Low-Power Direct-Sequence Spread-Spectrum Modem Architecture For Distributed Wireless Sensor Networks

Charles Chien
Rockwell Scientific Company
1049 Camino Dos Rios
Thousand Oaks, CA 91360
1-805-373-4106

cchien@rswc.com

Igor Elgorriaga
Rockwell Scientific Company
1049 Camino Dos Rios
Thousand Oaks, CA 91360
1-805-373-4332

ielgorriaga@rswc.com

Charles McConaghy
Lawrence Livermore National Lab
7000 East Avenue
Livermore, CA 94550-9234
1-925-422-8780

mcconaghy1@popeye.llnl.gov

ABSTRACT

Emerging CMOS and MEMS technologies enable the implementation of a large number of wireless distributed microsensors that can be easily and rapidly deployed to form highly redundant, self-configuring, and ad hoc sensor networks. To facilitate ease of deployment, these sensors should operate on battery for extended periods of time. A particular challenge in maintaining extended battery lifetime lies in achieving communications with low power. This paper presents a direct-sequence spread-spectrum modem architecture that provides robust communications for wireless sensor networks while dissipating very low power. The modem architecture has been verified in an FPGA implementation that dissipates only 33 mW for both transmission and reception. The implementation can be easily mapped to an ASIC technology with an estimated power performance of less than 1 mW.

Keywords

Low power, modem, spread spectrum, direct sequence, sensor.

1. INTRODUCTION

An important class of emerging networked systems for many military and commercial applications is wireless distributed microsensor networks that consist of a collection of communicating nodes, where each node incorporates a) one or more sensors for measuring the environment, b) processing capability in order to process sensor data into "high value" information and to accomplish local control, and c) a radio to communicate information to/from neighboring nodes and eventually to external users [1]. In the not-too-distant future, technology will advance to the point that miniature, ultra-low power CMOS chips integrating radios, digital computing, and MEMS sensors can be produced with low-cost [2]. This will permit large numbers of wireless distributed microsensors to be easily and rapidly deployed (e.g., airdropped into battlefields

or deployed throughout an aircraft or space vehicle) to form highly redundant, self-configuring, ad hoc sensor networks.

The current prototype microsensor node, shown in Figure 1, is based on an open, modular design using commercial-off-the-shelf (COTS) technology. These nodes combine sensing capabilities, such as seismic, acoustic, and magnetic, with a commercial digital cordless telephone radio and an embedded commercial RISC microprocessor in a small package. A detailed power measurement on the sensor node reveals that nearly half of the power is dissipated by the radio circuitry, which by itself consumes approximately 300 mW. The digital modem processing consumes about one third of the total radio power or 100 mW, which can be significant for the sensor network that needs to sustain on battery for over a period of several months. This paper presents a direct-sequence spread-spectrum modem architecture that enables low power communications for wireless sensor networks. Power measurements performed on an FPGA prototype shows a power dissipation of 33 mW when clocked at 32 MHz. When implemented using CMOS ASIC technology, the estimated power is less than 1 mW.
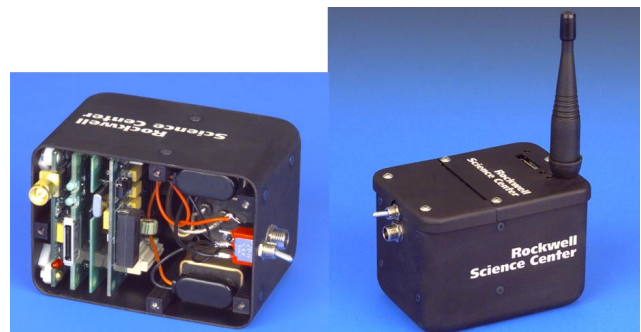
Figure 1 Distributed wireless microsensor nodes.

The paper is organized as follows. Section 2 gives an overview of the system design trade-offs that determined the modem architecture. Section 3 describes a time-shared architecture that provides rapid code acquisition while maintaining low processing power. Section 4 discusses the measured and estimated power performance of the modem when implemented using FPGA and ASIC technology. Section 5 concludes with some discussions on future work.

2. SYSTEM DESIGN

In contrast to most current systems, a sensor network requires low data rates, short range, and low power consumption in order to operate for long periods of time on batteries. These application specific requirements drive the design to reduce power consumption and the modem complexity. To determine trade-offs in complexity against performance, different modem types have been studied which include coherent and non-coherent detection. A coherent demodulator is costly in terms of complexity due to the need for phase and frequency tracking, typically implemented as a Costas loop [3]. However, it does achieve the highest SNR at the receiver for a given transmit power. A non-coherent demodulator, on the other hand, is substantially less complex and lower power but does have a degraded SNR performance. For instance, by applying a differentially coherent demodulator [4], the hardware is reduced by a factor of at least five relative to a coherent design based on the Costas loop. The reduction is due mainly to the elimination of a direct-digital frequency synthesizer and loop filter needed for the phase and frequency tracking function.

The savings in complexity and power reduction, however, result in SNR degradation of about 3-6 dB. Although the SNR reduction is large, such performance degradation can be tolerated in wireless sensor networks with a limited transmission range of 10-30 meters. For instance, assume an antenna gain of 7 dB, E_b/N_o of 10 dB, a carrier frequency of 900 MHz, 1 mW transmit power, bitrate of 7.87 kbps, and a partitioned model [5] that models the path loss exponential as a function of distance d, as shown below:

$$L(d) = \begin{cases} d^2, & 1 < d < 10 \text{ m} \\ \dfrac{d^3}{10}, & 10 < d < 20 \text{ m} \\ \dfrac{d^6}{6.572}, & 20 < d < 40 \text{ m} \\ \dfrac{d^{12}}{10^{10.912}}, & d > 40 \text{ m} \end{cases}$$

Analysis shows that for a BER of 0.001% the link margin at 30 meters is 50 dB which is sufficient to absorb the loss in SNR due to non-coherent demodulation as well as fading effects.

While the SNR loss due to non-coherent demodulation can be tolerated at short transmission ranges, a non-coherent demodulation scheme is vulnerable to frequency offsets. To accommodate large frequency offsets, differential encoding and decoding at the chip level is introduced as shown in Figure 2. In general, the differential decoder uses a complex multiplier that multiplies the received complex signal with the complex conjugate of a delay version. In contrast to traditional encoding and decoding which are applied to the data symbols, the encoder and decoder shown in Figure 2 operates on chips of the direct-sequence spread-spectrum waveform. Since the chip duration T_c is much shorter than the data symbol, the phase change due to a given frequency offset is small enough to achieve a sufficiently low SNR loss at the output of the demodulator.

Take for instance a transmitted waveform with a chipping rate of 1 Mchips/sec and a 127-chip spreading sequence. The data rate is 7.87 kbps. The system is designed for the 900 MHz ISM band. Since low cost is desirable for the deployment of a large number of sensors, 50-ppm crystals are used for frequency references. The resulting frequency offset is 100 ppm of 900 MHz or approximately 90 kHz. It can be shown that the SNR degradation due to frequency offset Δf over a period of T is approximately $20 \log |\pi T \Delta f / \sin(\pi T \Delta f)|$. Given the worst-case bitrate of 7.87 kbps, chip rate of 1 Mcps, and an offset of 90 kHz, the SNR loss is 0.1 dB and 31.3 dB for chip-level and symbol-level decoding respectively. Clearly, chip-level decoding mitigates the SNR loss adequately so that simpler circuits shown in Figure 2 can replace complex coherent demodulators that dissipate higher power.
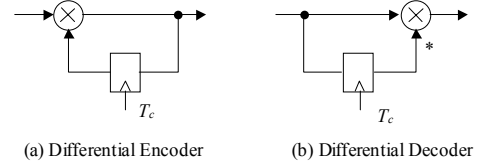

(a) Differential Encoder   (b) Differential Decoder

Figure 2 Chip-level differential encoder/decoder.

3. ARCHITECTURE

The modem consists of three main blocks: an acquisition loop, a timing recovery loop and a demodulator. The acquisition loop block makes a coarse alignment of the local PN-sequence with the transmitted PN-sequence to within half the chip duration. The timing recovery loop is used to reduce the remaining error from the coarse alignment and to track misalignments due to clock drift between the transmitter and the receiver. Once the locally generated sequence and the transmitted sequences are aligned, the received data can be demodulated.

The acquisition loop can be implemented serially or in parallel. The parallel implementation is performed with a matched filter, which is costly in both area and power [6]. To achieve a low power and low complexity design a serial-based implementation is chosen. For the serial acquisition loop, a serial correlator is used as shown in Figure 3. Note that I' and Q' are the differentially decoded signals of the I and Q channels, respectively.

The serial correlator multiplies the data with a locally generated PN-sequence. The integrate-and-dump (I&D) averages the result of the multiplication over a period equal to the PN-sequence length. The I&D output is then compared to a threshold voltage. If the I&D output value exceeds the threshold voltage, acquisition is declared. Otherwise, the acquisition control block skips ½ chip in the local PN generator until the threshold is exceeded. The main drawback of this type of acquisition loop is the long acquisition time, which with high SNR and q samples per chip may take up to $\tau_d (qN-1)T_s$, where N is the PN-sequence length, τ_d is the dwell time [7], and T_s is the sampling period. The acquisition time becomes large for a large N, which is needed for robust transmission. However, for a packet-switched system such as a sensor network where data is being transmitted on burst basis, a long acquisition time can substantially reduce the throughput of the network. To speed up the acquisition, K serial correlators can be placed in parallel to reduce the acquisition time by a factor of K. However, the increase in hardware results in both area as well as power penalty. It will be shown in Section 3.1 that by time-sharing the acquisition and timing recovery, such penalty can be eliminated.
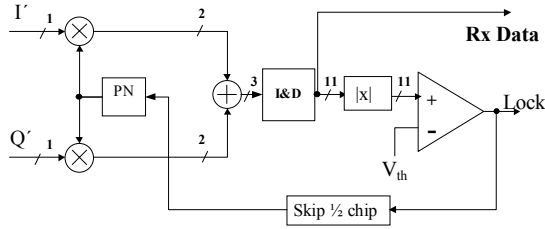
Figure 3 Serial acquisition architecture.

The time recovery loop eliminates any residual timing error after PN acquisition and keeps track of timing drift between the transmitted and received sequences. An early-late gate correlator is used for time tracking as shown in Figure 4. In order to measure the error, two locally generated sequences are compared. These two sequences are ½ chip early and late apart from the transmitted sequence. These sequences are then subtracted from each other to produce the error signal that is averaged with the loop filter. The averaged error controls a numerically controlled oscillator (NCO), which tunes its frequency to drive the timing error to zero.

Demodulation starts when the modem has completely acquired the timing of the incoming sequence, which means that the locally generated sequence is in phase with the transmitted sequence. Once timing has been acquired, the received data stream can be despread. The dispreading is performed serially to reduce power and complexity. Figure 3 shows the demodulator as part of the PN-acquisition loop whereby the sign bit of the integrate-and-dump is used as the decoded data bit.
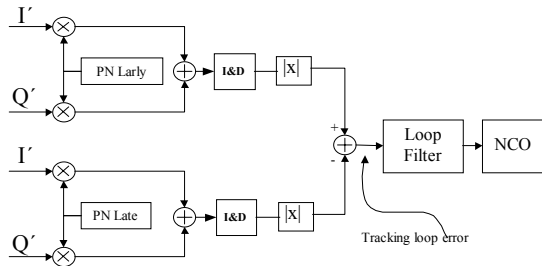

Figure 4 Timing recovery architecture.

Another source of performance degradation is due to quantization in the data paths. When quantized to insufficient number of bits, quantization noise can become appreciably large such that performance becomes unacceptable. The advantage of quantizing to fewest bits is in reduced hardware complexity and power dissipation. Simulations must be performed to determine an appropriate quantization level that provides adequate SNR as well as low hardware complexity and power. With respect to the latter, it would be desirable if one could use one bit quantization. Fixed-point simulation has determined that with one bit quantization, the required E_b/N_0 is 15 dB at 0.001% BER when using BPSK. The SNR degradation is about 5-dB with respect to coherent demodulation and full precision. Given the large link margin, such degradation can be tolerated. One bit quantization eliminates the need for highly complex multipliers and instead allows the use of a simple XOR gate.

The widths of the datapath used in the serial correlator are shown in Figure 3. With the one bit quantization at the inputs, the portion of the datapath responsible for chip rate processing has fewer bit slices so that extensive power saving is achieved. Although larger width

datapath is needed after the I&D, the resulting increase in power dissipation is miniscule since the dumped rate is much lower than the chip rate.

3.1 Time shared Architecture

Implementing a time-shared architecture of the building blocks shown in Figure 5, we can eliminate the drawback of using serial correlators for acquisition due to its long acquisition time without increasing complexity and power consumption. The time-shared architecture shown in Figure 5 takes advantage of the different states of the modem and the hardware that could be time shared in these states. Acquisition, time recovery and demodulation all use serial correlators but only one or two of these blocks are used concurrently. Using the idle logic of the time recovery loop during acquisition, the time to acquire can be reduced by a factor of three. The acquisition time reduction is possible by feeding three PN-sequences, off-set in time, to each of the serial correlator loops. The time offset is equally split among the sequences. This circuitry does not use extra power since although three serial correlators are used, the time to acquire is reduced by the same amount as well. During the time recovery state, two of the serial correlators form the early and late correlators while the remaining correlator is used as the despreader and demodulator. Thus, all three serial correlators are maximally used during the operation of the modem. With respect to a non time-shared architecture, our architecture uses 50% less area and power consumption.

4. MODEM PERFORMANCE

The modem is implemented in FPGA technology using one Xilinx Spartan series FPGA (XCS20XL-3-VQ100) with about 95% utilization of the 400 CLB's. The modem implements variable spreading codes of length 15, 31, 61 and 127 as well as a Barker code of length 11. The variable code lengths provide gains from 21dB down to 10dB. A chipping rate of 1Mhz is used, which results in variable data rates from 7.87 kbps for a spreading gain of 21dB to 90 kbps for the Barker-11 code. Both the PN acquisition and time recovery loops accept input samples at a rate of four samples per chip. The input sampling clock of the modem runs at 32Mhz and is generated by the NCO. This frequency is divided internally to provide a 1-MHz clock to the PN generator and a 4-Mhz clock to the serial correlators in the time recovery and PN-acquisition loops. Power supply to the modem is 3.3V.

4.1 Test Results

To test the PN acquisition, the starting time for a transmission is swept and the acquisition times are measured. With a dwell time of two bits, the results in Table 1 closely match the expected acquisition times.

Table 1 Measured acquisition times.

Length	Max. Acquisition (bits/µs)
11	8/88
15	10/150
31	21/651
63	42/2646
127	85/10795

During modem operation, the measured power consumption is 33 mW. To test the tracking performance of the modem, the transmitter frequency is fixed while the receiver frequency is varied. The time

recovery loop tracks frequency variations over 3.2 kHz or 100 ppm of the 32 MHz clock. For the demodulator tests, a known pseudo random pattern is transmitted and the received data is compared with this known pattern. A throughput of greater than 95% has been achieved over short-ranges (30 m) when tested with a 900 MHz RF front-end.

5. CONCLUSIONS AND FUTURE WORK

This paper describes a low-power direct-sequence spread-spectrum modem architecture for distributed wireless sensor networks. Using 1-bit chip-level differential decoding, a low complexity and low power demodulator is implemented that does not require highly complex phase and frequency tracking loops. Furthermore, through the time-sharing of three serial correlator, a modem has been implemented that achieves a 3X reduction in acquisition time but with no power or area penalty.

The modem has been implemented in a single Xilinx FPGA with less than 400 CLB's and a power dissipation of 33 mW at 3.3 V. The equivalent gate of the modem is approximately 8000. A wireless sensor prototype that uses the modem implemented in an FPGA is being developed at Livermore. The sensor node is a compact stack of modules consisting of a low frequency MEMS accelerometer, a digital signal processor, the FPGA-based spread-spectrum modem, a delay-line based RF front-end, and a Lithium battery. The entire sensor node fits in a 2 x 2 x 1 inch form factor.

Based on CMOS standard-cell library parameters, the power dissipation for the entire modem in 0.35 μm CMOS is estimated to be 600 μW. Such a low power wireless modem is suitable for future generation wireless sensors that may have sensors, microprocessor, modem, RF front-end, and a coin-cell battery all integrated in less than 0.5 cubic inches. With power management built into each node, it is expected that a sensor node can live on a coin-cell battery for as long as a year assuming a total power of 5 mW at 3V and a duty cycle of 0.5 %.

7. REFERENCES

[1] H. Marcy, J. Agre, C. Chien, and L. Clare, "Robust Microsensor Networking for Aerial Denial," *IEEE Milcom 99*, November 1999.

[2] G. Asada, M. Dong, T. S. Lin, F. Newberg, G. Pottie, W. J. Kaiser, H. O. Marcy, "Wireless integrated network sensors: low power systems on a chip," *IEEE ESSCIRC'98*, Hague, pp. 9-16, September 1998.

[3] C. Chien, R. Jain, E. Cohen, H. Samueli, "A12.7 Mchips/sec All-Digital BPSK Direct-Sequence Spread Spectrum IF Transceiver in 1.2- m CMOS", *IEEE Journal of Solid-State Circuits*, Vol. 29, No. 12, December 1994.

[4] J. G. Proakis, Digital Communications, McGraw Hill, New York, 1995.

[5] D. Akerberg, "Properties of a TDMA pico-cellular office communication system," *Proc. IEEE GLOBECOM*, Florida, pp. 1343-1349, December 1988.

[6] C. Deng and C. Chien, "A PN-acquisition ASIC for wireless CDMA systems," *Proceedings of the IEEE Custom Integrated Circuits Conference*, Orlando, Florida, May 21-24, 2000.

[7] D. M. DiCarlo and C. L. Weber, "Multiple dwell serial search: performance and application to direct-sequence code acquisition," *IEEE Transactions on Communications*, Vol. 31, No. 5, pp. 650-659, May 1983.
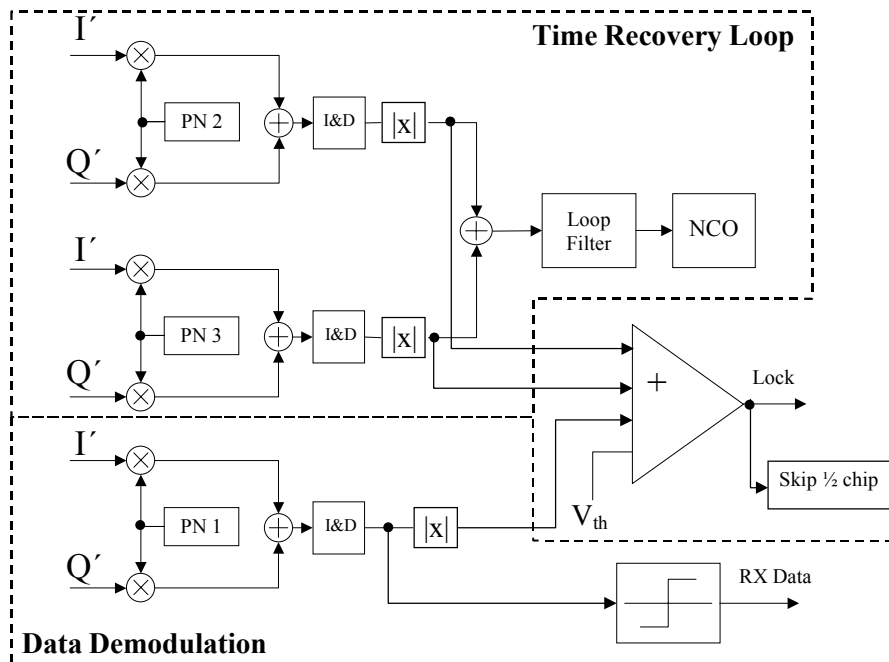
Figure 5 Time-shared modem architecture.