

A Flexible, Low Power, DC-1GHz Impulse-UWB Transceiver Front-end

Ian D. O'Donnell, Robert W. Brodersen

*University of California, Berkeley
Berkeley Wireless Research Center
{ian,bwb}@eecs.berkeley.edu*

Abstract

A flexible, low power, “mostly-digital”, dc-1GHz impulse-UWB transceiver front-end is presented. By duty-cycling nearly all of the circuitry with the pulse rate, power consumption of 570 μ W (RX) and 350 μ W (TX) is presented at 1Mpulse/s with 1-bit, 1.92GSample/s sampling, 50 Ω input matching, and 42dB of gain at 1.1V. Additionally, the transceiver allows for variable modulation (PAM, PPM or both), variable gain (0 to 42dB), independent timing control for each block, variable input impedance and pulse shape generation, limited pulling of the oscillator frequency, and gain stage offset trimming to 1mV.

1. Introduction

Since the FCC report in 2002, the majority of attention on ultra-wideband circuit and system design has been focused on high-rate UWB communication in the 3GHz to 10GHz range. Operation is also permitted below 960MHz for imaging, ranging, and communication systems such as surveillance sensor networks[1]. Operation at lower frequencies is desirable for lower power consumption, better material penetration, and ease of design, but comes at the cost of increased interference from pre-existing users, and increased passive and antenna sizes. A number of applications exist for low to moderate data-rate radios where ultra-low power consumption, low cost, and the ability to do channel sounding or ranging is strongly desired. The nature of impulse signaling promises a reduction in cost and power consumption owing to the use of a simplified RF front-end design that may be more easily integrated fully on chip. The generation and reception of short pulses also lends itself to a low-Q, duty-cycled approach.

The presence of a large amount of interference in the lower UWB band warns that the cost and power benefits of impulse signaling may be mitigated by the need to do complicated signal processing or generation, or high-resolution A/D conversion to combat interference. However, the excess of bandwidth available suggests a trade-off between throughput and complexity (and hence power consumption) is possible. By moving the A/D conversion closer to the antenna and processing the signal digitally, we can take advantage of the robustness, flexibility, scalability and lower power consumption that digital can offer. The expense of this action is potentially a large increase in A/D power consumption in addition to tighter sampling clock timing requirements. An analysis of the system requirements for such a “mostly-

digital” architecture was first published in [2] and more thoroughly described in [3]. The results concluded that using a 1-bit A/D converter, without any special interference cancellation techniques, can still provide adequate throughput (on the order of 100kbps to 1Mbps) over relatively short distances (<10m) in the presence of larger interferers with power consumption on the order of a milliWatt.

To demonstrate this concept, a front-end transceiver was designed and fabricated in 0.13 μ m CMOS and an overview of the system performance was published in [4]. This paper presents additional results and elaborates upon the circuit block implementation details. The power consumption of this front-end is over 50x lower than the front-end for a similar baseband, digital transceiver published in [5] and [6] which achieves approximately 200kbps transmission. In addition, the power consumption of this “mostly-digital” transceiver is comparable to recent low power analog approaches. One result based on a simple, clocked-correlator architecture[7] realizes 1Mbps and +/-2.5cm ranging over a 1m distance in 4.7mW (TX+RX). Another analog architecture boasting 299 μ W power consumption simply thresholds the incoming pulse[8]. However, this simple radio is confined to operate only in very high-SNR environments (< 35cm operation claimed at 25kbps). The digital approach promises improved performance in the presence of multipath and heavy interference and parallelizes more easily to reduce acquisition search time than these simple analog architectures.

2. System Architecture

A block diagram of the transceiver is shown in figure 1. Incoming pulses are received by the transimpedance amplifier whose primary task is to match impedance to the antenna. Four subsequent gain stages provide variable gain and a final buffer presents the signal to a time-interleaved bank of 32 1-bit A/D converters (slicers). A delay line of length 32 in a

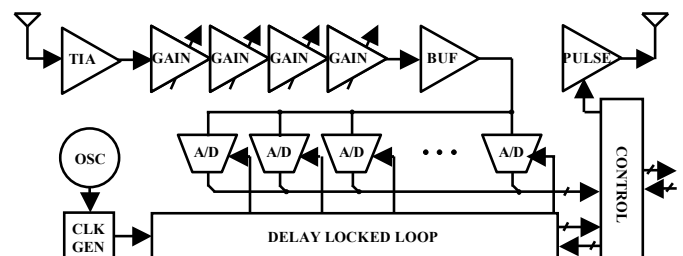


Figure 1. Block diagram of transceiver front-end.

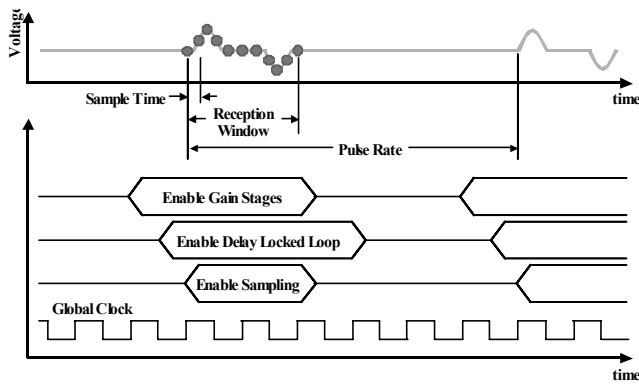


Figure 2. Duty-cycled operation of the transceiver

delay locked loop (DLL) generates the sampling clocks with approximately 520ps spacing for an effective 1.92GSample/s conversion rate. A third harmonic Pierce oscillator is used to generate the 60MHz system clock that drives the DLL and controls transceiver operation. Transmit pulses are generated using a simple H-bridge circuit at a controllable rate derived from the system clock. The front-end is highly integrated, requiring only an external crystal, LC-tank and bias resistor in addition to a power supply and antenna.

The transceiver may be operated in either a continuous or duty-cycled manner. The 60MHz global system clock is used to coordinate the timing of individual block activation and deactivation at the resolution of a clock cycle. All blocks, including the DLL and gain stages, may be duty-cycled except for the oscillator and bias circuits, which are always on. An example of duty-cycled operation for pulse reception is shown in figure 2. Note that the gain stages are powered a clock cycle prior to pulse reception to allow the circuitry to settle and the DLL is operated for a full cycle after the pulse to finish sampling and data re-alignment. The system control logic also supports pulse-position modulation by allowing for activation to occur a programmable number of cycles before or after the defined pulse rate. Pulse amplitude modulation is achieved through reversing polarity of the differential output driver. Additionally the transmitter and receiver are independently controlled, allowing for RADAR operation or communication.

3. Receiver Gain Stages

The gain stages dominate the power consumption in the receiver, consuming nearly half of the total power when continuously operated. Thus, conserving power in the gain stages is of great interest. To allow for efficient operation, the gain stages were designed to be duty-cycled; activating only during the time window in which we expect to receive a pulse. Due to their large operating bandwidth these amplifiers naturally have a low time constant, which allows them to settle quickly upon reactivation. Care must be taken with bias and offset cancellation, though, to ensure no slow time constants degrade this settling. The circuits used for the gain stages are shown single-ended in figure 3 along with the trimming details and modifications for duty-cycling. Note that although they are shown single-ended, all of the gain blocks were implemented differentially to reduce substrate and supply noise coupling.

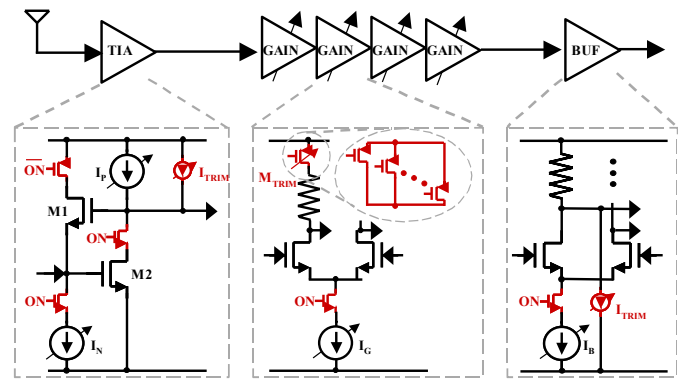


Figure 3. Gain Stage Circuits.

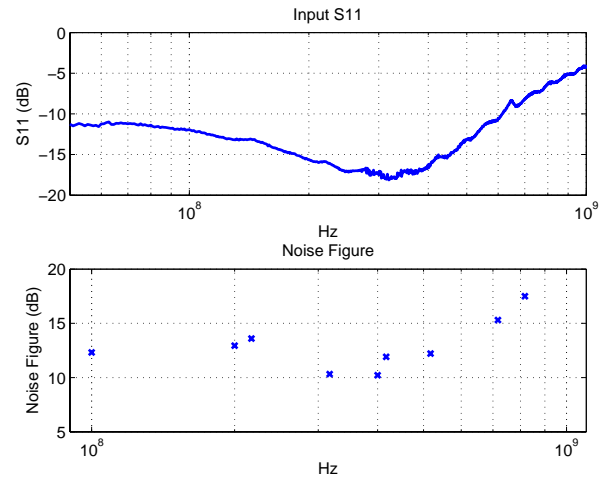


Figure 4. S11 and Noise Figure of Gain Stages.

The transimpedance amplifier dominates the power consumption of the gain circuits, taking itself fully one half the total gain stage power. As system simulations have indicated that the noise figure requirements are relaxed [3], the majority of this power consumption occurs from providing a 100Ω differential impedance at the input. This power consumption is reduced by employing shunt feedback to lower the input impedance at the expense of increased noise at the output. Measurements of the input impedance S11 parameter and the noise figure of the full gain front-end are shown in figure 4. The input impedance provides a better than -10dB S11 over 600MHz and the noise figure is approximately 12dB over the same range. The use of controllable bias allows the user to trade-off power for input impedance from 70Ω to 350Ω, (differential) providing for flexibility in the antenna/filter co-design.

To accommodate duty-cycling, switches were added to turn off the bias current sources. Also, to avoid a long recovery transient on the TIA, a duty-cycling switch is added to isolate transistor M1 and the potentially large input capacitance to which it is connected. Offset trimming is accomplished through a small digitally controlled current fed into the output. Coarse offset trimming using the digitally controlled bias current sources may be used as well, if necessary.

To achieve the gain range desired in [3], the transimpedance amplifier is followed by four variable gain stages with -2.5dB to 7.5dB of gain per stage in less than 1dB steps. These gain stages were implemented using a simple resistor loaded

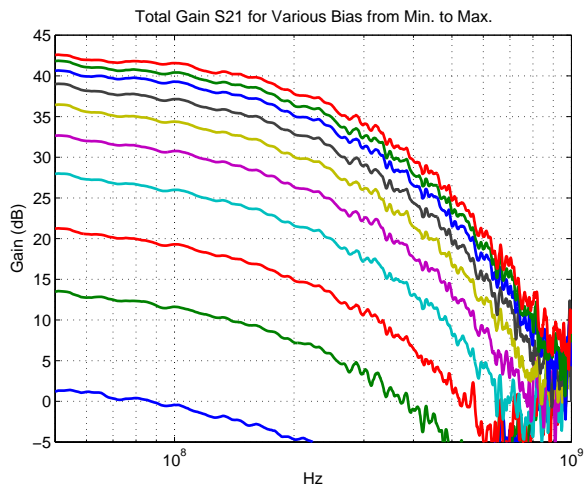


Figure 5. Total gain at various bias points from min. to max.

differential pair. A simple circuit was chosen because it is constant current when operating (reducing noise generation), clips in a friendly (i.e. non-saturating, memoryless) manner, provides adequate gain and signal swing, and is easily duty-cycled.

Offset in the variable gain blocks is trimmed by using 31 parallel PMOS triode devices at the top of the resistor to trim 10% of the resistor value. Capacitive coupling was not used between stages to cancel offset as the parasitic capacitance reduces the gain through a capacitive divider and increases the capacitive load at the output, resulting in the need for more power consumption for the same total gain value. The total gain is shown in figure 5 over the gain block bias range.

The final stage is a resistor-loaded unity-gain buffer that drives the ADC input. Resistor trimming could not be used for this stage as the resistance is much lower, requiring larger PMOS devices whose capacitance creates a frequency dependent offset when used to trim the dc offset. Hence current is pulled from the output in a manner similar to the transimpedance amplifier.

As the offset trimming circuitry is feed-forward, the transceiver requires an explicit gain trim prior to operation. This is realized by multiplexing the output of each stage to a precision comparator (off-chip for testability) and iterating through trim values until the differential output is within 1mV. All trim and bias blocks were designed to be monotonic and digitally controllable. Small pass transistor circuits are attached to each output to allow for examination without increasing the capacitive load during normal operation. To reduce potential coupling between stages through these pass transistors, the metal trace they share is grounded when not in use. Additionally, for debug a 50Ω output buffer was designed and connected through larger pass-transistors (to maintain signal bandwidth) at the TIA output and ADC input. This circuitry is shown in figure 6 and the measured available trimming range for each gain block is shown in figure 7 as a function of the digital trim code.

4. Receiver Sampling

The global system clock is created from a third harmonic Pierce oscillator circuit, requiring only an external crystal and LC tank. Two on-chip 6.4pF capacitor arrays may be used to




Figure 6. Gain stage observation/debug circuitry.

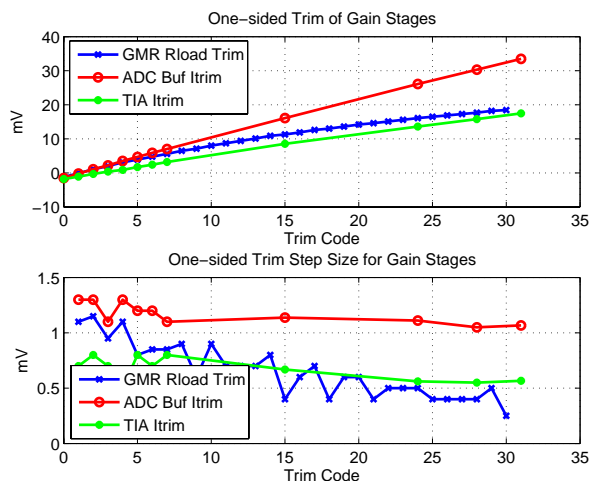


Figure 7. One-sided gain trim for each gain block.

pull the oscillation frequency in 25fF steps, allowing the oscillator to correct for crystal mismatch with < 1PPM accuracy. As tested, the frequency may be pulled +/-10PPM over the capacitor bank range. The measured cycle-to-cycle jitter standard deviation was 23ps.

The receiver sampling clocks are derived from the system clock with a delay locked loop. By combining the phases of the DLL, the A/D slicer control signals (i.e. reset, sample, evaluate) may be generated. Additionally, the phases are used to re-time the data samples; aggregating them into 32-bit words before handing them off to the digital backend.

The delay cell used is a current-starved inverter, as shown in figure 8, because it requires no static bias and there is no need to convert back to CMOS logic levels. To prevent pulse-swallowing within the delay line, the delay cell incorporates two current starved inverter blocks providing both a rising and falling edge for every incoming edge. In this manner, the duty-cycle of the incoming clock may be faithfully preserved throughout a long delay line even if there is a disparity between rising and falling delays. The DLL performance during continuous operation is shown in figure 9. Accuracy was generally measured to be +/-100ps, and per-tap cycle-to-cycle jitter was found to be less than 15ps relative to the first tap.

A DLL can be duty-cycled normally without problem simply by masking the input reference clock and delayed output appropriately. However, we also wish to duty-cycle the




Figure 8. Delay cell design.

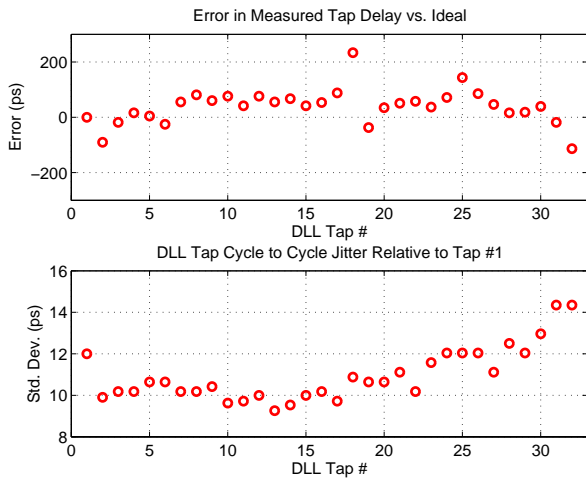


Figure 9. DLL error and jitter measurements.

DLL bias to save power at very low operational rates. A modification to the charge pump circuit, shown in figure 10, tri-states the charge pump output and separates the NMOS mirror bias to isolate the control voltages when the DLL is off. Simple logic is also used to drive the phase detector up/down outputs to ground when the DLL is off.

Measured performance of the DLL under duty-cycled operation is shown in figure 11 ranging from continuous operation to 2 updates every 128 cycles (1.6% update ratio). Less than 2% variation in delay across the delay chain was observed with the cycle-to-cycle jitter standard deviation (represented with an error bar) at less than 0.5%. The worst-case jitter, measured on the rising edge of the last tap, was also inspected up to 5 cycles after DLL reactivation to verify no transient perturbation due to re-activation is seen.

5. Transmitter

Pulse transmission is achieved through the use of a simple H-bridge circuit consisting of an NMOS pull-down and PMOS pull up, driving a differential output[9]. The pulse shape and edge rate are controlled in a limited manner by varying the timing and edge rate of the gate control voltages. Current is steered through the output in either a positive ("1") or negative manner ("0") for binary amplitude modulation. Pulse timing derives from the global system clock, and pulse position modulation also may be generated at the granularity of that clock period (16.67ns). Measured pulses are shown in figure 12. The pulse width may be varied from 1ns to 2ns and the edge rate from 150ps to 750ps.




Figure 10. Charge pump circuit.

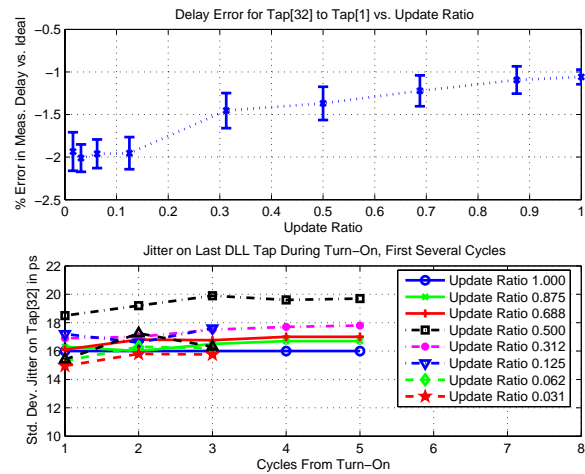


Figure 11. Duty-cycled DLL results.

6. System Results

To demonstrate system functionality pulses were generated with a 2ns width and 750ps edge and modulated with binary amplitude modulation according to an 11x11 concatenated Barker code at a 10MHz rate. These pulses were fed into a TEM horn, transmitted over 1 meter and received with another TEM horn. The resulting samples were post-processed in Matlab to obtain the cross-correlation. Figure 13 shows a snapshot of the received signal at the input of the A/D slicers over 8 pulse durations (800ns) where the large amount of interference also may be observed. The result of the cross correlation over 32 samples (~16ns) centered around the strongest response is shown below. The transceiver was controlled with a Xilinx Spartan-3 functioning as the digital backend and serial interface.

Power consumption for each block was also measured for various rates of duty-cycling based on the pulse transmission rate and is shown in figure 14. In continuous operation at 30mpulse/s the receiver consumes 3mW and the transmitter consumes 2mW. Duty-cycled at 1mpulse/s, the receiver consumes just 570μW and the transmitter 350μW.

A die photo is shown in figure 15. The die measures 2.8mm x 4.7mm and is padding dominated due to the wide parallel digital interface, separate block power supplies and test signals. Active circuit area is 2.04mm² for the receiver and 0.48mm² for the transmitter. A standard, single-poly 0.13μm digital CMOS process was used without any special mask steps (i.e. MIM cap, etc.)

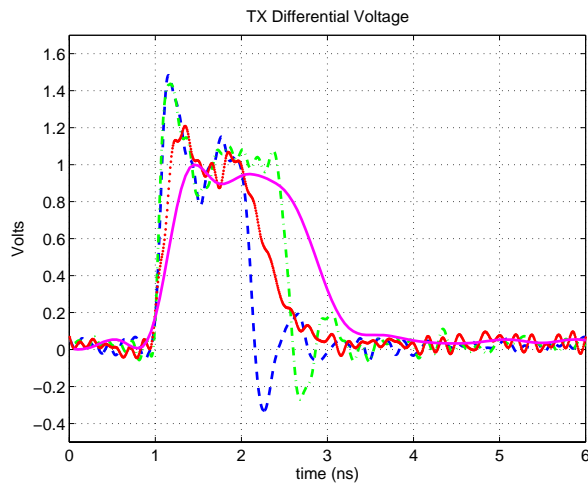


Figure 12. Measured transmit pulses.

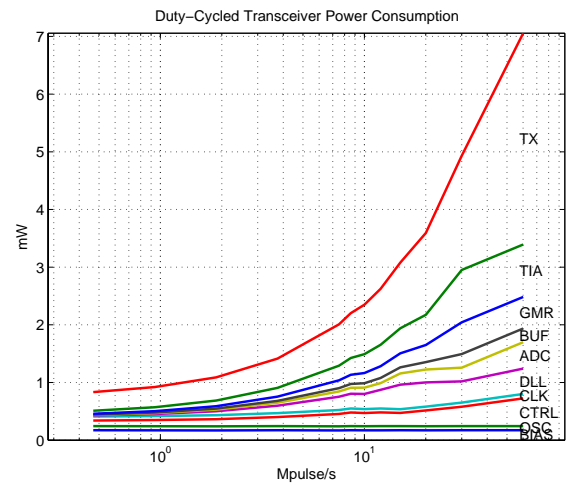


Figure 14. Power consumption vs. pulse rate.

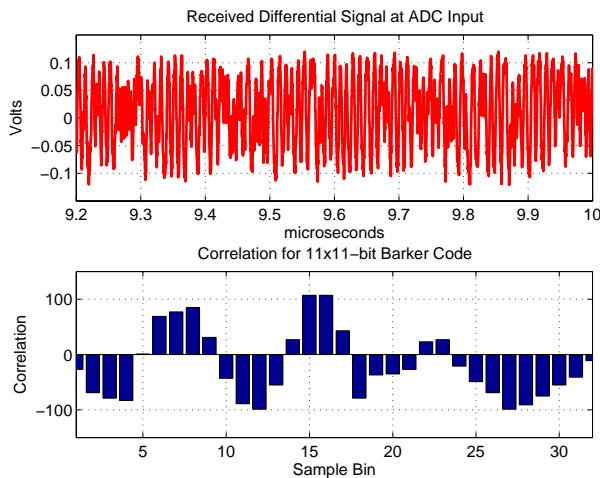


Figure 13. Received input at ADC and correlation results

7. Conclusions

A flexible, low power, highly-integrated impulse-UWB transceiver front-end was described. Implemented in $0.13\mu\text{m}$ CMOS, this transceiver may be duty-cycled to achieve a total power consumption of less than 1mW at a 1Mpulse/s rate while providing 1-bit, 1.92GSample/s sampling with 42dB of gain and a 50Ω input match at 1.1V . Additional control allows for variable gain, modulation, pulse shape generation, block timing, and trimming of gain offset and oscillator frequency.

Acknowledgement

The authors would like to acknowledge the support of the ONR (Award # N00014-00-1-0223) and the ARO (Award # 065861). The authors would also like to thank Stanley Wang for help with the transmitter design and layout, the industrial members of the BWRC, and in particular STMicroelectronics for chip fabrication.

References

[1] *First Report and Order*, Federal Communications Commission Std. FCC 02-48, Feb. 2002.

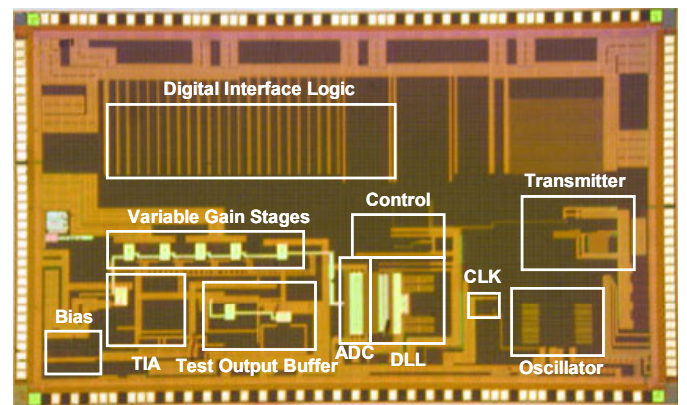


Figure 15. Die photo.

- [2] I. D. O'Donnell, "and R. W. Brodersen, "A highly-integrated, low-power, ultra-wideband transceiver for low rate, indoor wireless systems," Qualifying Exam, Dept. of Electrical Engineering, University of California at Berkeley, Nov. 2000.
- [3] I. D. O'Donnell and R. W. Brodersen, "An Ultra-Wideband Transceiver Architecture for Low Power, Low Rate, Wireless Systems," *IEEE Trans. Vehicular Technology*, vol. 54, no. 5, pp. 1623-31, Sept. 2005.
- [4] I. D. O'Donnell and R. W. Brodersen, "A 2.3mW Baseband Impulse-UWB Transceiver Front-end in CMOS," in *Proc. Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2006, pp.248-9.
- [5] F. S. Lee, D. D. Wentzloff, A. P. Chandrakasan, "An Ultra-Wideband Baseband Front-End," *Proc. of IEEE RFIC Symp.*, Jun. 2004, pp. 493-6.
- [6] R. Blazquez, P. P. Newaskar, F. S. Lee, and A. P. Chandrakasan, "A Baseband Processor for Pulsed Ultra-Wideband Signals," in *Proc. of CICC*, Oct. 2004, pp. 587-90.
- [7] T. Terada, S. Yoshizumi, Y. Sanada, T. Kuroda, "A CMOS Impulse Radio Ultra-Wideband Transceiver for 1Mb/s Data Communication and $\pm 2.5\text{cm}$ Range Findings," *Proc. Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2005, pp. 30-33.
- [8] A. Tamtrakarn, H. Ishikuro, K. Ishida, M. Takamiya, T. Sakurai, "A 1-V $299\mu\text{W}$ Flashing UWB Transceiver Based on Double Thresholding Scheme," in *Proc. Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2006, pp.250-1.
- [9] S. B. T. Wang, "Design of Ultra-Wideband RF Front-end," Ph.D. dissertation, University of California at Berkeley, Berkeley, CA, 2005.