive an interrupt. Specifically, it is possible to wake up a sensor network node in sleep mode with this circuit.

3.4. Improve the basic radio-triggered circuit

The calculation in the previous subsection assumes that the transmitting node is 10 feet from the radio-triggered circuit. In some applications this distance is acceptable. However, it is obviously desirable to extend this distance. With some engineering effort, this is possible.

Equation 1 indicates that the received power, P_r, increases as the transmission power P_t or antenna gains G_s and G_r increases. In sensor networks, usually we avoid increasing the transmission power unless necessary, because of concerns over power consumption and radio interference. The antenna gains may be increased by using more sensitive antennas. With other parameters unchanged, Fig. 2 shows the relation between the distance D and gain of the receiving antenna G_r.

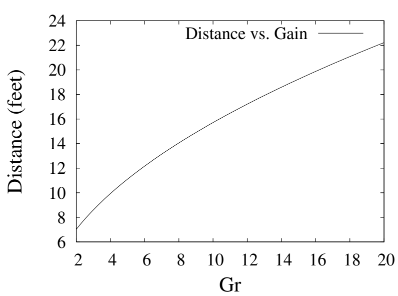

Figure 2. Distance and antenna gain

The antenna gain cannot increase infinitely. Both packaging issues and physical laws limit the sensitivity of antennas. So the increase of D while increasing G_r has a limit. A more promising way of increasing D is to decrease V_{out} needed to drive the load circuit. Since the main purpose of the radio-triggered circuit is to trigger interrupts, we can lower the interrupt threshold voltage. This requires redefinition of the electrical specification of the interrupt logic, and some modification of the interrupt-related circuit. So it involves a certain amount of engineering effort and production cost.

Engineering effort is needed in the modification of the interrupt related circuit. But the effort is likely to be an easy one. First, we do not have to modify the circuit of all types of interrupts. We can choose to only modify the interrupt circuit for the wake-up logic. Second, we need not modify the internal circuit of the interrupt handler. All we need is to use a small voltage to trigger the interrupt. So we can add a biased voltage on to V_{out} to make it a suitable interrupt source – apply a constant voltage supply to the output circuit so that V_{out} instantly exceeds the threshold when the wake-up packet triggers operation of the circuit. This makes the interrupt input voltage high enough to trigger the interrupt without modifying any circuit beyond the input connector. A more promising way is to add a comparator before the interrupt input. The comparator can serve as a device to measure the voltage output from the radio-triggered circuit and generate a much higher voltage when it exceeds a threshold. Finally, if the interrupt circuit needs a big gap to determine whether its input is 1 or 0, we can use transistors (or even amplifiers) to make the small radio-triggered signal strong enough to enable the interrupt circuit to clearly discern the input. Therefore, it is feasible to lower the interrupt voltage threshold to make it work with small radio-triggered signals.

Note that the comparators and amplifiers consume energy. One of the requirements of the radio-triggered circuit is to maintain approximately the same energy consumption level in sleep mode. We thus need to select comparators or amplifiers to make sure their stand-by current does not noticeably increase the overall power consumption of the system in sleep mode. For example, Dallas Semiconductor's MAX9119 comparator works at a current of 350nA. Compared to the 100μA sleep mode current we observed, it increases the energy consumption by 0.3%, which we consider negligible. Furthermore, with radio-triggered hardware, we can stop the wake-up timer which must keep running in sleep mode with existing power management schemes. This reduces energy consumption. Overall, the node should consume approximately the same amount of energy in sleep mode as with other power management schemes.

The production cost includes the add-on circuit for handling low-threshold interrupts and, possibly, the complexity of maintaining two sets of electrical specifications for interrupts. Yet the cost is worthwhile since as we can see, in Section 4, the performance gain is significant.

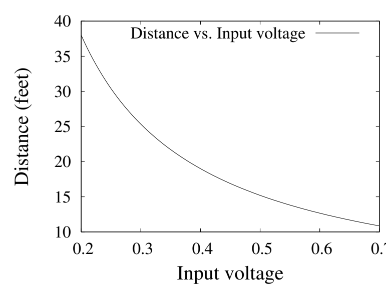

Figure 3. Distance and input voltage

With the output voltage V_{out} lowered, the input voltage V_{in} can also be lowered, hence the distance D increased. Fig. 3 illustrates the relation between the distance D and voltage output V_{in} required to trigger the interrupt. For example, if the input voltage was 0.25V, we can increase the distance between nodes to 30 feet (a 3 times improvement).

4. Performance evaluation of the basic radio-triggered circuit

We first evaluate the output voltage V_{out} of the basic radio-triggered circuit in Subsection 4.1, then evaluate the potential power-saving performance in terms of the prolonged system lifetime in Subsection 4.2.

4.1. Evaluation of the output voltage

In Subsection 3.3, we mentioned that the output voltage, V_{out}, will be more than $0.6V$. In this subsection, we use the SPICE simulator to determine the accurate value for V_{out}. SPICE is a circuit level simulator developed by Berkeley. It uses numerical methods to analyze the circuit and can trace the changes in the circuit against time.

For simulation purposes, the circuit is redrawn as shown in Fig. 4. The resistors, inductor, and diode are given reasonable parameters to make sure that they represent available products on the market. Specifically, the resistor R_s is 7000Ω and the inductor L_g is 2nH. As mentioned in Subsection 3.2, the diode D_{out} should have low bias voltage. Zero-bias diodes have been studied [13] and are available at an affordable price. To model diodes with low bias voltage, we specify the emission coefficient of diode D_{out} to be 0.3. A $800k\Omega$ resistor, R_l, is used to simulate the workload.

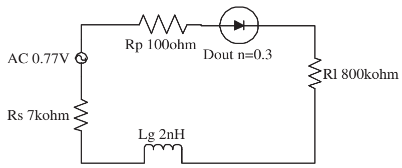

Figure 4. Circuit used in SPICE simulation

In Fig. 4, the voltage drop on R_l represents the output voltage V_{out}. The circuit simulation shows that V_{out} is $0.62V$. This supports the evaluation of the effectiveness of the basic radio-triggered circuit in Subsection 3.3.

4.2. Evaluation of the potential power saving

In this subsection, we evaluate the potential power saving in a representative sensor network application, and calculate the increase of the lifespan of the network when the basic radio-triggered power management is used.

We choose a tracking application as the representative application because tracking applications are well studied, practically deployed, and technically challenging and hence qualify as a good representative of sensor network applications.

We suppose a specific scenario for a sensor network with a tracking application. All the parameters and assumptions used in this scenario are based on our observation in experimental systems based on Berkeley Mica2 motes unless explicitly noted otherwise [9]. In this scenario, there are totally 1000 nodes randomly deployed. There are 10 events to be detected every day, and each event lasts 2 minutes. Also suppose that the each network node uses two 1600mAh AA batteries. [2] The average voltage of the AA batteries during the lifespan of the network is $1.45V$. We further assume the network nodes have two working modes – wake-up mode and sleep mode. The wake-up mode is the usual working mode with all the functional units ready to work, and the average wake-up mode current is 20mA. In sleep mode, a node shuts down all its components except the memory, interrupt handler, and the timer. The sleep mode current is $100\mu\text{A}$. When a network node changes from sleep mode to wake-up mode, there is a surge current on the average of 30mA for 5ms.

We compare the energy consumption of three schemes – always-on, rotation, and radio-triggered. The always-on scheme is a network without power management – all the nodes are on throughout the lifespan of the network. The rotation scheme is a wake-up/sleep schedule based scheme much like many current power management protocols. And the radio-triggered scheme uses the basic radio-triggered power management. We assume that the application requires the entire network to wake up within 4 seconds when an event occurs in order to provide more accurate tracking. We assume that power management message exchange in the network takes 200ms. [3]

First, let's look at the always-on scheme. With this scheme, the network nodes last until their batteries deplete. So the network lifespan L_{on} is

$$L_{on} = 1600\text{mAh}/20\text{mA} = 80\text{h} = 3.3\text{days}$$

// Why 3.3 days? fewer than Sudha's result ///////

Second, let's evaluate the rotation scheme. With this scheme, when there are no events, a node enters sleep mode for a period of time and then wakes up for another period of time. The length of the periods depends on how responsive the application requires the network to be and how reliable the network can deliver wake-up packets. As the network needs to respond to an event within 4 seconds, the maximum length of the sleep period is 4 seconds. Furthermore, if the network can reliably deliver wake-up packets within 200ms, the wake-up period can be as short as 200ms. Consequently, the optimal wake-up/sleep schedule is to have each node stay in sleep mode for 4 seconds (4000ms) and then stay in wake-up mode for 200ms. Thus one wake-up/sleep cycle is comprised of 4200ms. The 4-second time in sleep mode includes the 5ms sleep-to-wake-up transition time. This schedule corresponds to a duty cycle of 4.88%.

Each day has 1440 minutes, or 86400 seconds, among which a node spends 20 minutes, or 1200 seconds, on tracking. In the remaining 85200 seconds, a node enters wake-up and sleep cycles alternatively. Obviously, the total number of 4200ms wake-up/sleep cycles each day is $85200\text{s}/4200\text{ms} = 20286$. The energy that a non-sentry node consumes in one wake-up/sleep cycle is $((4000\text{ms} - 5\text{ms}) \cdot 100\mu\text{A} + 5\text{ms} \cdot 30\text{mA} + 200\text{ms} \cdot 20\text{mA}) \cdot 1.45V \cdot 2 = 13193.55\mu J$. Therefore, a nodes consumes $13193.55\mu J \cdot 20286 = 267.6J$ each day when there is no event of interest. For each event, there is one

2 A network node may only effectively use part of the energy in its batteries because the node may malfunction when the voltage drops below a certain level. Therefore, a 2000mAh battery may only effectively provide a capacity of 1800mAh. For simplicity, we assume that the network nodes are actually using 2000mAh batteries and the 1600mAh capacity is effectively available.

3 This is an assumption that differs from what we observe in real systems. Usually, there is a delay between the event and the first detection of the event. And the power management message can be delayed due to collision and multi-hop communication. We use this assumption for simplicity.

sleep-to-wake-up transition, and a node stays awake for 120 seconds. Hence, the energy consumption per event is $(5\text{ms} \cdot 30\text{mA} + 120\text{s} \cdot 20mA) \cdot 1.45V \cdot 2 = 6.96J$. With 10 events, this adds up to $69.6J$ each day. Adding the two parts of energy consumption, the total energy used per day is $267.6J + 69.6J = 337.2J$. Consequently, we can calculate the lifespan, L_{ro}, of the network as follows.

$$\begin{aligned} L_{ro} &= \tfrac{1600\text{mAh} \cdot 2 \cdot 1.45V}{337.2J} \\ &= \quad 49.5\text{days} \end{aligned}$$

Now let's evaluate the network lifespan with basic radio-triggered power management. We denote the lifespan as L_{rt}. With radio-triggered power management, nodes do not wake up until an event happens. There are totally $L_{rt} \cdot (10event/1day)$ events. Each day there are approximately 10 events which correspond to $10 \cdot 2min = 20min$, or 1.39% of the time in a day. For each event, a node is in the wake-up state for 2 minutes (120000ms), with 5ms in the 2 minutes experiencing higher surge current. So the energy consumed for each event is

$$\begin{aligned} &(120000 - 5)\text{ms} \cdot 20\text{mA} \cdot 1.45\text{V} \cdot 2 \\ &+5\text{ms} \cdot 30\text{mA} \cdot 1.45\text{V} \cdot 2 \\ &= 6960145\text{ms} \cdot \text{mA} \cdot \text{V} \end{aligned}$$

Consequently, L_{rt} can be solved by the following equation.

$$\begin{aligned} &1600\text{mAh} \cdot 2 \cdot 1.45\text{V} = \\ &\quad L_{rt} \cdot (1 - 1.39\%) \cdot 100\mu\text{A} \cdot 1.45\text{V} \cdot 2 + \\ &\quad L_{rt} \cdot (10event/1day) \cdot 6960145\text{ms} \cdot \text{mA} \cdot \text{V} \end{aligned} \quad (2)$$

Solving the equation, we get $L_{rt} = 2128h = 178\text{days}$.

Note that in these calculations, we do not consider the sentry service. The reason is that there are a variety of sentry devices with quite different implications on power management schemes, and usually sentries constitute only a small portion of the network.

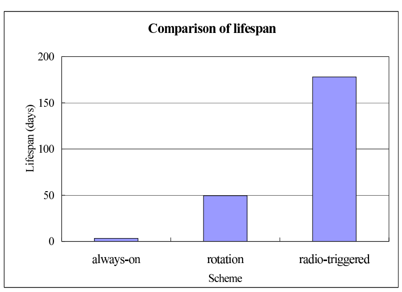

Figure 5. Comparison of the network lifespan of three power management schemes

Fig. 5 shows the comparison of the power saving in terms of the lifespan of the networks. The basic radio-triggered power management has a lifespan 2.6 times longer than the rotation-based scheme. If we calculate the average power of these schemes, the always-on scheme's average power is 58mW, the rotation scheme is 3.9mW, and the radio-triggered scheme is 1mW. The radio-triggered scheme saves 74% of the energy that the rotation scheme uses.

5. Store-energy radio-triggered circuit

The radio-triggered circuit described in Subsection 3.2 has a limited operating distance of 10 feet with current Berkeley motes. We discussed improvements to the circuit in Subsection 3.4. In this subsection, we present a more sophisticated design which uses a capacitor to store energy and enables the radio-triggered hardware to operate at a much longer distance up to 22 feet.

The major problem when operating at a long distance is the low energy level received by the receiving antenna. If the energy received is not sufficient to instantly drive an interrupt, we can still store energy for a period of time and then trigger the interrupt when the energy has accumulated to a certain level.

The schematic of the circuit is shown in Fig. 6. We use a step-up voltage transformer to increase the input voltage on the antenna. Suppose the voltage waveform on the antenna is a sinusoidal curve. The transformer increases the peak voltage of the sinusoidal curve to a certain level. Increasing the voltage level makes it easier to reach the voltage threshold required for triggering an interrupt. The transformer component does not necessarily need its own power supply, since it can be as simple as two coils coupled with each other. There is some energy loss during the voltage transforming process. But modern transformer technology has reached a point where this energy loss is negligible for our purpose – energy efficiency of more than 90% is common. Furthermore, the circuit is to store energy over a relatively long period of time (on the order of several milliseconds). The lost energy only makes the energy-storing process slightly slower, but does not affect its functionality.

We use a capacitor, C_{se} to store energy. When the antenna receives radio signals, the flow of current accumulates energy on C_{se} and the voltage across the capacitor increases accordingly.

Again, we use SPICE to simulate the circuit and evaluate how quickly the circuit responds to radio signals. Fig. 7 plots the change of the output voltage, V_{out}, versus time.

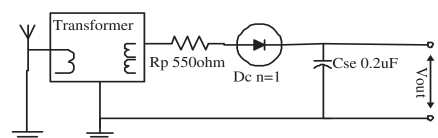

Figure 6. Store-energy radio-triggered circuit

The 433MHz wake-up radio signal begins at time 0. As we can see from Fig. 7, the output voltage V_{out} increases quickly in the first 1ms. In this period, the C_{se} is quickly accumulating energy. After 1ms, V_{out} grows very slowly. The simulation results show that V_{out} reaches 0.6V at around 2.8ms.

The store-energy design has a latency when responding to the wake-up signal. As shown in Fig. 7, the latency is about 2.8ms, in our example. In most applications, we believe this delay is acceptable.

To compute the operating distance of the store-energy radio-triggered circuit, we examine the energy accumulated

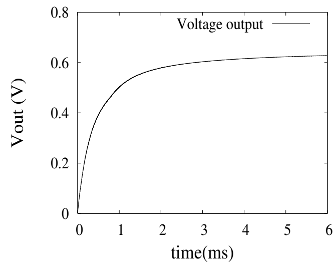

Figure 7. Voltage output of the store-energy radio-triggered circuit.

on the capacitor C_{se}. As mentioned above, an interrupt can be triggered by a signal higher than 0.6V voltage. If we use 0.6V as a threshold, and requires that C_{se} be sufficiently charged so that V_{out} is no less than 0.6V, we can compute the amount of energy needed to charge C_{se} as follows.

Using the capacitor energy formula $E = CV^2/2$, we know C_{se} has 36nJ energy when it reaches 0.6V. Note that the incoming power is given by Equation 1. Assuming that 80% of the received energy can be used to charge C_{se}, we have the following equation about the relation between received energy and the energy on C_{se}.

$$\frac{P_s G_s G_r \lambda^2}{(4\pi \cdot D)^2} \cdot 2.8ms \cdot 80\% = 36nJ \qquad (3)$$

With parameters assigned the same values as specified in Subsection 3.3, we solve the equation and find the distance D is 22 feet.

If we want a longer distance, we may choose to increase the step-up rate of the transformer. This enables the circuit to still charge C_{se} to a certain level, though it takes longer time, i.e., the latency is longer. Fig. 8 shows the relation between distance and latency. So as much as 30 feet (three times of the 10-foot operating distance with the basic radio-triggered circuit) can be accomplished if we can accept 5ms latency.

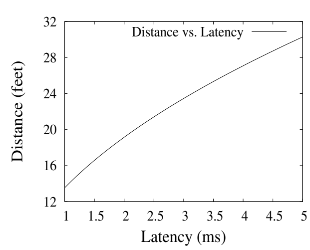

Figure 8. Distance and latency.

6. Advanced radio-triggered hardware

This section presents two ways to extend the basic radio-triggered hardware component to support better performance and stronger selectivity.

6.1. Amplified radio-triggered circuit

All the designs described above avoid using a depletable power supply – the wake-up radio signals are the only source of energy. If we relax this requirement, we can construct more powerful hardware components with longer operating distance and more sophisticated functions, such as false-alarm reduction and the ability to control more complicated hardware than the interrupt circuitry.

Fig. 9 shows the schematic of a radio-triggered circuit with amplification. An amplifier is added before the output of the radio-triggered circuit. The amplifier has internal power supply so it can generate an output signal that has more power than the received wake-up signals. By choosing an amplifier that has a suitable amplification capability, the amplified radio-triggered circuit generates a voltage output signal powerful enough to trigger an interrupt, or conduct other functions of choice. For example, it may perform some local computation to filter out false alarms, turn on some sensors, or initialize some components before the main processor resumes working.

The disadvantage of the amplifier circuit design is that the amplifier uses energy when there are no wake-up signals. To maintain a certain level of responsiveness, the amplifier must stay on and consume energy when there are no events of interest. This leakage energy in the idle state is what radio-triggered hardware is designed to avoid. Undesirable as it is, the leakage energy of an amplifier is often not a serious problem. For example, Texas Instruments' TLV240x amplifier typically draws 880nA, which is about 0.8% of the sleep mode current we observed [6]. Consequently, the amplified radio-triggered hardware can potentially provide more powerful functionality without compromising the energy efficiency. This makes it a useful option when the wake-up logic needs to support sophisticated functions. The amplifier circuit can also increase the operating distance of the radio-triggered circuit.

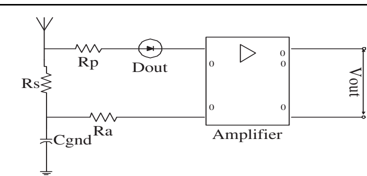

Figure 9. Amplified radio-triggered circuit

6.2. Radio-triggered ID

Radio-triggered circuits can serve as a basic structure to construct more advanced hardware components. In this subsection, we explore one such construction – a radio-triggered ID component.

In Section 3, we mentioned that the wake-up packets should be sent at a special radio frequency so that normal radio communication does not unintentionally wake up the idle part of the network. It is, however, impossible to prevent a node from waking up when the node is within the range of the wake-up signals, even though the node is not actively involved in the current event of interest. This means that selectivity provided by using a special wake-up radio frequency is still not strong enough to support fine granularity power management.

In this subsection, we explore the use of multiple frequencies to enhance the selectivity of radio-triggered hardware component. Physically, multiple frequencies can coexist and serve different circuits working at different frequencies. Technically, a node can employ multiple transceiver to transmit radio signals at different frequencies simultaneously. [4] Furthermore, recent progress in transceiver design has made it possible to transmit at multiple frequencies simultaneously from one transceiver and consume less energy than the sum of energy that would be used when transmitting signals at those frequencies separately [8].

Applying a multiple-frequency technique, we can design an ID system, radio-triggered ID (RTID), determined by the combination of a number of frequencies. Based on this RTID, we can implement better selectivity in a power management service, and, for some mission critical applications, it can even replace the traditional ID system because RTID is more robust and secure.

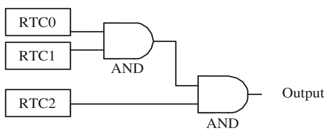

Figure 10. Radio-triggered ID component

Fig. 10 illustrates the design of the RTID component with three radio-triggered circuits, RTC0 - RTC2, that operate on distinct frequencies. Let's denote the corresponding frequencies Freq0 - Freq2. When and only when Freq0 is present, is RTC0 activated. If RTC0 is activated, it generates a positive output which is interpreted as "1" for the input of the first AND gate. Similarly, RTC1 generates "1" to the input of an AND gate when Freq1 is present, RTC2 generates "1" to the input of another AND gate when Freq2 is present. The output of the RTID component is "1" when and only when all Freq0 - Freq15 are present. The output of the RTID component can then be connected to the wake-up interrupt input, or a register bit that indicates whether the ID is matched. If the output is connected to the wake-up interrupt, the node is awakened when and only when all Freq0 - Freq2 are present. Similarly, if the output is connected to a register bit, the bit is set when and only when all Freq0 - Freq2 are present.

If there are 6 distinct frequencies in the network and the RTID component in Fig. 10 is used, the network can have 20 unique frequency combinations, and these frequency combinations can be assigned to nodes in the network. These combinations can serve as group IDs or locally unique IDs of network nodes. We call these combinations RTIDs.

When a node, say, node A, wants to wake up another node, say, node B, node A sends out wake-up signals simultaneously at 3 distinct radio frequencies corresponding to node B's RTID. This wakes up node B but not any other nodes with a different RTID.

To approximately assess the performance gain using RTIDs, let's examine a simple example. Suppose the network has a density that each node has 8 neighbors in its communication range. A node sends out a wake-up packet to wake up one specific neighbor. This wake-up packet also wakes up other neighbors with the same RTID. As a simple case, let's first assume that no sophisticated RTID assignment scheme is used and the RTIDs are evenly randomly distributed. In this case, the probability that exactly no other neighbor is awakened is 0.698, and the probabilities that exactly one, two, ..., seven other neighbors are awakened are 0.257, 0.041, 0.004, 0.0002, 0.000006, 1.04e-07, and 7.8e-10. If the energy consumed in one wake-up period, before the node enters sleep mode again, is E, the expected energy consumed by the neighbors on receiving one wake-up packet is

$$\begin{aligned}&0.698 \cdot E + 0.257 \cdot 2 \cdot E + 0.041 \cdot 3 \cdot E + \\ &0.004 \cdot 4 \cdot E + 0.0002 \cdot 5 \cdot E + 0.00006 \cdot 6 \cdot E + \\ &1.04 \cdot 10^{-7} \cdot 7 \cdot E + 7.8 \cdot 10^{-10} \cdot 8 \cdot E \\ &= 1.35 \cdot E\end{aligned}$$

Without RTID, a wake-up packet wakes up all the neighbors. In this case, the expected energy consumed is $8 \cdot E$. Therefore, the RTID technology enhances the selectivity of a radio-triggered wake-up, and reduces 83% of the energy consumption.

In the above example, RTIDs are randomly assigned to nodes. This corresponds to a random addressing scheme. With a more advanced addressing scheme that assigns different RTIDs to neighbors in a one-hop communication area, [5] we can achieve more energy saving. When there are no duplicate RTIDs among one's neighbors, a wake-up packet wakes up exactly one node. The energy consumed is E, which is 26% less than that with the random scheme and 88% less than that with no RTID support.

If the network has a higher density or needs finer selectivity, it is technically straightforward to add more frequencies and radio-triggered circuits to the node. For example, if we use 8 frequencies and 4 radio-triggered circuits on a node, the RTID system can provide 105 unique IDs.

An important feature of RTID is jamming-resistant. Traditional ID systems in network protocols embed the destination ID in the packet header. The receiving node parses the incoming packet and reads the ID in the header to determine the addressee. If there is radio jamming, the receiving node may not correctly read the ID, thus be unable to realize that the message is for itself. The RTID, on the other hand, is less vulnerable to radio-jamming. When node A uses RTID to identify that the recipient of the packet is node B, jamming radio signals can only add more energy to the frequencies used. Hence, node B is still able to be awakened. This meets the requirement that a node should not miss a wake-up

4 Antenna configuration is a concern with multiple frequencies. For simplicity, we assume that multiple transceivers (or one multi-band transceiver) use multiple antennas. Multiple antennas are likely to increase the size and packaging complexity of a node. Implementations may choose to share some antennas given that the antenna system provides adequate signal quality.

5 This is equal to a coloring problem. Whether or not this is feasible depends on application requirements. Technically, this type of advanced addressing scheme can be accomplished by frequency-tunable radio component or controlled deployment.

call. Note that a node may be awakened by the radio jamming signals if the signals are transmitted at the frequencies corresponding to its RTID. This is a "false positive" situation. In this case, a node can listen to the channel to estimate the network status. If it does not receive a correct packet in a period of time, node B can easily realize that jamming has possibly happened and try to solve the problem.

There is, obviously, hardware cost involved with RTID. For a component shown in Fig. 10, three radio-triggered circuits are needed. Besides, three radio transceivers or one radio transceiver capable of transmitting multiple-frequency signal is needed. Therefore, RTID might be used in mission critical systems where robustness and fine power management granularity are more important than cost.

7. Conclusions

In this paper, we present radio-triggered hardware and explore its applications in power management. By extracting energy from the radio signals, the radio-triggered hardware provides wake-up signals to the network node without using internal power supply. If adequate antennas are used, this wake-up mechanism does not respond to normal data communication signals and thus does not prematurely wake up the network node. More importantly, the radio-triggered circuit becomes energized and generates a wake-up signal exactly when suitable wake-up signals arrive. Because of the zero stand-by power consumption and timely wake-up capability, radio-triggered hardware can be used to construct highly flexible and efficient power management schemes.

Based on our observations in real-world experimental systems, we assess the performance of radio-triggered power management. Simulations and calculations show that power management schemes based on radio-triggered hardware can significantly enhance energy efficiency. It saves 98% energy used in a system without power management and over 70% energy used in a system with typical existing power management schemes.

We also explore extensions to radio-triggered hardware to enable longer operating distance and stronger selectivity. A store-energy radio-triggered circuit accumulates energy extracted from radio signals for a period of time before triggering the wake-up interrupt. This enables a longer operating distance. Relaxing the zero stand-by power requirement, an amplified radio-triggered circuit can also increase the operating distance. An RTID component is a combination of radio-triggered circuits at different frequencies. It features stronger selectivity and can further improve energy efficiency. Based on the radio-triggered hardware, efficient power management protocols can be constructed to ensure network nodes neither wake up too often nor miss wake-up calls.

8. Acknowledgments

This work is supported in part by NSF grant CCR-0098269 and CCR-0325197, the MURI award N00014-01-1-0576 from ONR, and the DAPRPA IXO offices under the NEST project (grant number F336615-01-C-1905).

Also thanks to Yingmin Li, Zhiyang Liu, Zhijian Lu, Ting Yan, Tian He, Anthony Wood, and the Usenet newsgroup community whose help makes this work possible.

References

[1] Atmel atmega128. *http://www.atmel.com*.

[2] Berkeley mote hardware. *http://webs.cs.berkeley.edu/tos/hardware/*.

[3] Chipcon Application Note AN003 – SRD Antennas. *http://www.chipcon.com/files/AN_003_Antennas_1_1.pdf*.

[4] Chipcon CC1000 Datasheet, Chipcon Inc. *http://www.chipcon.com/files/CC1000_Data_Sheet_2_1.pdf*.

[5] Holtek HT47C20L. *http://www.holtek.com.tw*.

[6] TI's nanopower amplifiers. *http://www.ti.com/*.

[7] Xemics xe8000. *http://www.xemics.com*.

[8] H. Hashemi and A. Hajimiri. Concurrent multi-band low-noise amplifiers: Theory, design, and applications. *IEEE Transactions on Microwave Theory and Techniques*, 50(1):288–302, Jan. 2002.

[9] T. He, S. Krishnamurthy, J. A. Stankovic, T. Abdelzaher, L. Luo, R. Stoleru, T. Yan, L. Gu, J. Hui, and B. Krogh. An energy-efficient surveillance system using wireless sensor networks. *MobiSys'04*, June 2004.

[10] J. Hill and D. Culler. Mica: A wireless platform for deeply embedded networks. *IEEE Micro.*, 22(6):12–24, Nov/Dec 2002.

[11] J. Hill and D. Culler. A wireless embedded architecture for system level optimization. *UC Berkeley Technical Report*, 2002.

[12] J. W. Hui, Z. Ren, , and B. Krogh. Sentry-based power management in wireless sensor networks. *The 2nd International Workshop on Information Processing in Sensor Networks (IPSN'03)*, 2003.

[13] J. Joe, M. Chia, A. Marath, and C. Ang. Zero bias schottky diode model for low power, moderate current rectenna. *DETS'97 Proc.*, 1997.

[14] A. Mainwaring, J. Polastre, R. Szewczyk, D. Culler, and J. Anderson. Wireless sensor networks for habitat monitoring. *ACM International Workshop on Wireless Sensor Networks and Applications*, Sept. 2002.

[15] C. T.-C. Nguyen, A.-C. Wong, and H. Ding. Tunable, switchable, high-q vhf microelectromechanical bandpass filters. *Digest of Technical Papers, IEEE ISSCC.*, Feb. 1999.

[16] J. Rabaey, J. Ammer, T. Karalar, S. Li, B. Otis, M. Sheets, and T. Tuan. Picoradios for wireless sensor networks: The next challenge in ultra-low-power design. *Proceedings of the International Solid-State Circuits Conference*, Feb. 2002.

[17] E. Riedy and R. Szewczyk. Power and control in networked sensors. *http://webs.cs.berkeley.edu/tos/papers/cs294-8.pdf*, May 2000.

[18] J. Silva., J. Shamberger, M. J. Ammer, C. Guo, S. Li, R. Shah, T. Tuan, M. Sheets, J. M. Rabaey, B. Nikolic, A. Sangiovanni-Vincentelli, and P. Wright. Design methodology for picoradio networks. *Proceedings of the Design Automation and Test in Europe*, March 2001.

[19] V. Stanford. Wireless sensor networks for habitat monitoring. *Pervasive Computing, IEEE*, 2(2):9–14, April-June 2003.

[20] W. Ye, J. Heidemann, and D. Estrin. An energy-efficient MAC protocol for wireless sensor networks. *Proceedings of the IEEE Infocom*, June 2002.