Energy Efficient Pulsed-UWB CMOS Circuits and Systems

David D. Wentzloff[1], Fred S. Lee[2], Denis C. Daly[3], Manish Bhardwaj[3],
Patrick P. Mercier[3], and Anantha P. Chandrakasan[3]

[1]University of Michigan, Ann Arbor, MI, 48109, USA
[2]Rambus Inc., Los Altos, CA, 94022, USA
[3]Massachusetts Institute of Technology, Cambridge, MA, 02139, USA

(*Invited Paper*)

Abstract — A custom UWB transceiver chipset is presented that communicates in three 550MHz-wide channels in the 3.1 to 5GHz band by using pulse position modulation (PPM). The transmitter uses an all-digital architecture and calibration technique to synthesize pulses with programmable width and center frequency. No analog bias currents or RF oscillators are required in the transmitter. The receiver performs channel-selection filtering, energy detection, and bit-slicing. The receiver circuits operate at 0.65V and 0.5V, and can turn on in 2ns for duty-cycled operation. The two chips are fabricated in a 90nm CMOS process, and achieve a combined 2.5nJ/bit at a data rate of 16.7Mb/s.

Index Terms — CMOS integrated circuits, ultra-wideband radio, UWB, pulse position modulation.

I. Introduction

Pulsed ultra-wideband (UWB) communication is an active field of research that has roots that can be traced back to the original Marconi spark gap radio. This field has gained momentum since a change in FCC regulations in 2002 that allows unlicensed communication using UWB. UWB signaling has many attributes that make it attractive for a wide range of applications; from ultra-low-power RFID tags and wireless sensors to streaming wireless multimedia and wireless USB at greater than 1Gb/s [1].

The amount of available bandwidth in the UWB band far exceeds the bandwidth required for low data rate transceivers. UWB radios are uniquely positioned to exploit the available bandwidth by trading off spectral precision and efficiency for other system specifications such as energy/bit and power consumption. This paper presents a custom pulsed-UWB chipset fabricated in a 90nm CMOS process. The chipset includes features to reduce power consumption such as an all-digital transmitter architecture, low-voltage RF and analog circuits in the receiver, and no RF local oscillators allowing the chipset to power on in 2ns for duty-cycled operation. The architecture and design of the custom chipset is described in Section II and III. The synchronization algorithm is described in Section IV. Measured results from the individual chips and the wireless system are presented in Section V.

II. Ultra-Low-Power UWB Architectures

Recent advances in semiconductor process scaling have enabled ultra-low power system-on-chip UWB transceivers requiring minimal off-chip components. This high level of integration has reduced UWB transceivers' performance sensitivity to parasitic capacitances and inductances of bondwires and pads [2][3]. The UWB chipset presented in this paper forms a robust UWB system with nearly all RF circuit blocks integrated on-chip other than a crystal oscillator, antenna, TX/RX switch, and transmit band-select filter.

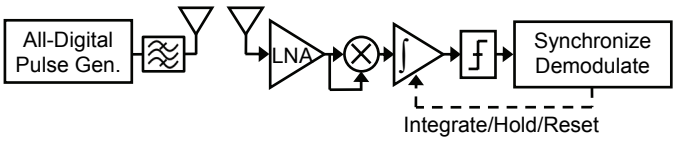

Fig. 1. Custom chipset transceiver architecture.

The architecture of the transceiver chipset presented in this paper is shown in Fig. 1 [4][5]. Binary pulse-position modulation (PPM) is used to encode the transmitted data. The PPM signal is transmitted in one of three channels in the 3.1 to 5GHz band, as shown in Fig. 2. Three channels are used to avoid potential in-band interferers and to add frequency diversity for multiple users. The receiver is a non-coherent energy detector that compares the received energy of two adjacent timeslots. The transmitter generates short pulses by combining edges from a digital delay line and amplifies the pulse with an all-digital power amplifier. The following sub-sections outline two high level ideas that were used to achieve

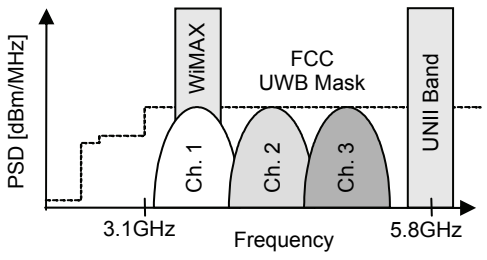

Fig. 2. Three-channel frequency plan and narrowband interferers.

low-power operation: non-coherent communication and low-Q highly integrated circuits.

A. Non-coherent Communication

For short ranges, non-coherent communication allows for energy and complexity savings over coherent communication. Non-coherent solutions suffer from loss in SNR, but for short-range links where transmitted output power does not dominate the power budget, the energy savings associated with non-coherent communication can overcome this loss. Non-coherent communication does not require precise phase control, which allows both the transmitter and receiver architecture to be simplified, particularly the high frequency circuits that can consume the majority of power in a wireless transceiver. A non-coherent energy-detection scheme can further reduce hardware complexity while providing resilience to multi-path fading without the cost of high frequency Rake-based techniques [6].

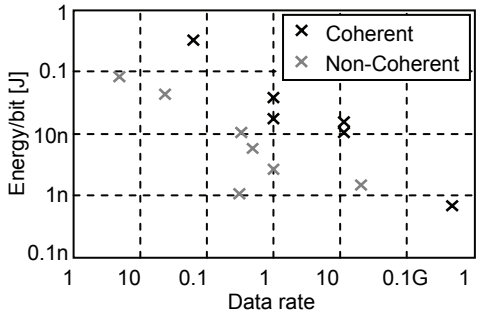

Fig. 3. Non-coherent vs. coherent energy-per-bit values for UWB and narrowband receivers recently published at ISSCC.

Fig. 3 presents energy-per-bit values for recently published coherent and non-coherent receivers. The figure shows a 10x energy/bit advantage of non-coherent over coherent receivers. The reduction in energy typically comes at the expense of system performance. The IEEE 802.15.4a standard [7] has recognized the advantages offered by non-coherent communication and includes support for it.

B. Low-Q Circuits

The RF circuit blocks in the receiver often require a significant amount of gain. For low-power realization of such gain, tuned circuits can be designed such that their bandwidth and center frequency matches the wireless signal. Single-stage amplifiers at a fixed current have a constant gain-bandwidth product, and by minimizing the bandwidth, gain can be maximized. For narrowband signals, this requires a very high quality factor resonant tank, which cannot be easily realized without costly off-chip resonators. Typical on-chip quality factors in the GHz frequency range are limited below 30, which results in a minimum circuit bandwidth of over a hundred MHz. Since UWB signals occupy many hundreds of MHz of bandwidth, low-Q circuits can be designed to match this bandwidth and center frequency. The receiver presented in this paper uses this bandwidth matching approach to minimize power while achieving 40dB of gain.

Even though tuned circuits allow for optimal energy efficiency, untuned circuits can be used without a significant increase in power consumption while allowing for significantly reduced area. While tuned circuits are typically limited to an octave of tuning range, untuned circuits can achieve much wider ranges. In recent years, many digital UWB transmitters have been presented that utilize untuned ring oscillators and power amplifiers to generate pulses [8]. In fact, when noise is not the limiting constraint, untuned circuits may offer superior energy efficiency to integrated tuned circuits [9].

III. Transceiver Circuit Design

A. Transmitter

Mostly-digital pulsed-UWB transmitters are driving down energy/bit using deep-submicron CMOS processes [1][10][11]. A block diagram of this transmitter is shown in Fig. 4 [4]. The transmitter uses an all-digital architecture, and no analog bias current or RF local oscillators are required. Power is therefore only consumed in subthreshold leakage currents and switching events. This transmitter architecture is made practical by the use of a non-coherent receiver, and by the wide available bandwidth. These attributes allow the frequency tolerances of the transmitter to be relaxed, enabling the digital architecture.

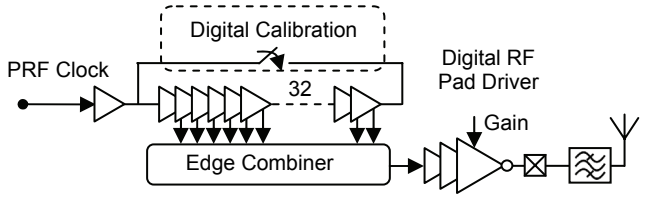

Fig. 4. Transmitter block diagram.

The transmitter is clocked by an external FPGA at the PRF input. Each edge of the PRF input propagates through a 32-stage delay line with an 8-bit digitally controlled delay. UWB pulses are generated by combining the equally-delayed edges to form a single RF pulse. The edge combination is similar in operation to a frequency multiplier, where the output of the combiner is toggled when an edge is received on any of its inputs. The center frequency of the pulse spectrum is varied by controlling the delay between edges. The width of the pulse is varied by making the number of edges combined programmable. By making both the delay and number of edges programmable, the pulse spectrum may be controlled with 6000ppm accuracy without requiring an RF local oscillator.

Twenty-five edges are required to synthesize UWB pulses in the highest channel, therefore a 32-stage delay line is used, from which 30 edges are available for combination.

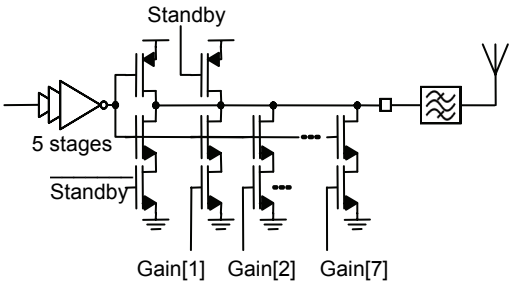

Fig. 5. Schematic of the transmitter output driver.

Individual edges are selected by ANDing them with a mask register. The 30 edges are combined using two time-interleaved 15-edge combiners, the outputs of which are XORed to complete the pulse synthesis. The pulse is buffered by a digital pad driver, and the output is filtered by an off-chip UWB band-select filter that directly drives a 50Ω antenna. A schematic of the output buffer is shown in Fig. 5. Stacked NMOS devices are used in the buffer to reduce subthreshold leakage current. The pull-down network is divided to implement a linear-in-dB adjustable transmit power. A standby mode is implemented in the buffer to reduce leakage during the interval between packets. The transmitter was fabricated in a 90nm CMOS process, and a die photo is shown in Fig. 6 [4].

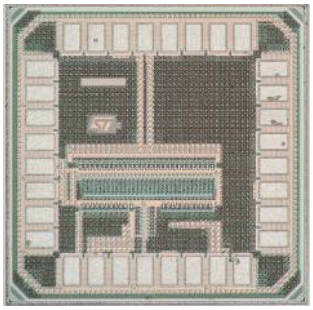

Fig. 6. Transmitter die photo.

B. Receiver

The chipset communicates using a binary PPM signaling scheme (Fig. 7), where a 2ns wide UWB pulse can arrive in one of two time slots during synchronized reception: T_{int1}, or T_{int2}. If the pulse arrives in T_{int1} or T_{int2} then a one or a zero is declared by the receiver, respectively. Though the UWB pulse is only 2ns wide, T_{int} is set to 30ns such that the worst-case multipath channel and crystal frequency offsets can be tolerated. The data rate is scaled by changing the time period between PPM symbols, T_{frame}. To maximize energy efficiency, the receiver is designed so that it can disable all bias currents during the time between symbols, thereby exploiting the duty-cycling inherent in pulsed-UWB signaling. Therefore, the amount of *energy* required to demodulate one bit does not change over a wide range of data rates. The maximum data rate of this architecture is 16.7Mbps when T_{frame}=60ns.

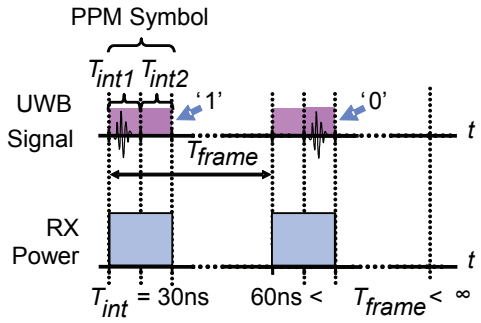

Fig. 7. PPM signaling scheme.

Fig. 8 shows a simplified block diagram of the non-coherent receiver. It is comprised of a 3-5GHz subbanded RF front-end, a passive self-mixer, and a low power mixed-signal baseband [5]. No RF PLL or oscillator is required. Only a 33MHz crystal is needed to operate the mixed-signal baseband. Channel selection is performed in the RF front-end for out-of-band noise/interference robustness. The RF and mixed-signal baseband operate at 0.65V and 0.5V respectively, for low power operation. Most of the receiver power is consumed in the RF gain stages. Energy is conserved at lower data rates by switching the receiver on only during the time that pulses are present.

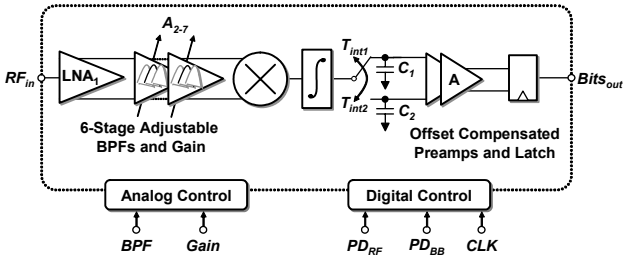

Fig. 8. Block diagram of the receiver.

The non-coherent receiver uses a mixed-signal relative-compare baseband to determine the bit. For bit-slicing, a low-power sample-and-hold capacitor network stores analog integration results during T_{int1} and T_{int2} onto separate capacitors C_1 and C_2, respectively. Thereafter, two cascaded offset-compensated preamplifiers and a latch perform a relative-compare on the two capacitor voltages to evaluate the received bit. This scheme inherently performs dc-offset compensation and pre-integrator signal-path normalization in

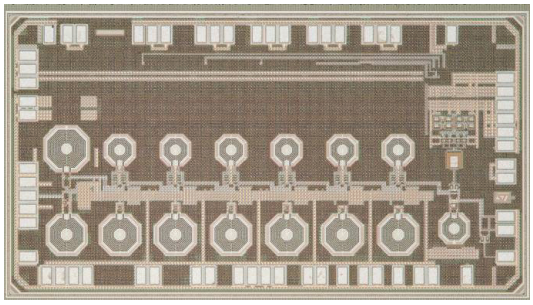

Fig. 9. Receiver die photo.

each bit decision. The receiver is implemented in a 90nm CMOS process, and a die photo is shown in Fig. 9.

C. Wireless Nodes

The transmitter and receiver UWB chips were tested using custom PCBs that mount on commercially available FPGA boards. Photographs of the transmitter and receiver nodes are shown in Fig. 10. Commercial omnidirectional UWB antennas were used for all wireless experiments.

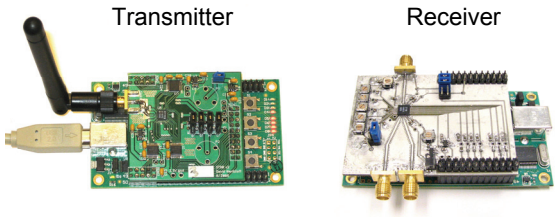

Fig. 10. Pictures of the transmitter and receiver nodes.

The FPGAs are used to implement the lower frequency digital functions while also providing USB communication with a PC. For the transmitter, the FPGA is used for digital configuration, calibration, and to implement PPM modulation. The transmitter FPGA uses instructions issued over the USB bus to configure internal registers in the transmitter IC. These registers control the pulse center frequency, number of edges combined, transmit power level, and power-down states. The receiver FPGA implements the synchronization and demodulation algorithms and is capable of accepting configuration instructions over the USB bus, such as correlator threshold levels.

IV. Synchronization

The receiver uses a preamble in each packet for acquisition and to synchronize the receiver and transmitter clocks before demodulation. When the receiver is synchronized and the PPM time reference known, the receiver compares the RF energy measured in intervals T_{int1} and T_{int2} (Fig. 7) to make a single bit decision. However, prior to synchronization the PPM time reference is unknown; therefore, the receiver compares the RF energy in consecutive 30ns time intervals until synchronization is declared. This results in a bit decision being made after every single integration time, or at twice the data rate. After synchronization is declared only one comparison is made in each PPM frame.

A block diagram of the synchronizer and demodulator implemented in the receiver FPGA is shown in Fig. 11. The data from the receiver arrives at twice the data rate from the consecutive comparisons. The data is parallelized into two 16.7Mb/s data streams, and sent to two identical correlator banks. A 16 bit code is used to identify each transmitter for acquisition. The beginning of the transmitted code is unknown by the receiver, therefore 16 correlators are used that correlate the received data with shifted versions of the 16 bit code. The 16 correlators and shifted codes account for all possible time differences between the transmitter and receiver. The incoming data is correlated with each bit of the code, and the results are accumulated over 16 cycles. After 16 cycles, each accumulator is reset and the correlations repeat. If any of the accumulator values exceed a programmable threshold, packet detection is declared and the receiver is synchronized.

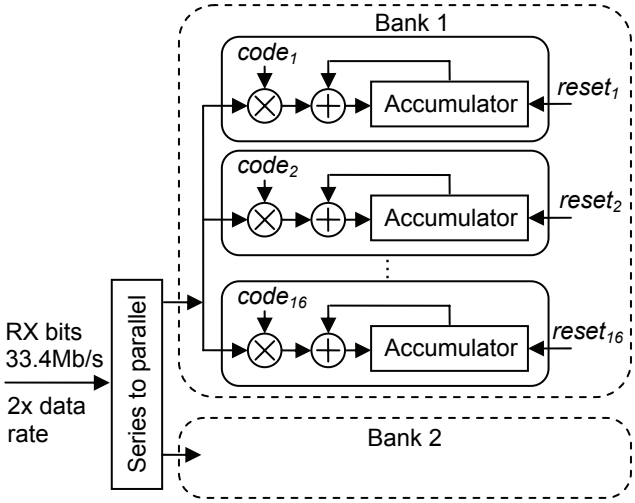

Fig. 11. Block diagram of synchronization correlator banks.

The structure of each packet is shown in Fig. 12. The synchronization code is repeated 4 times to increase the probability of detecting a packet. After synchronization is achieved, the receiver searches for an inverted 16 bit code that indicates the start of the payload data. After the inverted code is received, the receiver will begin filling the receiver FIFO with the header and payload data. If no inverted code is received after a short period of time, a false acquisition is declared and the receiver returns to synchronization mode.

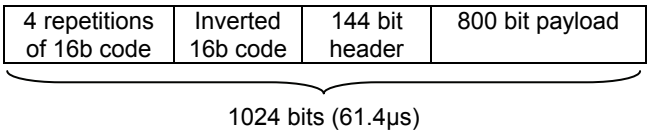

Fig. 12. Packet structure.

V. Experimental Results

A. Transmitter

The transmitter uses DB-BPSK [12] to scramble the transmitted spectrum. DB-BPSK is implemented in hardware by selectively delaying the pulse by half of an RF cycle. The two phases of the measured DB-BPSK pulses are shown in Fig. 13 when the pulse spectrum is centered in channel 2. The two pulses are superimposed for comparison. These pulses appear to be inversions of each other, however by focusing on the beginning of the pulses, it is apparent that they are actually delayed by half of the RF cycle period.

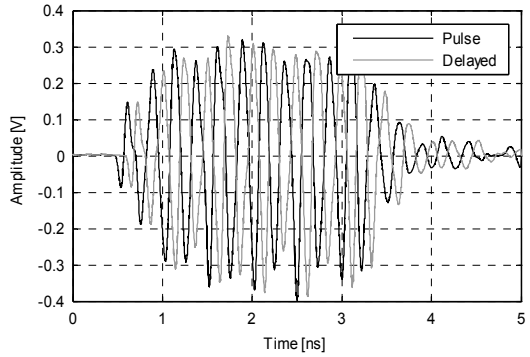

Fig. 13. Measured time domain pulses. The delayed pulse is delayed by half of an RF cycle period.

The measured RF spectra for the three channels are superimposed in Fig. 14 along with the FCC mask. The most difficult specification of the FCC mask to meet was the deep notch between 960MHz and 1.6GHz due to a sub-harmonic of the pulse center frequency introduced by the interleaved edge combiners. The sub-harmonic of the lowest channel lies at 1.7GHz, and spectral content is produced there with a -10dB bandwidth of 550MHz. The off-chip band-select filter sufficiently attenuates the sub-harmonic to meet the mask.

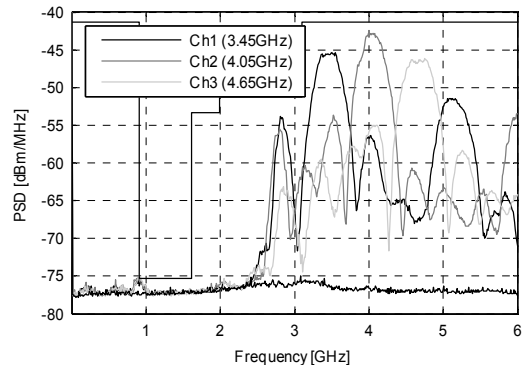

Fig. 14. Measured transmitted spectrum for three channels.

The transmitter uses an all-digital architecture, therefore no analog bias currents are required and power is consumed only in subthreshold leakage current and switching events. The standby power consumption due to subthreshold leakage current is 96µW. The power consumption while generating pulses at a data rate of 16.7Mb/s is 718µW, corresponding to an energy/bit of 43pJ/bit.

B. Receiver

The receiver front-end provides 40dB of gain and high-order roll-off for channel selectivity in each of the 3 bands, as shown in Fig. 15. A noise figure of 8.5-9.5dB is achieved, and better than -10dB S_{11} is achieved for all three bands. The worst-case measured input P_{1dB} is -45dBm.

At 100kb/s, the receiver achieves -99dBm sensitivity at a BER of 10^{-3}. Due to the fast turn-on and turn-off time, the receiver achieves 2.5nJ/bit across three orders of magnitude in data rate.

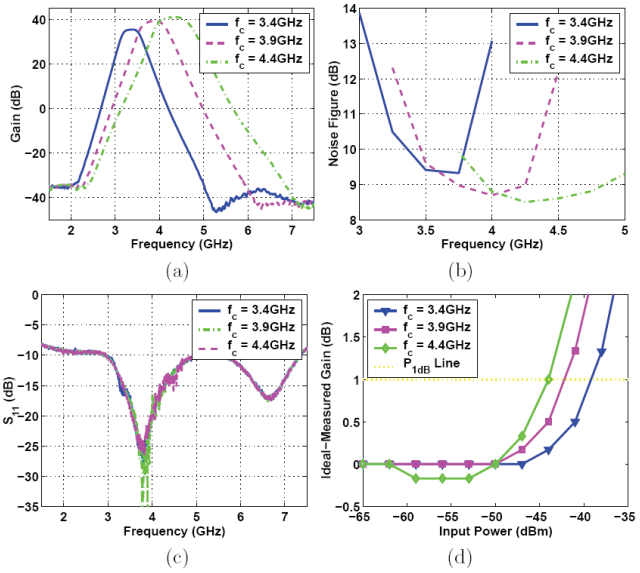

Fig. 15. Receiver front-end measurements.

C. Wireless Demonstration

A streaming video application is realized to demonstrate a unidirectional wireless link using the custom chipset. Streaming video is emulated by sending a sequence of images from one PC to another over the UWB link and displaying them in real time as they are received. The images are 120x160 pixel bitmaps with 12bit RGB color depth (28.8kByte images). The data for each image is divided into 288 packets of 800 bit payloads each. A detailed packet structure is shown in Fig. 12. The header contains an index indicating the appropriate location of the packet data in the overall image. This allows the receiver PC to reassemble each image from the series of packets even if received out of order, and allows the system to gracefully recover from a dropped packet. The index value in the header is repeated three times for redundancy and error correction purposes.

Before transmission, the packets are downloaded to the transmitter node over the USB bus. When enough packets are queued in the transmitter FPGA FIFO, the transmitter begins sending the packets to the receiver over the UWB link. A picture of the wireless demonstration setup is shown in Fig. 16 with the transmitter node and PC in the foreground, and receiver node and PC in the background. The link distance

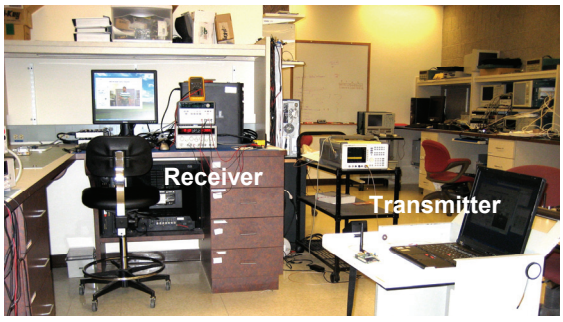

Fig. 16. Wireless test setup.

for this demonstration was 3m. The transmitter node is powered completely from the USB bus, therefore a laptop is used for maximum portability of the transmitter.

The PC on the receiver side continuously polls the receiver FPGA FIFO for received packets. Once a packet is received, it is sent over the USB bus to a PC. On the PC, the header is stripped and the index of the packet data into the image is read, correcting errors that may have occurred in the header only. The payload data is then inserted into the image buffer at the specified index location and the image is displayed on the screen. No error correction or coding is used on the image data so that errors are visible in the image, providing a visual indication of the quality of the wireless link.

VI. Future Directions

There are several opportunities to extend upon the UWB chipset and system presented in this paper. The IEEE 802.15.4a UWB low data rate standard [7] supports a non-coherent physical layer, and the transceiver architecture presented could be made standards compliant with some modifications. One such modification is the transmitter must support pulse bursting, where multiple UWB pulses are transmitted immediately after one another.

Synchronizing reliably at low SNRs using short preambles continues to be a challenge in non-coherent systems, and limits the overall energy efficiency of such systems. Increasing the resolution of the receiver ADC allows for a reduction in synchronization time. We are developing new synchronization codes that reduce preamble lengths by 2-3 times over traditional *m*-sequence codes. We are also investigating algorithms that reduce average lengths by another factor of 2-3 times by opportunistically observing the preamble.

A highly integrated UWB chipset is desired for commercial viability. The component count can be reduced by implementing a single chip transceiver with an integrated transmit/receive switch. Transmitter pulse shaping removes the need for an off-chip filter. Finally, all bias currents must be integrated and the architecture must include built-in self test algorithms and circuits resilient to power supply noise.

VII. Conclusions

Though commercial development of UWB radio has primarily focused on high-speed WPAN access, this work demonstrates that UWB is a viable candidate for low power and low data rate radio design, useful for energy-constrained applications. The all-digital transmitter achieves 43pJ/bit of energy consumption, and the duty-cycled receiver achieves 2.5nJ/bit in energy consumption over three orders of magnitude in data rate while exhibiting -99dBm of sensitivity at 100kbps for a 10^{-3} BER. Wireless unidirectional communication using the chipset was demonstrated with a streaming video application, and a complete 16.7Mb/s wireless link with a timing acquisition algorithm was established between two PCs.

Acknowledgements

This work was partially supported by the National Science Foundation under Grant ANI-0335256. The authors acknowledge the support of C2S2, the Focus Center for Circuit & System Solutions, one of five research centers funded under the Focus Center Research Program, a Semiconductor Research Corporation program. Chip fabrication was provided by STMicroelectronics.

References

[1] S. Iida, et al., "A 3.1 to 5 GHz CMOS DSSS UWB Transceiver for WPANs," ISSCC. Dig. Tech. Papers, 2005, pp. 214-594.

[2] G. R. Aiello, "Challenges for Ultra-wideband (UWB) CMOS Integration," IEEE RFIC Symposium, pp. 8-10, June 2003.

[3] A. Ismail and A. Aibidi, "A 3-10GHz LNA with wideband LC ladder matching network," ISSCC. Dig. Tech. Papers, pp. 384-385, Feb. 2004.

[4] D.D. Wentzloff, A.P. Chandrakasan, "A 47pJ/pulse 3.1-to-5GHz All-Digital UWB Transmitter in 90nm CMOS," ISSCC. Dig. Tech. Papers, February 2007, pp. 118-119.

[5] F.S. Lee, A.P. Chandrakasan, "A 2.5nJ/b 0.65V 3-to-5GHz Subbanded UWB Receiver in 90nm CMOS," ISSCC. Dig. Tech. Papers, February 2007, pp. 116-117.

[6] L. Stoica, A. Rabbachin, H. Repo, S. Tiuraniemi, and I. Oppermann, "An ultrawideband system architecture for tag based wireless sensor networks," IEEE Trans. Veh. Technol., vol. 54, pp. 1632-1645, Sep. 2005.

[7] IEEE 802.15.4 WPAN Low Data Rate Alternative PHY Layer: http://www.ieee802.org/15/pub/TG4a.html.

[8] H. Kim, D. Park, and Y. Joo, "All-digital low-power CMOS pulse generator for UWB system," Electronics Letters, pp. 1534-1535, Nov. 2004.

[9] D.C. Daly, A.P. Chandrakasan, "An Energy Efficient OOK Transceiver for Wireless Sensor Networks," IEEE J. Solid-State Circuits, vol. 42, no. 5, pp. 1003-1011, May 2007.

[10] J. Ryckaert, et al., "A 0.65-to-1.4nJ/burst 3-to-10GHz UWB Digital TX in 90nm CMOS for IEEE 802.15.4a," ISSCC. Dig. Tech. Papers, 2007.

[11] L. Smaini, et al., "Single-Chip CMOS Pulse Generator for UWB Systems," IEEE J. of Solid-State Circuits, Vol 41, 2006, pp. 1551-1561.

[12] D.D. Wentzloff, A.P. Chandrakasan, "Delay-Based BPSK for Pulsed-UWB Communication," IEEE International Conference on Speech, Acoustics, and Signal Processing, 2007.