System-Level Characterization of Single-Chip Radios for Wireless Sensor Network Applications

Yuxin He
Department of Electrical Engineering
Northern Arizona University
Flagstaff, AZ 86011 USA

Paul G. Flikkema
Department of Electrical Engineering
Northern Arizona University
Flagstaff, AZ 86011 USA
Email: paul.flikkema@nau.edu

***Abstract*—We tested three single-chip ISM-band radio transceivers with an eye toward communication system-level performance in wireless sensor network applications. We compared the performance of an older-generation chip (Texas Instruments TRF6900A) with two recent chips (Texas Instruments CC1100 and Analog Devices ADF7020). To understand packet-level sensitivity and the potential gains of forward error-control coding, we evaluated the packet error rate and bit error rate statistics as a function of received signal power at two data rates. We characterized features—automatic frequency control, digital received signal strength indication (RSSI), and digital transmit power control—of the two newer chips. We also modeled and evaluated their current consumption and energy efficiency (in terms of per-bit energy use). We found that the new-generation radio chips have significantly higher integration and overall performance, and that radio chip selection for wireless sensor node design is application-dependent. Our work can be used as a design pattern for further testing of additional and future radios for wireless sensor network nodes.**

I. Introduction

The two factors that dominate wireless sensor network (WSN) performance are link quality and energy consumption. Link quality determines the available data rate and channel reliability, which in turn drive channel code design, network throughput and the dynamics of network topology. Each of these also couples to energy efficiency, and it is well-known that radio communication is often the sensor node function with the highest energy consumption and hence the largest impact on battery lifetime. Yet little is known about the systems-level performance characteristics of single-chip ISM-band radios (SIRs) that may, at least in many applications, dominate the network and energy-consumption performance of WSNs.

Our goal was to bridge the gap between data-sheet specifications and the performance observed by many WSN investigators at the LLC or MAC level [1], [2], [3], [4]; there is a large gap between BER curves and goodput/delay characteristics that view radios as a black boxes underlying a MAC protocol. We were especially interested in performance measures that will help systems-level designers in evaluation of error detection, correction, and recovery strategies, transmit power control algorithms, attenuation-based localization schemes, and characterization of communication energy use that can inform models of energy consumption. Among the questions we attempted to answer were: What is the sensitivity of modern SIRs as measured their ability to detect and synchronize to incoming packets? What precision and accuracy are achievable in the measurement of link losses? What is the energy cost of successfully communicating a bit of information?

This paper's organization reflects our objectives. We first give an overview of three SIRs, followed by a summary of our methods. Then we describe test results, including physical-layer link quality in terms of sensitivity as a function of received power and frequency offset, power control and RSSI measurement, and power consumption

Y. He is currently with Qualcomm, Inc., San Diego, CA, USA.

II. Single-Chip ISM-Band Radios

Driven by a diverse set of large and growing markets, SIRs are available from several manufacturers, including Analog Devices, Nordic Semiconductor, Semtech, and Texas Instruments.

We chose three SIRs that operate in the 902-928 MHz US ISM band, since it provides sufficient bandwidth and tolerable propagation losses for the environmental sensing application of interest to us. The 2.4-2.4835 GHz ISM band is a global standard and has a large available bandwidth, but is compromised by larger attenuation losses in outdoor (particularly forest and grassland) environments. We tested the Texas Instruments TRF6900A [5], the Texas Instruments (formerly Chipcon) CC1100 [6], and the Analog Devices ADF7020 [7]. We refer to these as the TRF, CC, and ADF in what follows. Their performances should be similar in the 868-870 MHz European ISM band.

A summary of the manufacturer's specifications for the three SIR's is shown in Table I. The TRF has been in production since at least 2001, and represents a previous generation of low-cost chips. It is performs frequency acquisition, downconversion to baseband, and hard-decision sampling, and therefore requires the remaining functions, including bit/frame

Feature	TRF6900A	CC1100	ADF7020
Supported Modulation Formats	FM, FSK	BFSK, GFSK, MSK, OOK, ASK	BFSK, GFSK, OOK, GOOK, ASK
Date Rate	100 Kbps max	1.2 - 500 Kbps	0.15 - 200 Kbps
Sensitivity @ 1.2 Kbps	-110 dBm	-111 dBm	-119 dBm
Output Power	4.5 dBm	10 dBm	13 dBm
Tx Current (max pwr)	37 mA	31 mA	28 mA
Rx Current (max sens.)	26 mA	15 mA	21 mA
Number of External Parts (typical)	68	23	34

TABLE I
SUMMARY OF MANUFACTURER'S SPECIFICATIONS.

synchronization, demodulation, and bit tracking be performed in software by an MCU (microcontroller). The CC and ADF are more recent designs, and share several characteristics. Both perform frame synchronization, bit tracking, and detection, and provide a "black box" serial interface for loading of transmit payloads and readout of received payloads. Also in contrast to the TRF receiver, which uses two stages of downconversion both of these SIRs use a single I/Q downconversion to a low IF, where all subsequent processing occurs.

III. METHODS

As indicated in Table I, each SIR must be integrated onto a printed-circuit board with passive components, an antenna, an MCU interface, and a regulated power supply. We used two samples of the each manufacturer's evaluation boards for each of the CC and ADF SIRs, and a representative sample of an in-house board for the TRF. This hardware, along with performance-optimized MCU-based real-time software for bit/frame acquisition, bit tracking, and detection, is used in the second-generation WiSARDNet wireless sensor network implementation [8].

The RF FSK test signals were generated using an Agilent E4432B RF signal generator, externally driven in some cases by an MCU to achieve the desired frequency deviation. Power was measured with an Agilent E4402B spectrum analyzer. All cable losses were measured with an Agilent 8714E RF network analyzer and compensated for accordingly. For power consumption, we measured current using an Agilent 3458A programmable DMM.

IV. LINK QUALITY

Each SIR receiver requires two fields—a preamble and and a sync or sync/ID field—at the start of a packet for automatic frequency control (AFC) pull-in, frame sync, and bit tracking initialization. TI recommends a 4-byte preamble followed by a 4-byte sync field (8 bytes for a 500 Kbps data rate) for the CC, while AD recommends for the ADF a 6-byte preamble (for AFC pull-in) followed by a 16-24 bit sync/ID field. In our testing, we used a 48-bit preamble for all three SIR's. We used a 13-bit Barker code sync field for the TRF. The sync field used for the CC was 4 bytes, giving it a theoretical advantage (at the expense of energetic overhead) over the 3-byte sync field used for ADF testing.

The 256-bit data payload consisted of a packet number (ID) and data generated by a linear feedback shift register pseudo-random number generator.

An MCU was used to generate the packets transmitted by the RF signal generator, and another MCU received the packets from the SIR. In each case, an MCU was used to setup and turn on the SIR; as mentioned earlier, algorithms in software were used for the TRF.

A. Packet-Level Sensitivity

One systems-level figure of merit is overall "goodput"; this measure is normally used for networks, but captures link-level performance with one criterion: the *packet error rate* (PER). In our PER testing, we tracked the fate of packets in the receiver, including (i) packets the receiver was unable to synchronize on, and (ii) packets that passed the synchronization threshold but had at least one bit error. The PER is the ratio of the number of packets that failed the synchronization criterion plus those that had at least one bit error to the total number of packets transmitted. In this data and the BER data following, error bars indicate the two-sided standard deviation of the estimate based on tests using 500K bits at each input power level, and results were averaged over the two sample boards for the ADF and CC. The results (Figure 1) for a data rate of 1250 bps show that the TRF/software radio (which was highly optimized) was better in performance that the CC. At the higher 4800 bps data rate, the data confirms the manufacturer's data sheets with a sensitivity advantage of 5 dB for the ADF over the CC.

B. Bit Error Rate Statistics

Once the receiver has achieved synchronization on a packet, the bit error rate (BER), the ratio of the number of bit errors to the total number of bits in packets for which the receiver achieved synchronization, becomes important for the evaluation of automatic-repeat request coding schemes or forward error-control codes. Accordingly, we measured the BER (Figure 2), again using 500K bits transmitted in packets with 256-bit data payloads as for the PER testing. Here we see the ADF has a sensitivity advantage of 6.5 dB over the CC, which corresponds reasonably well to the 8 dB advantage from the manufacturers' data sheets.

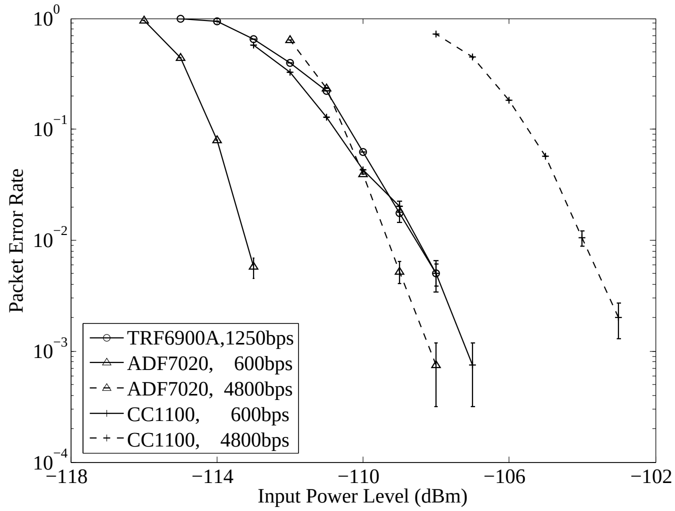

Fig. 1. PER as a function of received signal power for data rates of 600 bps (1250 bps for TRF) and 4800 bps.

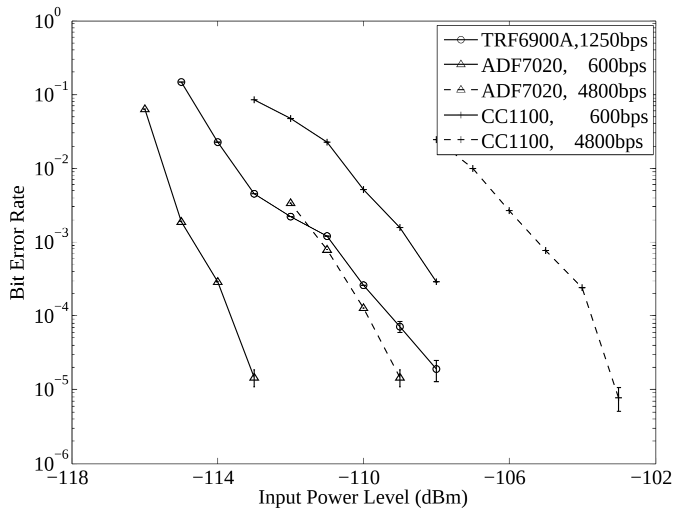

Fig. 2. BER as a function of received signal power for data rates of 600 bps (1250 bps for TRF) and 4800 bps.

C. *Automatic Frequency Control*

If a receiver has poor frequency acquisition performance, then high sensitivity in the ideal case (where there is zero frequency difference between the transmitted carrier and the receiver center frequency) can be nullified in real-world links where the crystal oscillators in the transmitter and receiver have a frequency offset. This offset can be caused by manufacturing variations, temperature differences between transmit and receive radios, and aging. We tested frequency pull-in (acquisition) performance of the CC and ADF receivers by measuring the packet error rate (PER) as a function of a transmit-receive frequency offset programmed into the RF signal generator, and found the received power required for a 1% PER at each offset. The results (Figure 3) show that the sensitivity advantage of the ADF over the CC is *robust* with respect to frequency offset—the advantage over the CC

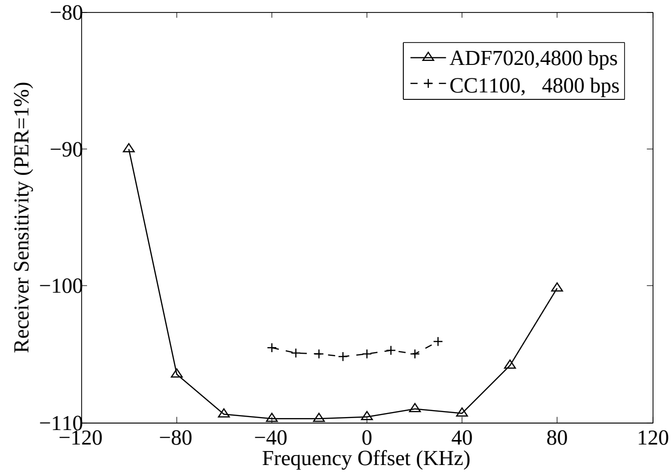

Fig. 3. Receiver sensitivity (at 1% PER) as a function of Tx-Rx frequency offset.

at zero frequency offset (see Figure 3) is maintained over all frequencies within the pull-in range of ± 40 KHz tested for the CC. In addition, the ADF has a significantly wider pull-in range. This advantage can exploited in the design of sensor network nodes: a larger pull-in range allows for a combination of lower-cost crystals and/or longer node hibernation intervals between communication transactions (and greater energy savings).

V. Link Characterization and Optimization

In WSN's where channel capacity is not a limitation and energy efficiency is paramount, the optimum solution may be to use the highest possible channel data rate (and hence higher transmit power); that is, high-rate bursts between sleeping intervals is often preferable to lower rate transmissions for a given real-time average data rate. This maximizes the efficiency of the power amplifier and possibly the system's voltage regulator (especially if the latter is of the linear type). However, other constraints may exist, such as a data rate that was fixed at design time. In this case, minimization of power usage under the constraint of a maximum allowed link BER/PER quality may be desired. It may also be useful to use measured link losses for channel characterization and node localization. Hence both measurement of received signal power and control of transmit power may be of great interest in some applications.

A. *Received Signal Measurement*

We characterized the accuracy and precision of the measured received power provided by the RSSI (received signal strength indicator) function of the ADF and CC. As shown in Figure 4, the ADF's measurement precision decreases monotonically with the true received signal power. This is expected, since thermal noise is constant. The CC's precision is less well-behaved. The accuracy is reported as an offset of the sample mean from the true value. In our tests, the ADF and CC were respectively below and above the true value. We

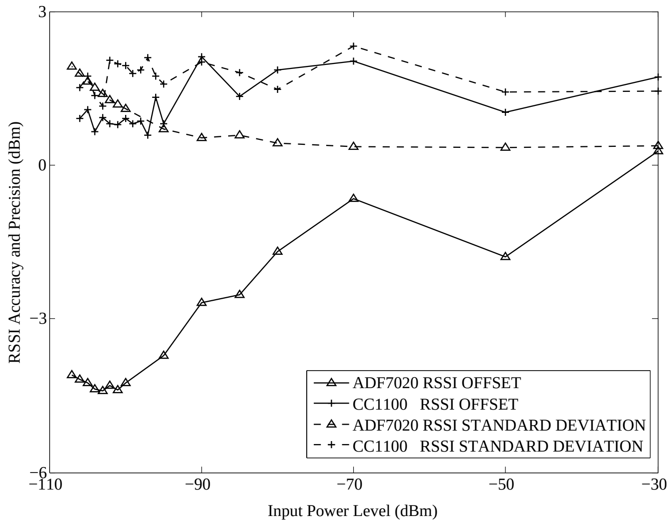

Fig. 4. RSSI accuracy and precision for the ADF and CC SIR's as a function of true received signal power.

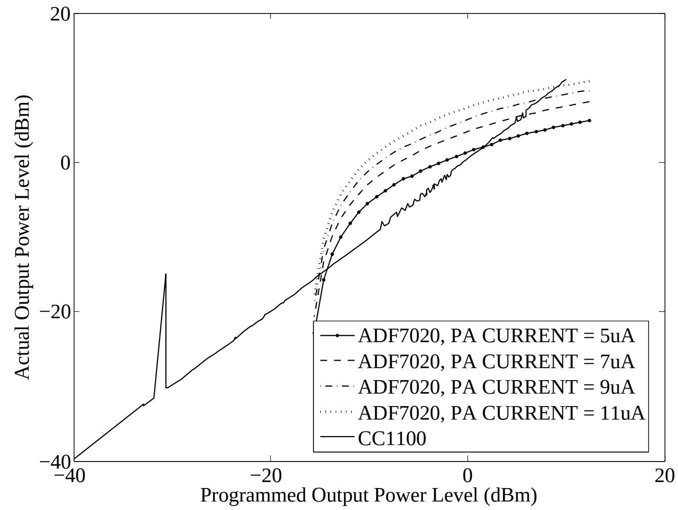

Fig. 5. Transmit power control accuracy for CC1100.

also observed enough variations between samples that both would likely require unit-level calibrations and compensation if accuracy of ≤ 1 dB was required.

B. Transmit Power Control

According to the manufacturer, the ADF provides programmable output power from -16 to +13 dBm in 0.3 dBm steps via a 6-bit register field and PA bias current control via a 2-bit field. Similarly, the output power level of the CC has two levels of programmability via a 3-bit value to select one of eight programmable registers; the combination provides a specified range of -30 to +10 dBm. For these SIRs, we tried to provide some initial answers to the following questions: When the transmit power is set according to the formula published by the manufacturer, what is the actual transmitted power? If not, is it possible to generate a calibration curve?

For brevity, we averaged the results of our two sample boards for each SIR. In our tests using the manufacturers' evaluation boards, both SIRs achieved essentially the same maximum transmit power (Figure 5). The CC is superior in terms of linearity; for accurate power control, the ADF would require a calibration curve. Interestingly, it is possible to set the transmit power on the CC to levels much lower than -30 dBm. However, both samples of our CC boards showed a glitch at a programmed power output of -31 dBm—but this is below the manufacturer's specified minimum of -30 dBm. We note that extremely low transmit powers are not particularly useful since (as can be computed from the data sheets) power efficiency decreases with output power.

VI. Power and Energy Efficiency

A. Current Consumption

We measured and computed the current consumption data for the CC and ADF for a data payload of $m = 256$ bits. Both SIR's were configured for maximum transmit power and sensitivity. Modulation was FSK with a frequency deviation of 50 KHz. We used Manchester coding for an information bit rate of 4800 bps and a baud rate of 9600.

We characterized the current consumption for three phases of SIR operation: transmission/reception of the data payload, packet overhead (including preamble and sync fields), and warmup/shutdown. To determine each of these components, we approximated the supplied (continuous) current $i(t)$ from measurements at a 2 Ks/s sampling rate.

Current was measured for the ADF (Table II) and the CC (Table III) on one sample of each. Because our measured data differed somewhat from the manufacturers' data, we averaged the two to get reported results. Based on our experience in programming the two SIR's we assumed a programming time of 0.1 ms for writing one control register on the ADF, and 0.6 ms for writing all 36 control registers on the CC. Note that programming the ADF for operation typically requires writing about 6 registers, while programming the CC requires more register writes.

The CC uses significantly less current in receive mode, perhaps due to less complexity in the frequency tracking subsystem. Warmup time for both SIRs includes time for programming them via serial communication, and we also assumed a 1 ms time for radio shutdown. These will vary depending on the serial clock rate used, but does not appreciably affect energy consumption (next section) due to the relatively short time in this mode compared with the preamble/sync and data portions of the packet.

B. Overall Energy Efficiency

Building on the current measurement tests,we computed the energy using $v\int_{\tau} i(t)dt$ where v was the supply voltage and τ was the relevant measurement interval. The total communication energy consumption per bit is

$$\phi = \frac{\Phi(m)}{m} = \frac{\Phi^{\mathrm{Tx}}(m) + \Phi^{\mathrm{Rx}}(m)}{m},$$

Function	ADF7020			
	Tx		Rx	
	Time (ms)	Current (mA)	Time (ms)	Current (mA)
Warm up	1.3	11	3	15
Packet Overhead	7.5	31.5	7.5	22.5
Data	53.3	31.5	53.3	22.5
Shut down	1	17.5	1	15

TABLE II
ADF7020 CURRENT CONSUMPTION.

Function	CC1100			
	Tx		Rx	
	Time (ms)	Current (mA)	Time (ms)	Current (mA)
Warm up	1.6	5.7	1.6	5.6
Packet Overhead	8.3	32.7	8.3	16.2
Data	53.3	32.7	53.3	16.2
Shut down	1	16.5	1	9

TABLE III
CC1100 CURRENT CONSUMPTION.

Energy, μJ		CC1100		ADF7020	
		Tx	Rx	Tx	Rx
Warmup/Shutdown	$\Phi^{?x}_{WS}$	85	60	105	193
Packet Overhead	$\Phi^{?x}_{PO}$	896	444	780	557
Data	$\Phi^{?x}_{D}(256)$	5752	2849	5541	3958
Total Packet Energy	$\Phi(256)$	10086		11134	
Total Bit Energy	ϕ	39		44	

TABLE IV
TRANSMIT, RECEIVE AND TOTAL ENERGY CONSUMPTION.

where $\Phi^{\mathrm{Tx}}(m)$ ($\Phi^{\mathrm{Rx}}(m)$) is the total transmit (receive) energy used for the packet with a length-m data payload. Each of these packet-level energies is the sum

$$\Phi^{?\mathrm{x}}(m) = \Phi^{?\mathrm{x}}_{\mathrm{WS}} + \Phi^{?\mathrm{x}}_{\mathrm{PO}} + \Phi^{?\mathrm{x}}_{\mathrm{D}}(m),\ ? \in \{\mathrm{T}, \mathrm{R}\}$$

where WS and PO respectively denote the fixed costs of warmup/shutdown and packet overhead (preamble), and D denotes the energy used for the data payload.

As described in the previous section, the energy cost of programming the SIRs is application-dependent. However, this cost is small relative to the overhead costs of transmitting and receiving the preamble and sync fields. More importantly, since these costs are amortized over the data bits, longer packets increase energy efficiency, possibly at the expense of delay depending on the true bandwidth of the sensed data.

The per-bit energy use ϕ (Table IV) indicates an approximately 10% advantage for the CC under these test conditions. However, the joint effects of the channel data rate, transmit power, and media access control algorithm used are extremely important. For example, the actual per-bit energy cost in applications must also account for the time the receiver is on prior to the arrival of the packet. These considerations have motivated algorithms to minimize listening time [9], and transmit strategies that maximize channel data rates in short bursts, e.g., the ShockBurst technology used by Nordic Semiconductor. However, there is yet another trade-off: lower data rates imply higher E_b/N_0 values and hence greater range, which can be important in environmental sensing applications.

VII. CONCLUSIONS

The goal of this work was to better understand the performance of SIR's in the context of approaches to energy-use optimization in wireless sensor networks that integrate variable data rates, link power control, forward and ARQ channel coding, and overall node hardware cost. Because the two current-generation SIR's have significantly higher levels of digital-circuit integration, they require far less software-radio technology performance than the last-generation chip. We also found that selection of an SIR is application-dependent. The ADF7020's lower noise floor gave it a higher sensitivity than the CC1100, an advantage in applications where channel losses are high. On the other hand, the CC1100 used less current while transmitting, a potential advantage in applications where energy efficiency is paramount. These two SIR's differ significantly in other areas: for example, the CC1100 has a built-in 1/2-rate convolutional encoder/decoder, while the ADF7020 can output soft bit information for implementation of other codes and decoders. Radio selection for sensor node design is a challenging task that requires detailed knowledge of the SIR capabilities and the application context.

ACKNOWLEDGMENTS

The authors thank S. Wang for work on processing, interpreting and plotting the test data. This work was supported in part by NSF grant CNS-0540414.

REFERENCES

[1] J. Zhao and R. Govindan, "Understanding Packet Delivery Performance in Dense Wireless Sensor Networks," ACM Sensys'03, Los Angeles, CA, Nov. 2003.

[2] D. Lal, A. Manjeshwar, F. Herrmann, E. Uysal-Biyikoglu, and A. Keshavarzian, "Measurement and Characterization of Link Quality Metrics in Energy Constrained Wireless Sensor Networks," Proc. IEEE 2003 Global Communications Conference (Globecom), San Francisco, California, Dec. 2003.

[3] A. Woo, T. Tong and D. Culler, "Taming the Underlying Challenges of Reliable Multihop Routing in Sensor Networks," ACM SenSys 2003, Los Angeles, CA, Nov. 2003.

[4] M. Rondinone et al., "Designing a Reliable and Stable Link Quality Metric for Wireless Sensor Networks," REALWSN, Glasgow, UK, April 2008.

[5] Texas Instruments, TRF6900A Single-Chip RF Transceiver Data Sheet, SLAS258E, Revised Aug 2005.

[6] Texas Instruments, CC1100 Low-Cost Low-Power Sub-1 GHz RF Transceiver Data Sheet Rev. C, SWRS038C, 27 May 2008.

[7] Analog Devices, ADF7020 ISM Band Transceiver IC Data Sheet Rev. B, August 2007.

[8] Z. Yang et al., "WiSARDNet: A System Solution for High Performance In Situ Environmental Monitoring," 2nd International Workshop on Networked Sensor Systems (INSS 2005), San Diego, 2005.

[9] R. Jurdak, P. Baldi, and C. Lopes, "Adaptive Low Power Listening for Wireless Sensor Networks," *IEEE Transactions on Mobile Computing,* Vol. 6, No. 8, 2007.