MIT Open Access Articles

Low-power impulse UWB architectures and circuits

Citation: Chandrakasan, A.P. et al. "Low-Power Impulse UWB Architectures and Circuits." Proceedings of the IEEE 97.2 (2009): 332-352. © Copyright 2009 IEEE

As Published: http://dx.doi.org/10.1109/JPROC.2008.2008787

Publisher: Institute of Electrical and Electronics Engineers

Persistent URL: http://hdl.handle.net/1721.1/58940

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Massachusetts Institute of Technology

Low-Power Impulse UWB Architectures and Circuits

Pulsed ultrawide band provides the energy needed for a high data-rate battery-operated transceiver, and a low data-rate transceiver suitable for battery-less operation.

By Anantha P. Chandrakasan, *Fellow IEEE*, Fred S. Lee, *Member IEEE*,
David D. Wentzloff, *Member IEEE*, Vivienne Sze, *Student Member IEEE*,
Brian P. Ginsburg, *Member IEEE*, Patrick P. Mercier, *Student Member IEEE*,
Denis C. Daly, *Student Member IEEE*, and Raúl Blázquez, *Member IEEE*

ABSTRACT | Ultra-wide-band (UWB) communication has a variety of applications ranging from wireless USB to radio-frequency (RF) identification tags. For many of these applications, energy is critical due to the fact that the radios are situated on battery-operated or even batteryless devices. Two custom low-power impulse UWB systems are presented in this paper that address high- and low-data-rate applications. Both systems utilize energy-efficient architectures and circuits. The high-rate system leverages parallelism to enable the use of energy-efficient architectures and aggressive voltage scaling down to 0.4 V while maintaining a rate of 100 Mb/s. The low-rate system has an all digital transmitter architecture, 0.65 and 0.5 V radio-frequency (RF) and analog circuits in the receiver, and no RF local oscillators, allowing the chipset to power on in 2 ns for highly duty-cycled operation.

KEYWORDS | Analog–digital conversion; circuits; communication systems; digital communication; digital integrated circuits; integrated circuits; radio communication; receivers; transmitters

I. INTRODUCTION

Pulsed ultra-wide-band (UWB) communication, also referred to as impulse-radio (IR) UWB, is an active field of research with roots that can be traced back to the original Marconi spark gap radio. UWB signals are defined by the Federal Communications Commission (FCC) as having −10 dB bandwidths greater than 500 MHz. Transmitted power density is limited to be less than −41.3 dBm/MHz in the 3.1–10.6 GHz band [1]. This field has gained momentum since a 2002 change in FCC regulations 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 radio-frequency identification (RFID) tags to wireless USB at greater than 1 Gb/s [2]. For many of these applications, energy is critical because the radios are situated on battery-operated or even batteryless devices.

For high-rate systems there are several ways of modulating a UWB signal, of which IR-UWB and orthogonal frequency-division multiplexing (OFDM) are the most popular [3]. In IR-UWB, data is encoded by modulating a short pulse's position, amplitude, and/or polarity. For OFDM systems, a collection of bits is transmitted per symbol, using synchronized orthogonal carriers. WiMedia Alliance is currently working on the OFDM approach.[1] This paper focuses on IR-UWB for both high- and low-rate systems and describes circuits and architectures that enable energy-efficient operation.

Advanced complementary metal–oxide–semiconductor (CMOS) processes can efficiently implement the wide-bandwidth circuitry required for UWB with minimal area

---

[1]http://www.wimedia.org.

and few off-chip components. In comparison to narrow-band signaling, where highly tuned structures are typically required for maximum RF power efficiency, UWB circuits can be implemented using wide-band "digital" structures. This paper presents several analog and RF circuits implemented using highly digital architectures in CMOS processes.

UWB signaling can be used for low- and high-data-rate systems, and both systems are described in this paper. In addition to describing high-level tradeoffs and surveying research in the field, this paper focuses on two custom energy-efficient IR-UWB systems. The high-rate system uses binary phase-shift keying (BPSK) to achieve data rates of 100 Mb/s using direct conversion at the receiver, whereas the low-rate system achieves a maximum of 16.7 Mb/s using binary pulse-position modulation (PPM) with noncoherent energy detection. The latter can efficiently scale to lower data rates. While data rate requirements and potential applications for the two systems are different, the fundamental goal of maintaining low power remains the same. Techniques for extending the communication distances for the low-rate system are also discussed.

II. HIGH-RATE IR-UWB ARCHITECTURES

The consumer electronics industry is exploring the use of UWB communications as a short-range high-data-rate radio technology, to complement longer range radio technologies such as Wi-Fi, WiMAX, and cellular wide-area communications. UWB communications can be used in a wireless personal-area network to achieve ad hoc connectivity between portable devices or to host devices within the immediate area, eliminating the need for wires [4]. Using UWB for high-data-rate last-meter wireless links requires that UWB radios be integrated into battery-operated devices such as mobile phones and handheld devices. Consequently, there is a need for an energy-efficient UWB transceiver. Section III will discuss the use of parallelism in various blocks of the baseband, specifically the analog-to-digital converter (ADC)

and digital baseband processor, as highlighted in Fig. 1, to reduce energy consumption while maintaining the 100-Mb/s performance required for high-rate UWB communications.

The high-data-rate baseband ADC and digital processor presented in this paper target a custom IR-UWB system [3], [5]. BPSK-modulated Gaussian pulses are transmitted at a pulse repetition frequency (PRF) of 100 MHz in one of fourteen 500-MHz-wide channels within the UWB band [6]. The transmitter utilizes the exponential properties of a bipolar junction transistor to generate pulses that accurately approximate a Gaussian shape [7]. An alternative approach includes the use of a programmable pulse modulator at the transmitter [2].

The direct-conversion RF front-end of the receiver, shown in Fig. 1, is composed of a 3.1–10.6 GHz unmatched low-noise amplifier (LNA), on-chip UWB filter, RF single-ended-to-differential converter with integrated RF notch filter, wide-band mixers, and baseband 250-MHz buffers and filters [8]. After down-conversion, the received complex baseband signal is sampled at a Nyquist rate of 500 MS/s by two 5-bit ADCs. For real-time demodulation, the digital baseband processor must meet this throughput.

An important consideration in the receiver architecture is determining the partition between the analog and digital domain. Acquisition for synchronization, channel estimation, and demodulation is done entirely in the digital domain. This mostly digital architecture was chosen rather than a partial analog approach [9], since it allows for deep voltage scaling and greater flexibility. Other advantages include the simplification of the analog elements in the transceiver, and the possibility of exploring digital channel adaptability and recovery. Furthermore, performing synchronization in the digital domain eliminates the need to feed a signal from the digital domain to the clock generation subsystem [10]. Only the automatic gain control is fed back to the analog domain. This is unlike other systems [2], where the digital baseband processor controls the phase of the ADC clock. Mostly digital architectures have also been used in systems [11], [12] that focus on low-rate IR-UWB communications in the lower


Fig. 1. *UWB direct conversion receiver block diagram for high-data-rate system. Parallelized baseband is enclosed in dash-line box.*

UWB band from 0 to 960 MHz. Reference [11] uses a high-speed 1-bit ADC to sample the entire band with no mixers in the front end.

An interesting property of UWB communications is that the transmitted power is quite low, and thus the overall transceiver power tends to be dominated by the receiver. Accordingly, the next section focuses on how to reduce the power consumption of the blocks in the receiver baseband.

III. HIGH-RATE IR-UWB BASEBAND

Parallelism is exploited in both the ADC and the digital baseband processor in order to achieve the 100-Mb/s throughput with minimum power consumption. Time-interleaving allows the use of an energy-efficient successive approximation register (SAR) architecture for the ADC [13], while in the digital baseband processor, time-interleaving enables operation using an ultra-low-voltage supply [14].

A. 5-bit Analog-to-Digital Converter

The 250-MHz down-converted pulses require a 500-MS/s Nyquist converter, but the required resolution is limited to 4–5 bits [15]. Flash ADCs are the typical choice for this high-speed low-resolution regime. A flash converter compares the input, in parallel, to every possible threshold voltage and determines the binary output in a single clock cycle. This use of voltage parallelism enables the highest speed ADC operation, but requires an exponential growth in the number of comparators with the resolution. This undesirable complexity characteristic has long motivated the choice of other architectures. The SAR topology has only a linear growth in the number of comparisons with the resolution; however, it computes each bit of the digital output sequentially and therefore requires multiple clock periods to resolve a conversion, limiting conversion speed. Time interleaving [16] uses parallel channels, sampling at fixed time-intervals to increase

the conversion time of any single channel, permitting use of the energy-efficient SAR architecture for this high-speed application [17]. An energy comparison between the flash and time-interleaved SAR architecture is presented in [18].

One limitation to the general use of parallelism is the requirement of independent processing from sample to sample. Successive samples of any true Nyquist converter, however, should be assumed to be completely independent of each other. Thus, processing the samples in parallel should give identical results to processing them serially. In practice, however, mismatches between channels can negatively impact ADC performance. The three primary mismatch concerns are offset, gain, and timing skew. The design of the time-interleaved SAR ADC is presented below, specifically addressing our solutions to mismatches.

1) Top-Level Architecture: The SAR algorithm requires one period to decide each of the output bits plus one period for sampling. With six time-interleaved channels, the internal channel clock period matches the overall sampling clock. Thus, only one clock needs to be generated and distributed. Besides easing clock distribution requirements, this also minimizes timing skew between channels. A balanced layout for this single sampling clock is sufficient to reduce errors arising from timing skew to below the 5-bit level.

The top-level block diagram of the six-channel ADC is shown in Fig. 2. Synchronization is performed by passing a start token that signals when a channel should begin sampling. This keeps the overhead associated with time-interleaving to a minimum.

2) Channel Circuits: Each channel is composed of a capacitive digital-to-analog converter (DAC), a comparator, and digital control logic, often referred to as the SAR itself (Fig. 2). The DAC is the split capacitor array [18], which features decreased switching energy and faster
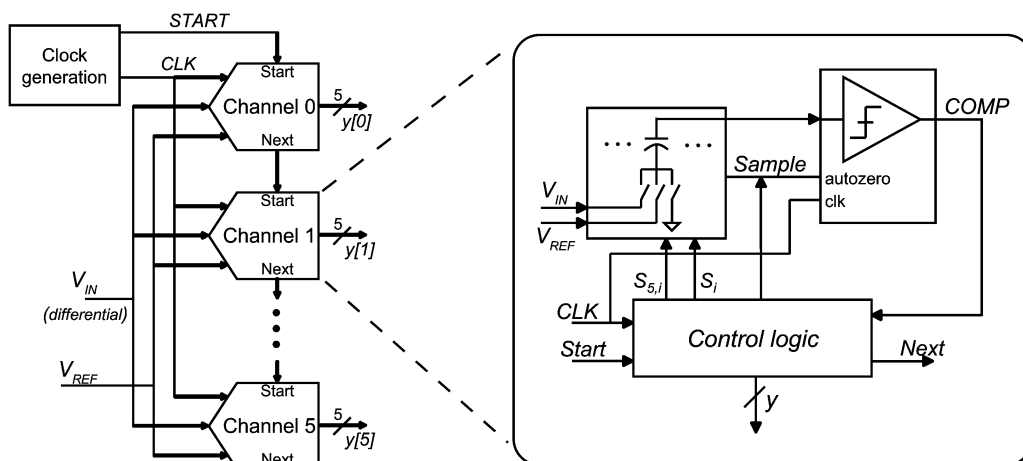

Fig. 2. *Block diagram of six-way time-interleaved SAR ADC and SAR channel.*

Fig. 3. *Die photograph of ADC.*

Technology	65-nm CMOS 1P6M
Supply Voltage	1.2 V
Sampling Rate/Resolution	500 MHz/5 bit
SNDR (f_{in} = 239 MHz)	26.1 dB
Power (analog/digital)	2.86/3.06 mW
Active Area	0.65 mm × 1.4 mm

switching speed than the conventional binary weighted capacitor array. Gain mismatch between channels is limited by capacitor matching, and the unit capacitor size is thereby chosen conservatively.

The comparator uses a two stage auto zeroed preamplifier and a regenerative latch. The preamplifiers reduce the large offset voltage of the latch to below one-quarter of the least significant bit (LSB) voltage when referred to the input of the entire comparator chain, sufficient to limit offset mismatch. All of the transistors in the comparator have longer than the minimum channel length in order to improve matching and output impedance.

3) Implementation and Measured Results: The ADC has been fabricated in a 65-nm CMOS process. A photograph of the 1.9 × 1.4 mm die is shown in Fig. 3. The input and clock paths use a fully balanced layout in the middle of the die. The effect of mismatch between channels can be seen in the fast Fourier transform (FFT) in Fig. 4. Distortion from the measured $0.3V_{LSB}$ offset variation appears as spurs (e)–(f) at multiples of the channel sampling frequency. Spurs (a)–(d) arise from timing and gain errors, dominated by the former in this implementation. The
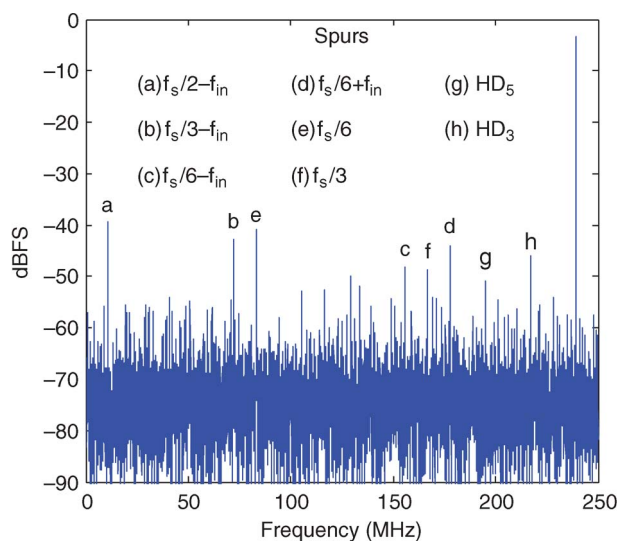
measured gain error is 0.9%. The ADC achieves full Nyquist operation, with the effective number of bits dropping from 4.5 at dc to 4 at Nyquist. The measured 6-mW power consumption is split roughly evenly between the analog and digital supplies. The ADC performance summary is listed in Table 1.

B. Digital Baseband Processor

The digital baseband processor receives two 500-MS/s signals from the ADCs and performs acquisition and demodulation. The received packet, shown in Fig. 5, is composed of a sequence of 500 MHz bandwidth BPSK pulses and can be divided into two sections: preamble and payload. The preamble contains repetitions of an $N_C = 31$-bit Gold code. A Gold code is a type of pseudo-noise sequence with desirable autocorrelation properties for packet synchronization. The number of samples per pulse in the preamble (N_s) is given by the sampling rate (500 MS/s) divided by the preamble PRF. Ideally, the preamble PRF, and consequently N_s, are selected to enable channel estimation without having to compensate for intersymbol interference (ISI) [19]. For the channel models provided in [20], $N_s = 20$ allows ISI free channel estimation for CM1 to CM3 and negligible ISI errors in CM4 [21]. The payload contains the actual data and is sent at a PRF of 100 MHz for a 100-Mb/s data rate with no channel coding.
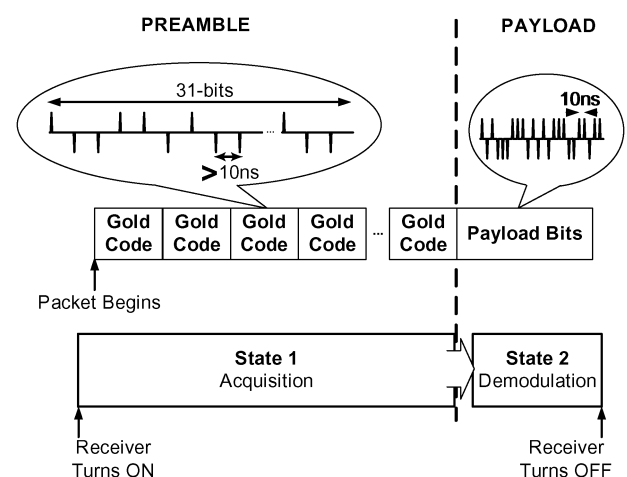

Fig. 4. *FFT of 239 MHz input with dominant spurs labeled.*

Fig. 5. *UWB physical-layer packet format.*

1) Algorithm: Since the BPSK modulation can be interpreted as a direct-sequence code-division multiplex signal, IR-UWB and code-division multiple access (CDMA) systems have many common characteristics. The main difference is that the duty cycle of the UWB signal is less than 100%. For that reason, the architecture of a UWB transceiver utilizes components from classic CDMA transceivers such as the acquisition, synchronization, and tracking algorithms.

Acquisition involves the computation of the correlation between the received noisy preamble and a local clean template of the Gold code sequence. Initially the local Gold code sequence is not aligned with the received preamble, resulting in a low correlation value. Different delays of the local template are correlated against the received preamble until the result of the correlator meets a predetermined threshold. This threshold is chosen to ensure a small probability of false detection, leveraging the good autocorrelation properties of the Gold code (a large correlation is achieved only when the local template and the received preamble are aligned within a granularity of the 2-ns sampling period). Details of this acquisition algorithm are described in [22]. Upon synchronization, the correlation result is used to provide the channel estimation.

Following synchronization, the baseband processor can demodulate the payload using a partial RAKE receiver to collect and optimally combine the signal energy received on the multiple echo paths using the tap gains determined by the channel estimation. The number of taps can be fixed or determined by a programmable threshold [19]. The signal is tracked in the payload using a delay-locked loop (DLL) to account for offsets between the ADC clock frequencies at the transmitter and receiver. Carrier frequency offsets can be estimated and corrected using a Costas loop, as described in [5].

The total energy spent on receiving the UWB signal can be divided into two components: acquisition (preamble) energy and demodulation (payload) energy. The acquisition energy can be seen as overhead energy since it does not contribute during data demodulation. The amount of overhead energy per packet should be minimized particularly for packets with small payloads. In the baseband processor, the majority of this overhead energy is consumed during acquisition by the correlation computation. Since each point of the correlation can be computed independently, this overhead energy can be significantly reduced by exploiting parallelism.

2) Ultra-Low-Voltage Operation: The correlation is computed using the correlator shown in Fig. 6. The energy of the baseband processor can be reduced by aggressively scaling down its voltage supply (V_{DD}) such that the correlator operates near its minimum energy point [23]. The minimum energy point occurs because the total energy


Fig. 6. *Architecture of correlator used during acquisition.*

per operation is composed of active energy (E_{active}) and leakage energy (E_{leakage})

$$E_{\text{total}} = E_{\text{active}} + E_{\text{leakage}}$$
$$= C_{\text{eff}}V_{DD}^2 + I_{\text{leak}}V_{DD}T_{\text{delay}} \qquad (1)$$

where C_{eff} is the average effective switched capacitance of the entire circuit, including the average activity factor. From (1), we see that lowering V_{DD} decreases the active energy. While reducing V_{DD} reduces the leakage power ($I_{\text{leak}}V_{DD}$), it also increases the delay (T_{delay}) of the gates. When V_{DD} is above the threshold voltage of the device, T_{delay} increases linearly with decreasing V_{DD}, and there is no significant change in the leakage energy; however, when V_{DD} drops below the threshold voltage of the device, both T_{delay} and leakage energy increase exponentially with decreasing V_{DD}. Since the active energy and leakage energy scale in opposite directions as V_{DD} decreases, a minimum energy point occurs in the subthreshold region.

The energy profile of the correlator, obtained from simulation, is shown in Fig. 7 and indicates that the minimum energy point occurs at 0.3 V. For real-time


Fig. 7. *Energy profile of the correlator over varying voltage supplies.*

acquisition and demodulation of a UWB packet, the baseband processor must perform signal processing with a throughput of 500 MS/s. This can be achieved by a single correlator operating at a frequency of 500 MHz with a much higher voltage than 0.3 V, but Fig. 7 shows that this is not energy efficient. To reduce energy, the correlators should operate with a voltage supply as close as possible to its minimum energy point and utilize a bank of parallel correlators to maintain the required throughput.

Furthermore, to minimize the complexity of parallelism, the number of parallel correlators in the bank should be chosen such that an integer number of pulses (f), or $f \cdot N_s$ samples, are processed per clock cycle. If this restriction is not met, the Gold code cannot be hardwired to the correlators (Fig. 6). Instead, additional complexity is necessary to associate the Gold code input to the correlators. To process $f = 1$ pulse per cycle requires a bank of $1 \cdot N_s = 20$ correlators. These correlators only need to operate at 25 MHz, which allows the voltage supply to be scaled down to 0.4 V. Although this is slightly above the optimum V_{DD} of 0.3 V, it does not cause a significant energy penalty since the minimum energy point is shallow. By operating at an ultra-low-voltage of 0.4 V rather than 1 V, the energy per operation is reduced by approximately 6× [14].

Additional parallelism can be used to shorten the acquisition time. The bank of $1 \cdot N_s = 20$ correlators can be replicated by N_C, to compute all points of the corre-lation simultaneously [24]. This reduces overhead energy of the *entire* receiver by minimizing the on-time of the RF front-end [8] and ADC [25].

3) Architecture: To demonstrate the impact of parallelism and aggressive voltage scaling on power consumption, a simplified version of the baseband processor described in [19] was implemented. This simplified baseband processor uses a fixed five-finger partial RAKE with a maximum delay of five samples for demodulation, capturing the energy in a 10-ns interval that includes the most powerful multipath components. The maximal-ratio combining (MRC) in the RAKE receiver is also parallelized by a factor of four so that it operates with the same ultra-low-voltage supply and operating frequency as the correlators. This also reduces the energy required for demodulation. Since much of the computation required for the Costas loop is already done by the correlators, specifically the accumulation, it requires minimal additional computation. Although it was not included in the simplified design, the estimated impact of the Costas loop on power is negligible. This architecture was verified on the field-programmable gate array (FPGA) of the discrete prototype described in [6] and demonstrated real-time demodulation.

The highly parallelized implementation with a total of 620 correlators (31 banks of 20 parallel correlators) and four RAKE MRCs is shown in Fig. 8. The number of correlators per bank is dictated by the maximum frequency
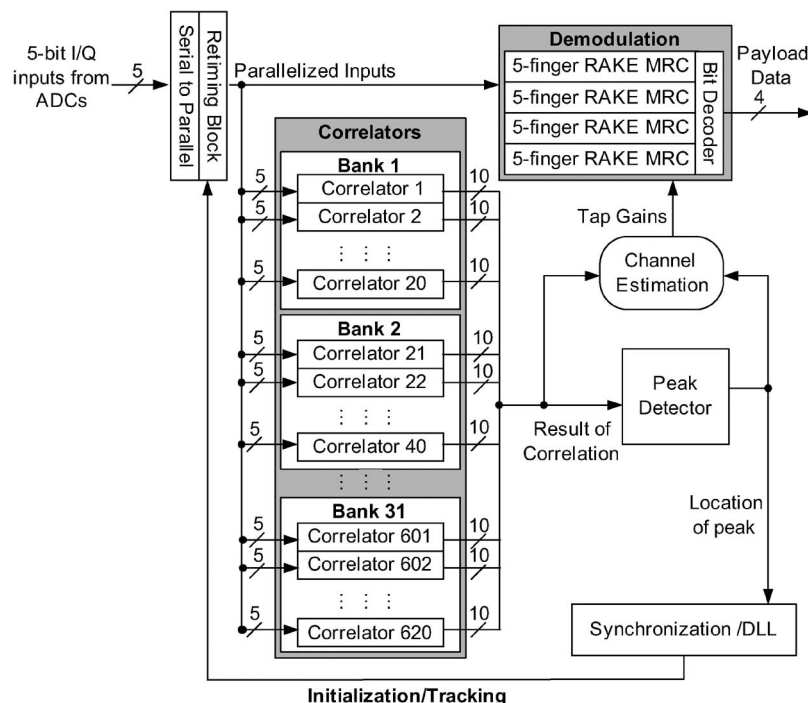

Fig. 8. *Architecture of the highly parallelized energy-efficient digital baseband processor.*

| Technology	90-nm CMOS 1P7M
Supply Voltage	0.4 V
Frequency	25 MHz
Standard cell count	281,260
Power (acq/dem)	7 mW / 1.7 mW

Fig. 9. *Die photo and performance summary of digital baseband processor.*

of the correlators near its minimum energy point, while the number of banks is dictated by the length of the Gold code sequence (N_C).

4) Implementation and Measured Results: The baseband processor is implemented in a standard-V_T 90-nm CMOS process and demonstrates 100-Mb/s operation at 0.4 V with an operating frequency of 25 MHz. A photograph of the die photo and summary of the performance metrics are shown in Fig. 9. Note that only 23% of the die area is active as the design was pad-limited.

The breakdown of the energy per bit consumed by the digital baseband processor is shown in Fig. 10. The average energy overhead consumed during acquisition is fixed for a packet. Thus, the shorter the payload, the greater the overhead energy per bit as the overhead energy is amortized over fewer bits. For a 4-kbit packet, using the measured power for acquisition and demodulation provided in Fig. 9, the average energy per bit consumed by the digital baseband processor is 20 pJ, with 3 pJ for acquisition and 17 pJ for demodulation.

C. Summary

An energy-efficient baseband for a UWB radio has been presented. Parallelism has enabled very low power consumption in this high-performance application. The low complexity but high latency successive approximation reg-

Fig. 10. *Breakdown of energy per bit consumed by the digital baseband processor.*

ister ADC architecture is combined with time-interleaving to achieve the desired throughput and performance in deep-submicrometer CMOS. The highly parallelized packet acquisition leads to a significant reduction in operating voltage. The baseband presented here can be integrated in a single-chip solution in a highly energy-efficient manner by increasing the number of time-interleaved ADC channels. With 20 time-interleaved channels, each channel would directly feed one set of correlator banks, and each channel could operate at a lower conversion rate with reduced supply voltage for further energy savings. A highly interleaved SAR ADC with the optimum degree of parallelism to maximize energy efficiency is presented in [26].

IV. LOW-RATE IR-UWB ARCHITECTURES

Many low-power applications such as sensor networks and industrial monitoring require only low-rate (below 1 Mb/s) communication and have strict average power constraints. Key standards for this application space include Bluetooth and IEEE 802.15.4. In March 2007, an amendment to the low-data-rate 802.15.4 standard was approved, adding UWB signaling as a physical layer option [27]. Any UWB pulse shape is supported by the standard, so long as it matches sufficiently close to the provided reference pulse. The UWB physical layer operates at a mandatory data rate of 851 kb/s and optional data rates of 110 kb/s, 6.81 Mb/s, and 27.24 Mb/s. The amendment is designed specifically for ultra-low-power radios and includes an optional mode to allow for noncoherent communication.

The choice of noncoherent versus coherent modulation is a key system-level tradeoff. Although coherent signaling schemes utilize bandwidth more efficiently and achieve better sensitivity than noncoherent schemes, for low-data-rate systems these benefits come at the cost of degraded energy efficiency when normalized by data rate. This is due to the power cost of phase tracking hardware for coherent architectures. For instance, coherent systems typically require frequency accuracies less than 100 ppm, whereas noncoherent signaling allows for frequency accuracies over 1000 ppm. This enables the use of highly digital architectures for both transmitter and receiver.

A key opportunity of high signal bandwidth communication is that high instantaneous data rates are inherent to the system. For low-data-rate systems, this means that the high-rate transceiver can be aggressively duty cycled. This allows for a general trend of improved energy efficiency compared to narrow-band low-rate radios. This is because for high-rate radios, the fixed cost of analog/RF bias currents is amortized over more bits/second as data rate increases, resulting in a tradeoff between energy efficiency and instantaneous data rate. Fig. 11 shows recently published receiver energy/bit as well as the energy/bit of the receiver presented in this paper [28]. The energy per bit for this work was calculated as the sum of the receiver energy/bit plus the leakage power component, which

Fig. 11. *Receiver energy/bit values versus data rate for UWB and narrow-band receivers recently published at ISSCC [28].*

Fig. 13. *Three-channel frequency plan and narrow-band interferers.*

causes the energy/bit to rise at lower data rates. For data rates below 100 kb/s, our work on low-rate UWB systems targets energy efficiency that is orders of magnitude better than historical low-rate narrow-band systems.

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 [29], [30]. 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, transmit/receive (Tx/Rx) switch, and transmit band-select filter.

The architecture of the transceiver chipset presented in this paper is shown in Fig. 12 [28], [31]. The transceiver chipset is not 802.15.4a compliant, but the target system shares many specifications with the standard. Binary PPM is used to encode the transmitted data. The PPM signal is transmitted in one of three channels in the 3.1-to-5 GHz band, as shown in Fig. 13. Three channels are used to avoid potential in-band interferers and to add frequency diversity for multiple users. The receiver is a noncoherent energy detector that compares the received energy of two adjacent timeslots. This type of receiver squares the incoming signal at RF; therefore no local oscillator or phase-locked loop (PLL) is required for down-conversion.

V. LOW-RATE IR-UWB TRANSMITTER

In general, low-rate and high-rate IR-UWB transmitters share many common attributes and are typically inter-


Fig. 12. *Custom chipset transceiver architecture for the low-data-rate system.*

changeable from a functional perspective. However, in practice, low-rate UWB transmitters often use different circuits and systems that allow for reduced power consumption compared to high-rate transmitters but at the cost of reduced performance. This section first outlines commonly used techniques to synthesize UWB pulses. Next, an all-digital transmitter for low-rate IR-UWB systems is presented [31]. Finally, this section describes FCC constraints on peak and average power and presents techniques to reduce peak power constraints on low-rate UWB systems.

A. Architecture

Pulse-based UWB transmitters generally synthesize pulses in the 3.1-to-10.6 GHz band using one of two techniques.

- *Up-Conversion Pulse Synthesis:* The first technique involves generating a pulse at baseband and up-converting it to a center frequency in the UWB band by mixing with a local oscillator (LO) [2], [32]. Note that the transmitter may not have an explicit mixer that performs the up-conversion operation. For instance, a simple switch can either enable or disable the output of a LO, thus effectively mixing the RF signal with a rectangular baseband pulse. If used in an active LC oscillator, the turn-on and turn-off transients may be exploited to produce nonrectangular pulse shapes [33].

- *Carrier-Less Pulse Synthesis:* The second technique involves generating pulses that directly fall in the UWB band without requiring frequency translation [34]–[36]. The pulse width for these types of transmitters is usually defined by delay elements that may be tunable or fixed. A baseband impulse may excite a filter that shapes the pulse [37], [38], or the pulse may be directly synthesized at RF with no additional filtering required [39], [40].

One advantage of carrier-less techniques over traditional mixer-based architectures is that the carrier frequency generation is inherently duty cycled; that is, RF energy is only generated when it is required. A disadvantage of this approach is that an integrated down-converting receiver typically cannot share the RF generation circuits, and therefore must have its own LO.

Pulse-based transmitters can additionally be categorized in terms of how pulses are delivered to an antenna for propagation.

- *Analog Power Amplification:* This technique involves amplifying generated pulses with an analog power amplifier, often biased as class A or class AB [7], [41]. This technique is the most robust from a communication viewpoint, as amplitude modulation and pulse shaping are readily achievable. Unfortunately, this technique typically dissipates large static bias currents and can overwhelm the power budget of energy-starved systems.
- *Digital Buffering:* Instead of amplifying small signal pulses, digitally generated pulses can be buffered to drive a 50 Ω load and associated parasitics using static CMOS logic gates [42], [43]. This technique offers excellent energy efficiency, as no static bias currents are consumed. However, linear amplitude modulation and pulse shaping are more difficult to achieve.

The transmitter considered in this work combines a series of equally delayed edges to form a single RF pulse directly in the band of interest. The combined edges are then buffered through a series of digital inverters to drive a 50 Ω antenna. All blocks, including the RF pad driver, use full-swing static CMOS digital circuits, and no analog bias currents are required. A block diagram of the transmitter is shown in Fig. 14 [31].

The transmitter is designed to support PPM so that a noncoherent, energy detection receiver may be implemented. A tapped delay line is used to generate a series of edges after each rising edge of the *start pulse* signal. Pulses are synthesized by combining a programmable number of the edges from the delay line. Twenty-five edges are required to synthesize pulses in the highest channel, therefore a 32-stage delay line is used, from which 30 edges are made available for combination. Individual edges are selected by ANDing them with a 30-bit mask register. The 30 edges are combined using two time-interleaved 15-edge combiners based on the work in [44].

The edge combiner block toggles its output when an edge is received on any of its inputs. The generated pulse appears like a digital clock that is on for only a few cycles. By making both the delay-per-stage and number of combined edges programmable, the pulse spectrum may


Fig. 15. *Timing diagram of the edge combination.*

be precisely controlled without requiring an RF local oscillator. The pulse is buffered by a digital pad driver, and the output is filtered by a UWB band-select filter that directly drives a 50 Ω antenna. An FPGA performs the PPM modulation of the *start pulse* signal, in addition to the PN sequence generation for the scrambling of the output spectrum. It is also used to implement the digital algorithm for calibrating the 8-bit delay-per-stage code of the delay line. The entire transmitter is clocked at 33 MHz.

A timing diagram of the edge combination is shown in Fig. 15. The output edges *out[1:30]* are masked, and only the selected edges are combined by the two interleaved combiners. The *Comb 1* and *Comb 2* signals are XORed, synthesizing the up-converted pulse with a spectrum centered in the desired channel.

B. Delay-Based BPSK Scrambling

BPSK is the most common modulation found in literature for IR-UWB systems, as it offers an inherent 3 dB advantage in signal-to-noise ratio (SNR) over amplitude and position modulation schemes. However, binary-PPM transmitters can often reduce circuit complexity over their BPSK counterparts, as pulse generation circuits do not necessarily require differential analog amplifiers, mixers, or large-area baluns to generate the required biphase pulses [7], [42]. Furthermore, noncoherent receiver implementations can be used, which often have higher energy efficiencies than coherent architectures (Fig. 11) [45].

Unfortunately, a binary-PPM pulse spectrum will contain large spectral lines, even when modulated with random data [46]. This results in a PPM transmitter having to lower its power relative to a BPSK transmitter in order to


**Fig. 14.** *Block diagram of the all-digital transmitter [31].*

Fig. 16. *BPSK inversion and delay-based inversion. The reference four-cycle pulse is overlaid with the inverted BPSK and DB-BPSK pulses for comparison.*

meet the FCC spectral mask. Typically higher order PPM, or BPSK scrambling in addition to PPM, is used to eliminate these tones. Because BPSK decouples the scrambling problem from the modulation, it is typically preferred over higher order PPM, which adds complexity to the receiver hardware and synchronization algorithm. It would therefore be advantageous to combine the spectral properties of BPSK signals with the implementation simplicity of an all-digital PPM transmitter. This can be achieved with delay-based BPSK (DB-BPSK) modulation [47].

An illustration comparing DB-BPSK with conventional BPSK is shown in Fig. 16 with the reference pulses in black and inverted pulses in gray. Conventional BPSK pulses are generated by inverting the reference pulse as shown on the left. DB-BPSK pulses are generated by instead delaying the reference pulse by half of the RF tone period, as shown on the right. Conceptually, this delay appears to have the same effect as inversion, except for the half-cycles at the beginning and end of the pulses. As the number of cycles per pulse increases, the half-cycle extensions become a smaller fraction of the total pulse energy. Therefore, one would expect the performance of DB-BPSK to approach that of BPSK for an increasing number of cycles per pulse. For more than nine cycles per pulse, the spectrum of DB-BPSK pulses is similar to that of BPSK pulses in the main lobe of the

spectrum, making DB-BPSK suitable for scrambling a PPM pulse spectrum or communicating data.

A delay-line-based architecture is ideal for implementing DB-BPSK modulation with minimal overhead, as the delay per stage is already tuned to be half of an RF period. Therefore, only one additional delay element and the ability to select edges from the delay line are required.

C. Transmitter Circuit Design

Each delay stage of the transmitter uses a differential architecture, as shown in Fig. 17. The delay is digitally controlled with an 8-bit code using a combination of binary-weighted current starving and a capacitor bank. During normal pulsed operation, the final stage of the delay line is disabled. Enabling the last stage configures the delay line as a free-running oscillator, used only for calibration of the delay per stage.

Cross-coupled inverters are used for regeneration of the edges, serving two purposes: 1) to reduce the mismatch between rise and fall times of the differential signals and 2) to suppress common mode latching in the ring when configured as an oscillator because there are an even number of stages.

The output edges of the delay line are fed to the edge combiner, which directly feeds the RF pad driver. The RF pad driver is essentially a series of digital inverters with added functionality for varying the gain and reducing leakage currents. A schematic of the pad driver is shown in Fig. 18. The transistors in the final stage are sized to maximize the efficiency of the driver when driving a 50 Ω antenna. Because the transistors are large, leakage is a concern in the final stage. A well-known technique of stacking NMOS devices is used to reduce leakage. In this case, the stacked NMOS transistors result in a leakage current five times lower than that of a single NMOS device with equal pulldown strength. This is due to the intermediate node between stacked transistors floating above ground, increasing the threshold voltage of the top
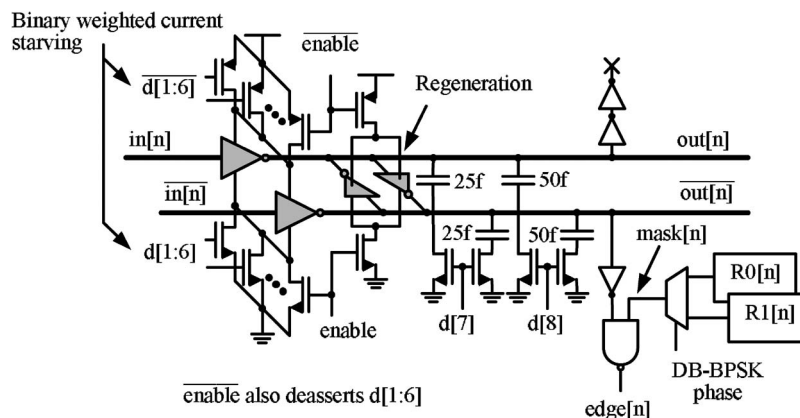

Fig. 17. *Schematic of the differential delay stage with binary weighted current starving and capacitor bank networks.*
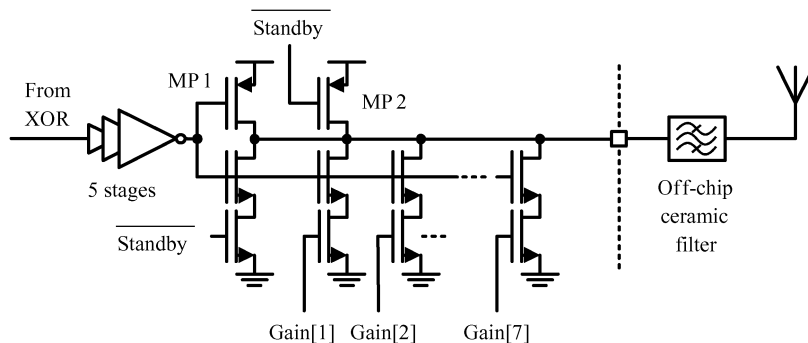
Fig. 18. *Schematic of the digital pad driver with linear-in-dB gain setting and standby mode.*

transistor in the stack due to the body effect, thus reducing the off-current.

During pulsed operation, the output is pulled high by a strong PMOS pullup (*MP1*) and pulled low by the stacked NMOS pulldown network. The power of the output pulse is varied by adjusting the drive strength of the pulldown network through the digital *Gain[1:7]* controls. The widths of the NMOS devices in each of the branches of the pulldown network are weighted in order to produce a linear-in-dB power adjustment. The driver incorporates a high-impedance standby mode, where the output is pulled to V_{DD} through a weak PMOS device (*MP2*). The output is weakly held high in order to eliminate transients that would otherwise occur when coming out of standby mode.

The pulses generated by this transmitter have spectral content centered around both the desired channel center frequency and at dc. In order to eliminate the spectral content at dc, an off-chip 3 x 1.5 x 1 mm ceramic filter with the equivalent frequency response of a seventh-order high-pass Butterworth filter is used. This filter could also be shared with the receiver, serving as the band-select filter when the transmitter driver is tri-stated.

D. Measurement Results

The transmitter is fabricated in a standard digital 90 nm CMOS process with a metal–insulator–metal (MIM) capacitor option. The design uses a combination of components from a standard cell digital library, as well as full-custom layout. The receiver is wire bonded in a 24-lead quad flat no-lead (QFN) package. All measurement results presented in this section are of the packaged chips with the off-chip ceramic filter.

In order to center the pulse spectrum in the desired channel, the delay/stage is calibrated offline by configuring the 32-stage delay line as a free-running oscillator and measuring the frequency of oscillation. This is performed by disabling the driving stage and enabling the last stage in the delay line which feeds back to the input. Ideally, the ring oscillates at the pulse spectrum center frequency divided by 32. The pulse center frequency is thus predicted by

measuring the frequency of the ring oscillator by counting the ring cycles for a fixed period of time. To calibrate the delay, the digital control is adjusted in a successive approximation algorithm, which completes in a maximum of 62 μs. The difference between the center frequency of the pulse spectrum and the ring frequency multiplied by 32 for each value of the delay control was measured and is less than ±25 MHz, or 6000 ppm, which is sufficient for communication with a noncoherent receiver.

A transient waveform illustrating the effects of DB-BPSK is shown in Fig. 19. The two phases of the pulses are superimposed for easy comparison. While these pulses appear to be inversions of each other, by focusing on the beginning of the pulses, it is apparent that one is delayed by half of the RF cycle period relative to the other.

The measured spectrum for PPM alone, and PPM with DB-BPSK scrambling enabled, is shown in Fig. 20. The PPM signal is modulated by random data; however, a line spectrum is still produced as predicted. By applying DB-BPSK scrambling to the same PPM modulated signal without making any other changes in the transmitter, the lines are completely eliminated in the main lobe. Notice


Fig. 19. *Measured DB-BPSK pulses with a center frequency of 4.05 GHz. The two phases of the DB-BPSK pulses are superimposed, highlighting the (1/2)f_{RF}-cycle delay.*

Fig. 20. *Measured power spectral densities resulting from PPM pulses modulated by random data with and without DB-BPSK scrambling enabled.*

Fig. 22. *Die photo of the all-digital transmitter.*

that if scrambling is not used, the PPM spectrum exceeds the −41.3 dBm/MHz FCC mask and the transmitted energy/pulse would have to be reduced by 10 dB in order to be FCC compliant.

The three-channel pulse spectrum was measured, and the results are superimposed in Fig. 21 along with the FCC mask. DB-BPSK scrambling was enabled for these three measurements.

A die photo of the transmitter is shown in Fig. 22. The 0.8×0.8 mm^2 chip is pad limited, and the active area is 0.2×0.4 mm^2. Most of the active area is consumed by the delay line and control logic that configures the chip.

E. Extending Communication Distance

1) Physical Limits: Communication distance in a UWB system is maximized when the SNR seen at the receiver is maximized. This occurs when the transmitter generates maximum total output power under regulatory limits. At low PRFs, prohibitively large amplitude transmitted pulses are required to maximize power under FCC spectral

masks. For example, a peak-to-peak voltage swing of 19.3 V is required to maximally satisfy FCC spectral masks at a PRF of 10 kHz. Not only would this likely violate peak power limits but this is also impractical in deep submicrometer CMOS where supply voltages are on the order of 1 V.

An alternative approach at low data rates to maximize total power under FCC masks is to reduce output voltage swings and increase the PRF (i.e., multiple pulses transmitted per bit). In addition, many pulses can be transmitted immediately after one another to form a burst of pulses. We define the burst repetition frequency (BRF) as the frequency of pulse bursts, which is always less than or equal to the PRF. The IEEE 802.15.4a standard incorporates pulse bursting into the UWB PHY [27]. Illustrative transient waveforms are shown in Fig. 23.

Bursting multiple pulses back-to-back can allow for improved receiver and transmitter performance. For the receiver presented in Section VI, each integration period
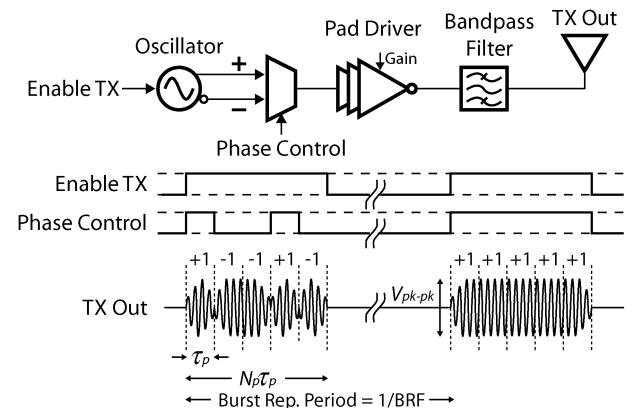

Fig. 21. *Measured spectra for the three channels.*

Fig. 23. *Transmitter architecture incorporating pulse bursting techniques. N_p is the number of pulses per burst and $PRF_{avg} = N_p BRF$.*

can accommodate multiple pulses. With an integration window of 30 ns, the proposed receiver is capable of integrating up to 15 2-ns pulses without suffering any energy/bit overhead. Integrating multiple pulses at once allows for reduced energy compared to integrating individual pulses separately with an equally long integration period. For highly duty cycled transceivers where bias currents are enabled only during transmission and reception of pulses, increasing the number of pulses per burst decreases the overall turn-on and turn-off energy cost for a given PRF.

2) Regulatory Limits: FCC regulations limit the efficacy of pulse bursting and must be considered in system design. The FCC limits the output power in the 3.1-to-10.6 GHz UWB band in two ways [1].

1) The *average* power spectral density (PSD) must be less than or equal to −41.3 dBm/MHz. This corresponds to a theoretical maximum total power of −13.9 dBm for a 550 MHz bandwidth signal. In practice, this number is reduced by 2–4 dB due to nonideal pulse generation.

2) The *peak* power may not exceed 0 dBm at the UWB signal's center frequency in a 50 MHz resolution bandwidth (RBW). Since most spectrum analyzers are not equipped with a 50 MHz intermediate-frequency (IF) filter, the peak power measurement is typically performed at a lower RBW and the limit conservatively set to be $P_{pk} \leq 20log_{10}(\text{RBW}/50 \text{ MHz})$.

High-data-rate IR-UWB transmitters are typically average power limited, while low-data-rate transmitters are typically peak power limited [7]. Low-data-rate transmitters can thus forsake a considerable amount of total power in order to remain peak power compliant. In other words, the peak-to-average power ratio (PAPR) of low-pulse-rate transmitters is generally large and thus violates the FCC peak power limit well before the average power limit. In this paper, PAPR refers to the ratio of peak power in a given RBW to the average power. This definition of PAPR differs from the standard definition of PAPR, which is equal to the peak amplitude of a waveform divided by the root mean square value of the waveform.

The average power of low-BRF transmitters has a $10log()$ dependence on BRF, yet the peak power does not depend on BRF for a fixed number of pulses per burst [48], [49]. As a result, the PAPR can be reduced by decreasing the peak-to-peak voltage swing and increasing the BRF while maintaining constant total power. Since in this case bits comprise several burst repetitions, the number of receiver observations required per bit is increased as a tradeoff for enhanced communication distance. More sophisticated coding techniques may be applied in lieu of simple burst repetitions to further increase communication distance at the cost of increased hardware complexity.

An additional method to reduce PAPR is to prohibit long bursts of pulses with no phase inversions. In coherent systems, this can be achieved by data coding such that long strings of identical bits are reduced and/or eliminated. Noncoherent systems, which often use a pseudorandom binary sequence (PRBS) as the phase scrambling data, can insert run length limits (RLLs) to actively eliminate long strings of identical bits.

For example, a linear feedback shift register (LFSR)-based scrambler may output a run of 8 '+1' phase bits at some point in its sequence, which results in a considerable amount of peak power at the carrier frequency. An RLL of three would ensure that in this situation, the fourth and eighth phase bits are inverted, as illustrated in Fig. 24. This technique spreads peak power away from the carrier frequency while leaving the average PSD undisturbed for reasonable run length limits. However, small-valued RLLs (such as run length limits of three when a burst contains 15 pulses) can distort the average PSD by spreading a noticeable amount of average power away from the carrier frequency.

Fig. 25 illustrates the reduction of peak power in noncoherent systems by decreasing voltage swings,


Fig. 24. *Example of run-length limiting. The simulated peak PSDs emulate the results of a spectrum analyzer operating in peak-hold mode.*

Fig. 25. *Peak power versus V_{pk-pk} with $P_{avg} = -41.3$ dBm/MHz fixed. Here, $R_{nom} = 10$ kb/s, RBW = 8 MHz, $\tau_p = 2$ ns, $N_p = 16$, and RLL = 4 (when used). Measured results are from a custom pulse generator [50].*

increasing BRF, and applying run-length limiting. The maximum average PSD is fixed at -41.3 dBm/MHz across all voltages and BRFs. In the case for a data rate of 10 kb/s, the communication distance is maximized under peak and average power constraints without run-length limiting at $V_{pk-pk} = 0.26$ V and BRF = 3.3 MHz. Sixteen pulses are generated per burst, and 330 burst-repetitions are required per bit. Applying an RLL of four reduces peak power by 3 dB, giving $V_{pk-pk} = 0.37$ V, BRF = 1.7 MHz, and 170 burst-repetitions per bit. This represents a $1.9\times$ improvement in receiver energy efficiency with the relatively small hardware overhead cost of a run-length limiting block at the output of an LFSR.

VI. LOW-RATE IR-UWB RECEIVER

A. Architecture

As described in Section IV, noncoherent modulation is employed to reduce energy consumption and complexity of the overall system. There have been several published noncoherent IR-UWB receivers. In [51], PPM modulated pulses are mixed to baseband and demodulated through quadrature analog correlation. Incoming pulses are correlated with windowed sines in the analog domain at baseband. A second, popular receiver architecture is an energy detection architecture, whereby the received signal is squared with itself and integrated to obtain a representation of the received energy [28], [33], [52], [53]. In contrast to quadrature analog correlation, no windowed sines need to be generated at baseband; however, a significant amount of gain is required before the squaring operation to overcome its inherent nonlinear attenuation of small inputs. This squaring can be done either at RF [52] or at baseband [41]. The focus of this section is on the noncoherent UWB receiver presented in [28].

Fig. 26 shows a simplified block diagram of the noncoherent receiver. It is composed of a 3-to-5 GHz subbanded RF front-end, a passive self-mixer, and a low-power relative-compare baseband [28]. By performing channel selection at RF, no RF PLL is required. Only a 33-MHz crystal is needed to operate the relative-compare baseband. The receiver measures the amount of energy received in two consecutive 30 ns integration windows and compares them in the analog domain to demodulate PPM symbols. Between PPM symbols, the receiver can be rapidly duty cycled to support scalable data rates up to 16.7 Mb/s. The entire receiver can operate at 0.65 or 0.5 V and is implemented in a 90 nm CMOS process.

B. RF/Analog Circuit Design

The RF front-end consists of a low-noise amplifier (LNA), six RF gain stages, and a passive self-mixer (Fig. 26). Each of the six amplifiers contain a second-order resonant bandpass filter for channel selection. The front-end provides 40 dB of gain at RF to overcome the $V_{out} = k * V_{in}^2$ transfer characteristic of the self-mixer for small inputs. The circuit schematic for the LNA is shown in Fig. 27(a). This block operates at 0.65 or 0.5 V, depending on the required sensitivity. This circuit is based upon a common-gate common-source single-to-differential conversion amplifier (with core transistors M_2 and M_3) to generate a differential


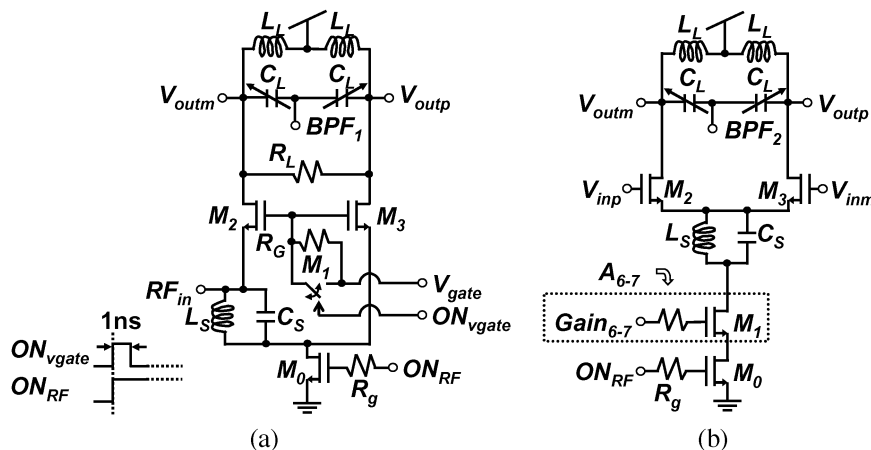**Fig. 26.** *Block diagram of the self-mixing receiver.*

Fig. 27. *(a) LNA and (b) A_{2-7}.*

output [54]. For efficient duty cycling, the LNA is dynamically biased and switched on within 2 ns.

Fig. 27(b) shows the schematic for A_{2-7}. Each of these resonant amplifiers acquires biasing from the preceding amplifier's dc output voltage. Each stage provides 6 dB of gain and can be tuned to any one of three subband channels at 3.4, 3.9, and 4.4 GHz. The final two gain stages A_{6-7} have power scalable, bandwidth-independent gain control with the inclusion of M_1. Because the input gate voltage of each amplifier is dc biased to the power supply through the load inductor of the prior stage, large capacitances are not charged/discharged during power on/off. Only the parasitic capacitance at the drain of M_0 is charged/discharged.

The passive self-mixer implements a squaring operation through the use of transistors biased in their triode region acting as voltage-controlled resistors. The biasing

for the source and drain of these transistors is derived from the dc output voltage of the previous amplifier stage. The maximum gain through the mixer occurs when the gate-to-source voltages of the transistors are biased close to the threshold voltage. This passive self-mixer requires zero voltage headroom for operation. The single-ended output of the mixer lends itself to the baseband circuitry that follows.

C. Mixed-Signal Baseband Circuit Design

To achieve rapid coarse acquisition during the preamble search, uninterrupted time-adjacent integration windows for continuous bit decisions are needed [55]. For rapid synchronization, a receiver having multiple overlapping integration windows is ideal [45], but in this work a single integrator is employed. Fig. 28 shows the implemented baseband block diagram and the corresponding operation schedule. Four capacitors are rotated

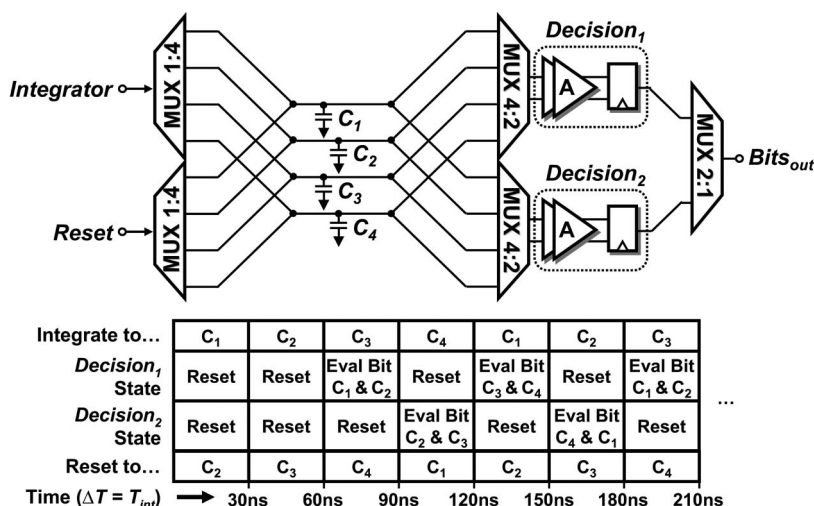| Integrate to…	C_1	C_2	C_3	C_4	C_1	C_2	C_3	
Decision₁ State	Reset	Reset	Eval Bit C_1 & C_2	Reset	Eval Bit C_3 & C_4	Reset	Eval Bit C_1 & C_2	…
Decision₂ State	Reset	Reset	Reset	Eval Bit C_2 & C_3	Reset	Eval Bit C_4 & C_1	Reset	
Reset to…	C_2	C_3	C_4	C_1	C_2	C_3	C_4	
Time ($\Delta T = T_{int}$) →	30ns	60ns	90ns	120ns	150ns	180ns	210ns	

**Fig. 28.** *Detailed diagram of relative-compare baseband.*

Fig. 29. *Integrator and S/H capacitors.*

among three states: reset, integrate, and evaluate to achieve integrations every T_{int}. Two offset-compensated relative-compare paths (Decision$_1$ and Decision$_2$) swap between evaluate and reset modes to provide a continuous stream of bit decisions every $T_{int} = 30$ ns. The pipeline delay for a bit decision to appear at the output from the end of a PPM symbol integration time is 45 ns. The fastest clock required to operate the entire receiver is $1/T_{int}$, or 33.3 MHz.

At baseband, the integrator, sample-and-hold (S/H) capacitor bank (C_1, C_2, C_3, and C_4), offset-compensated preamplifier stages, and latch are all designed for 0.5 V operation and work together to perform signal demodulation. The integrator, shown in Fig. 29, is an inverter externally biased at its switching threshold for 38 dB of dc gain. This is a similar approach to [33], where a dynamic inverter was used as an integrator. If additional gain is required, positive feedback can be employed [51].

The integrator is connected to one of four capacitors that are rotated through at each integration time according to the operation schedule in Fig. 28. Although the integrator integrates in increments of 30 ns, the capacitors hold the final integration value for an additional 60 ns so that adjacent bit decisions can be made. For each

capacitor, the 30 ns integration time, two consecutive 30 ns hold times, and 30 ns reset time are staggered relative to each other. V_{eval1} and V_{eval2} connect the S/H capacitors to the appropriate offset-compensated pre-amplifiers for downstream bit evaluation. The total gain for the cascaded preamplifiers stages is 20 dB, which reduces the latch offset voltage by an order of magnitude. The latch generates a single bit output per relative compare. Although this is sufficient for PPM data demodulation, additional bits of information would assist in automatic gain control and rapid synchronization.

D. Measurement Results

The receiver is wire bonded in a QFN 28-pin leadless package and mounted on an FR4 printed circuit board (PCB). A commercial FPGA to USB2.0 interface board provides digital I/O control to the chip and is where a back-end timing acquisition algorithm is implemented [55]. The RF front-end provides up to 40 dB of gain and consumes more than 99% of the total receiver power. In the 3.4, 3.9, and 4.4 GHz bands, the -3 dB bandwidth varies from 430 to 715 MHz. The filters roll off rapidly and offer channel selection for out-of-band noise and interference suppression. In the 4.4 GHz band, the noise figure (NF) is measured to be 8.6 dB. Without external matching networks, the packaged chip exhibits -10 dB of matching across the entire 3-to-5 GHz channel, regardless of sub-band configuration.

The receiver achieves a sensitivity of -99 dBm at a bit error rate (BER) of 10^{-3} at 100 kb/s in the 4.4 GHz band [Fig. 30(a)], corresponding to a input pulse amplitude of approximately 250 μV. The measured sensitivity is close to a calculated sensitivity of -101 dBm given a front-end NF of 8.6 dB and assuming the squarer and baseband circuits are ideal and contribute no noise [56]. The sensitivity degrades in the other frequency bands due to reduced RF gain. The BER waterfall curves have a more dramatic


Fig. 30. *(a) Sensitivity versus BER and (b) in-band interferers versus BER at 100 kb/s.*

Fig. 31. *Out-of-band interference test.*

rolloff for this PPM signal than for a raw coherent signal because the gain path contains a squared term, which links signal amplitude to achievable gain.

The impact of both in-band and out-of-band interferers on BER is measured. The receiver is initially configured such that the BER is at 10^{-5} without any interferers. To achieve 10^{-5} BER, the received power is set to -98 dBm in the 4.4 GHz band at a data rate of 100 kb/s, translating to a UWB pulse amplitude of 282 μV. Fig. 30(b) shows how the receiver degrades with an in-band sinusoidal interferer. Fig. 31 shows the tests performed for out-of-band interferers at the known 802.11 frequencies and the corresponding results. In most cases, the receiver can tolerate up to -15 and -20 dBm of out-of-band interferer power, corresponding to 56 and 32 mV sinusoidal amplitudes, respectively.

When operating from a 0.65 V supply, the receiver consumes 2.5 nJ/bit for data rates above 10 kb/s. For rates below 10 kb/s, energy/bit increases as data rate decreases, as the fixed leakage power becomes a more significant portion of the energy consumed per bit. The receiver also operates at 0.5 V, and subnanojoule/bit of operation is achievable; however, sensitivity degrades by 15 dB due to reduced front-end gain at the lower supply voltage. The chip area is 1 by 2.2 mm (Fig. 32). Isolation between circuit blocks is accomplished through thick p+ substrate guard rings.


Fig. 32. *Receiver die photo.*

VII. LOW-RATE IR-UWB SYSTEM RESULTS

A. Overview

A unidirectional wireless link is constructed to demonstrate system-level functionality of the transmitter presented in Section V and the receiver presented in Section VI. The transmitter and receiver UWB chips are mounted on custom PCBs that connect to commercially available FPGA prototyping boards. Photographs of the transmitter and receiver nodes are shown in Fig. 33. Commercial omnidirectional UWB antennas are used for all wireless experiments.

FPGAs are used to implement the low-frequency digital functions while also providing USB communication with a PC. For the transmitter, the FPGA is used for digital configuration, for 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.

B. Synchronization

The receiver uses a preamble in each packet for acquisition and synchronization of 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. 26) 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 30 ns time intervals until synchronization is declared. This results in a bit decision's 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.

The synchronization code is repeated four times to increase the probability of detecting a packet. After synchronization is achieved, the receiver searches for a

Fig. 33. *Photos of (a) transmitter and (b) receiver nodes.*

start of frame delimiter (SFD) code that indicates the start of the payload data. After the SFD code is received, the receiver demodulates the payload data. If no SFD code is received after a short period of time, a false acquisition is declared and the receiver returns to synchronization mode.

C. Wireless Demonstration

A streaming video application is realized to demonstrate a unidirectional wireless link using the custom chipset. Streaming video is emulated by transmitting a sequence of images from one PC to another over the UWB link and displaying them in real time as they are received. A wireless link is demonstrated at distances up to 10 m.

VIII. CONCLUSION

In this paper, we presented two impulse-based transceivers demonstrating high-data-rate coherent and low-data-rate noncoherent UWB communication, respectively. In both cases, the focus was to minimize the overall energy consumption at both the architecture and the circuit level. The total energy/bit of a UWB radio is dominated by power consumption of the electronics, as opposed to radiated power of the transmitter, and furthermore the receiver power dominates the transmitter. Therefore, efforts to minimize total energy/bit of UWB systems should focus on reducing power in the circuits, even at the expense of additional radiated power. The high-data-rate system applied parallelism to the baseband ADC and the baseband processor of the receiver to achieve a low-power implementation without sacrificing performance. Specifically, parallelism enabled the used of an energy-efficient SAR architecture in the ADC and enabled deep-voltage scaling

down to 0.4 V in the baseband processor while maintaining a data rate of 100-Mb/s.

The low-data-rate system includes an all-digital transmitter architecture and low-voltage RF and analog circuits in the receiver to reduce power consumption. Additionally, no RF local oscillators are required in either the transmitter or the receiver, allowing the chipset to power on in 2 ns for aggressively duty-cycled operation. Pulse bursting techniques were shown as an extension to the chipsets presented as a means for increasing the communication distance or relaxing the specifications of the electronics.

Like all radio systems, the selection of a UWB implementation greatly depends on the target application. What these results show is that IR-UWB radios will excel in energy-constrained applications with low to moderate data rates, such as implantable devices and RFID tags, particularly when the devices are transmit-only. ∎

REFERENCES

[1] FCC, "Revision of part 15 of the commission's rules regarding ultra-wideband transmission systems," FCC 02-48, Feb. 2002.

[2] S. Iida, K. Tanaka, H. Suzuki, N. Yoshikawa, N. Shoji, B. Griffiths, D. Mellor, F. Hayden, I. Butler, and J. Chatwin, "A 3.1 to 5 GHz CMOS DSSS UWB transceiver for WPANs,"

in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2005, pp. 214–215.

[3] R. Blazquez, F. S. Lee, D. D. Wentzloff, P. P. Newaskar, J. D. Powell, and A. P. Chandrakasan, "Digital architecture for an ultra-wideband radio receiver," in *Proc. IEEE Veh. Technol. Conf.*, Oct. 2003, pp. 1303–1307.

[4] L. Yang and G. B. Giannakis, "Ultra-wideband communications: An idea whose time has come," *IEEE Signal Process. Mag.*, pp. 26–54, Nov. 2004.

[5] F. S. Lee, B. P. Ginsburg, J. D. Powell, M. Scharfstein, D. D. Wentzloff, and A. P. Chandrakasan, "A 3.1 to 10.6 GHz 100 Mb/s pulse-based ultra-wideband radio

receiver chipset," in *Proc. IEEE Int. Conf. Ultra-Wideband*, Sep. 2006, pp. 185–190.

[6] D. D. Wentzloff, R. Blazquez, F. S. Lee, B. P. Ginsburg, J. Powell, and A. P. Chandrakasan, "System design considerations for ultra-wideband communication," *IEEE Commun. Mag.*, vol. 43, pp. 114–121, Aug. 2005.

[7] D. D. Wentzloff and A. P. Chandrakasan, "Gaussian pulse generators for subbanded ultra-wideband transmitters," *IEEE Trans. Microwave Theory Tech.*, vol. 54, pp. 1647–1655, Jun. 2006.

[8] F. S. Lee and A. P. Chandrakasan, "A BiCMOS ultra-wideband 3.1–10.6 GHz front-end," *JSSC*, vol. 41, pp. 1784–1791, Aug. 2006.

[9] M. Verhelst and W. Dehaene, "System design of an ultra-low power, low data rate, pulsed UWB receiver in the 0–960 MHz band," in *Proc. IEEE Int. Conf. Commun.*, May 2005, vol. 4, pp. 2812–2817.

[10] R. Blazquez, P. P. Newaskar, F. S. Lee, and A. P. Chandrakasan, "A baseband processor for impulse ultra-wideband communications," *IEEE J. Solid-State Circuits*, vol. 40, pp. 1821–1828, Sep. 2005.

[11] I. D. O'Donnell and R. W. Brodersen, "An ultra-wideband transceiver architecture for low power, low rate, wireless systems," *IEEE Trans. Veh. Technol.*, vol. 54, pp. 1623–1631, Sep. 2005.

[12] C.-H. Yang, K.-H. Chen, and T.-D. Chiueh, "A 1.2 V 6.7 mW impulse-radio UWB baseband transceiver," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2005, vol. 1, pp. 442–608.

[13] B. P. Ginsburg and A. Chandrakasan, "A 500 MS/s 5b ADC in 65 nm CMOS," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2006, pp. 140–141.

[14] V. Sze and A. Chandrakasan, "A 0.4-V UWB baseband processor," in *Proc. IEEE Int. Symp. Low Power Electron. Design*, Aug. 2007, pp. 262–267.

[15] P. P. Newaskar, R. Blazquez, and A. P. Chandrakasan, "A/D precision requirements for an ultra-wideband radio receiver," in *Proc. IEEE Workshop Signal Process. Syst.*, Oct. 2002, pp. 270–275.

[16] W. Black and D. Hodges, "Time interleaved converter arrays," *IEEE J. Solid-State Circuits*, vol. SC-15, pp. 1022–1029, Dec. 1980.

[17] D. Draxelmayr, "A 6b 600 MHz 10 mW ADC array in digital 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2004, pp. 264–265.

[18] B. P. Ginsburg and A. P. Chandrakasan, "Dual time-interleaved successive approximation register ADCs for an ultra-wideband receiver," *IEEE J. Solid-State Circuits*, vol. 42, pp. 247–257, Feb. 2007.

[19] R. Blazquez and A. P. Chandrakasan, "Architectures for energy-aware impulse UWB communications," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Mar. 2005, pp. 18–32.

[20] J. Foerster, "Channel modeling sub-committee report final, IEEE P802.15-02/368r5-SG3a," IEEE P802.15 Working Group for WPAN, Feb. 2002, Tech. Rep.

[21] R. Blazquez, "Ultra-wideband digital baseband," Ph.D. dissertation, Massachusetts Inst. of Technology, Cambridge, 2006.

[22] R. Blazquez, P. Newaskar, and A. Chandrakasan, "Coarse acquisition for ultra-wideband digital receivers," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Apr. 2003, vol. 4, pp. 137–140.

[23] B. Calhoun, A. Wang, and A. Chandrakasan, "Modeling and sizing for minimum energy operation in sub-threshold circuits," *IEEE J. Solid-State Circuits*, vol. 40, pp. 1778–1786, Sep. 2005.

[24] V. Sze, R. Blazquez, M. Bhardwaj, and A. Chandrakasan, "An energy efficient sub-threshold baseband processor architecture for pulsed ultra-wideband communications," in *IEEE Int. Conf. Acoust., Speech Signal Process.*, May 2006, pp. III 908–III 911.

[25] B. P. Ginsburg and A. P. Chandrakasan, "500-MS/s 5-b ADC in 65-nm CMOS with split capacitor array DAC," *IEEE J. Solid-State Circuits*, vol. 42, pp. 739–747, Apr. 2007.

[26] B. P. Ginsburg and A. P. Chandrakasan, "Highly interleaved 5b 250 MS/s ADC with redundant channels in 65 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2008, pp. 240–241.

[27] *IEEE Standard for Information Technology—Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific Requirement Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs)*, IEEE Std. 802.15.4a-2007, 2007, *(Amendment to IEEE Std. 802.15.4-2006)*, pp. 1–203.

[28] F. Lee and A. Chandrakasan, "A 2.5 nJ/b 0.65 V 3-to-5 GHz subbanded UWB receiver in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, pp. 116–590.

[29] G. Aiello, "Challenges for ultra-wideband (UWB) CMOS integration," in *Proc. IEEE Radio Freq. Integr. Circuits Symp.*, 2003, pp. 497–500.

[30] A. Ismail and A. Abidi, "A 3 to 10 GHz LNA using a wideband LC-ladder matching network," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2004, vol. 1, pp. 384–534.

[31] D. Wentzloff and A. Chandrakasan, "A 47 pJ/pulse 3.1-to-5 GHz all-digital UWB transmitter in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, pp. 118–591.

[32] Y. Zheng, K.-W. Wong, M. Annamalai Asaru, D. Shen, W. H. Zhao, Y. J. The, P. Andrew, F. Lin, W. G. Yeoh, and R. Singh, "A 0.18 μm CMOS dual-band UWB transceiver," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, pp. 114–115.

[33] T.-A. Phan, V. Krizhanovskii, and S.-G. Lee, "Low-power CMOS energy detection transceiver for UWB impulse radio system," in *Proc. IEEE Custom Integr. Circuits Conf.*, Sep. 2007, pp. 675–678.

[34] Y. Zheng, Y. Tong, C. W. Ang, Y.-P. Xu, W. G. Yeoh, F. Lina, and R. Singh, "A CMOS carrier-less UWB transceiver for WPAN applications," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2006, pp. 378–387.

[35] V. Kulkarni, M. Muqsith, H. Ishikuro, and T. Kuroda, "A 750 Mb/s 12 pJ/b 6-to-10 GHz digital UWB transmitter," in *Proc. IEEE Custom Integr. Circuits Conf.*, Sep. 2007, pp. 647–650.

[36] T. Norimatsu, R. Fujiwara, M. Kokubo, M. Miyazaki, A. Maeki, Y. Ogata, S. Kobayashi, N. Koshizuka, and K. Sakamura, "A UWB-IR transmitter with digitally controlled pulse generator," *IEEE J. Solid-State Circuits*, vol. 42, pp. 1300–1309, Jun. 2007.

[37] S. Bagga, W. A. Serdijn, and J. R. Long, "A PPM Gaussian monocycle transmitter for ultra-wideband communications," in *Proc. IEEE Conf. Ultra Wideband Syst. Technol.*, May 2004, pp. 130–134.

[38] Y. Jeong, S. Jung, and J. Liu, "A CMOS impulse generator for UWB wireless communication systems," in *Proc. IEEE Int. Symp. Circuits Syst.*, May 2004, vol. 4, pp. 129–132.

[39] K. Marsden, H.-J. Lee, D. Ha, and H.-S. Lee, "Low power CMOS re-programmable pulse generator for UWB systems," in *Proc. IEEE Conf. Ultra Wideband Syst. Technol.*, Nov. 2003, pp. 443–447.

[40] H. Kim, D. Park, and Y. Joo, "All-digital low-power CMOS pulse generator for UWB system," *Electron. Lett.*, vol. 40, no. 24, pp. 1534–1535, Nov. 2004.

[41] Y. Zheng, M. A. Arasu, K.-W. Wong, Y. J. The, A. P. H. Suan, D. D. Tran, W. G. Yeoh, and D.-L. Kwong, "A 0.18 μm CMOS 802.15.4a UWB transceiver for communication and localization," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2008, pp. 118–119.

[42] L. Smani, C. Tinella, D. Hlal, C. Stoecklin, L. Chabert, C. Devaucelle, R. Cattenoz, N. Rinaldi, and D. Belot, "Single-chip CMOS pulse generator for UWB systems," *IEEE J. Solid-State Circuits*, vol. 41, pp. 1551–1561, Jul. 2006.

[43] T. Terada, S. Yoshizumi, M. Muqsith, Y. Sanada, and T. Kuroda, "A CMOS ultra-wideband impulse radio transceiver for 1 Mb/s data communications and ±2.5 cm range findings," *IEEE J. Solid-State Circuits*, vol. 41, pp. 891–898, Apr. 2006.

[44] C. Kim, I.-C. Hwang, and S.-M. Kang, "A low-power small-area ±7.28-ps-jitter 1-GHz DLL-based clock generator," *IEEE J. Solid-State Circuits*, vol. 37, pp. 1414–1420, Nov. 2002.

[45] L. Stoica, A. Rabbachin, and I. Oppermann, "A low-complexity noncoherent IR-UWB transceiver architecture with TOA estimation," *IEEE Trans. Microwave Theory Tech.*, vol. 54, pp. 1637–1646, Jun. 2006.

[46] Y.-P. Nakache and A. Molisch, "Spectral shaping of UWB signals for time-hopping impulse radio," *IEEE J. Sel. Areas Commun.*, vol. 24, pp. 738–744, Apr. 2006.

[47] D. Wentzloff and A. Chandrakasan, "Delay-based BPSK for pulsed-UWB communication," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Apr. 2007, vol. 3, pp. III-561–III-564.

[48] Agilent Technologies, "Spectrum analysis . . .pulsed RF," Palo Alto, CA, Application Note 150-2, May 2005.

[49] P. P. Mercier, D. C. Daly, M. Bhardwaj, D. D. Wentzloff, F. S. Lee, and A. P. Chandrakasan, "Ultra-low-power UWB for sensor network applications," in *Proc. IEEE Int. Symp. Circuits Syst.*, May 2008, pp. 2562–2565.

[50] P. P. Mercier, D. C. Daly, and A. P. Chandrakasan, "A 19 pJ/pulse UWB transmitter with dual capacitively-coupled digital power amplifiers," in *Proc. IEEE Radio Freq. Integr. Circuits Symp.*, Jun. 2008.

[51] J. Ryckaert, M. Badaroglu, V. De Heyn, G. Van der Plas, P. Nuzzo, A. Baschirotto, S. D'Amico, C. Desset, H. Suys, M. Libois, B. Van Poucke, P. Wambacq, and B. Gyselinckx, "A 16 mA UWB 3-to-5 GHz 20 Mpulses/s quadrature analog correlation receiver in 0.18 μm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2006, pp. 368–377.

[52] Y. Zheng, Y. Tong, J. Yan, Y.-P. Xu, W. G. Yeoh, and F. Lin, "A low power

noncoherent CMOS UWB transceiver ICs," in *Proc. IEEE Radio Freq. Integr. Circuits Symp.*, Jun. 2005, pp. 347–350.

[53] L. Stoica, S. Tiuraniemi, I. Oppermann, and H. Repo, "An ultra wideband low complexity circuit transceiver architecture for sensor networks," in *Proc. IEEE Int. Symp. Circuits Syst.*, May 2005, vol. 1, pp. 364–367.

[54] R. Bagheri, A. Mirzaei, S. Chehrazi, M. Heidari, M. Lee, M. Mikhemar, W. Tang, and A. Abidi, "An 800 MHz to 5 GHz software-defined radio receiver in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2006, pp. 1932–1941.

[55] D. Wentzloff, F. Lee, D. Daly, M. Bhardwaj, P. Mercier, and A. Chandrakasan,

"Energy efficient pulsed-UWB CMOS circuits and systems," in *Proc. IEEE Int. Conf. Ultra-Wideband*, Sep. 2007, pp. 282–287.

[56] J. G. Proakis, *Digital Communications,* 4th ed. New York: McGraw-Hill, 2001.

ABOUT THE AUTHORS

Anantha P. Chandrakasan (Fellow, IEEE) received the B.S, M.S., and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1989, 1990, and 1994, respectively.

Since 1994, he has been with the Massachusetts Institute of Technology (MIT), Cambridge, where he is currently the Joseph F. and Nancy P. Keithley Professor of Electrical Engineering. He also is Director of the MIT Microsystems Technology Laboratories. His research interests include low-power digital integrated circuit design, wireless microsensors, ultra-wide-band radios, and emerging technologies. He is a coauthor of *Low Power Digital CMOS Design* (Norwell, MA: Kluwer Academic, 1995), *Digital Integrated Circuits,* 2nd ed. (Englewood Cliffs, NJ: Pearson Prentice-Hall, 2003), and *Sub-threshold Design for Ultra-Low Power Systems* (New York: Springer, 2006). He is also a coeditor of *Low Power CMOS Design* (New York: IEEE Press, 1998), *Design of High-Performance Microprocessor Circuits* (New York: IEEE Press, 2000), and *Leakage in Nanometer CMOS Technologies* (New York: Springer, 2005).

Prof. Chandrakasan was a corecipient of several awards including the 1993 IEEE Communications Society's Best Tutorial Paper Award, the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997, the 1999 DAC Design Contest Award, the 2004 DAC/ISSCC Student Design Contest Award, the 2007 ISSCC Beatrice Winner Award for Editorial Excellence, and the 2007 ISSCC Jack Kilby Award for Outstanding Student Paper. He was a Technical Program Cochair for the 1997 International Symposium on Low Power Electronics and Design (ISLPED), VLSI Design '98, and the 1998 IEEE Workshop on Signal Processing Systems. He was the Signal Processing Subcommittee Chair for ISSCC 1999–2001, the Program Vice-Chair for ISSCC 2002, the Program Chair for ISSCC 2003, and the Technology Directions Subcommittee Chair for ISSCC 2004–2008. He was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001. He served on SSCS AdCom from 2000 to 2007 and was the Meetings Committee Chair from 2004 to 2007. He is the Technology Directions Chair for ISSCC 2009.

David D. Wentzloff (Member, IEEE) received the B.S.E. degree in electrical engineering from the University of Michigan, Ann Arbor, in 1999 and the S.M. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, in 2002 and 2007, respectively.

In summer 2004, he was with the Portland Technology Development group, Intel Corporation, Hillsboro, OR. Since 2007, he has been with the University of Michigan, where he is currently an Assistant Professor of electrical engineering and computer science. He is serving on the Technical Program Committee for ICUWB.

Prof. Wentzloff is a member of IEEE and Tau Beta Pi. He received the 2002 MIT Masterworks award and the 2004 Analog Devices Distinguished Scholar award.

Vivienne Sze (Student Member, IEEE) received the B.A.Sc. (Hons.) degree in electrical engineering from the University of Toronto, Toronto, ON, Canada, in 2004 and the S.M. degree from the Massachusetts Institute of Technology, Cambridge, in 2006, where she is currently pursuing the doctoral degree.

Her current research interests include energy-efficient architectures and algorithms for video and image processing.

Ms. Sze received the 2007 DAC/ISSCC Student Design Contest Award. She received the Natural Sciences and Engineering Research Council of Canada (NSERC) Julie Payette fellowship in 2004, the NSERC Postgraduate Scholarships in 2005 and 2007, and the Texas Instruments Graduate Woman's Fellowship for Leadership in Microelectronics in 2008.

Fred S. Lee (Member, IEEE) received the B.S., M.Eng., and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 2002 and 2007, respectively.

He is currently with SiTime, Sunnyvale, CA, developing quartz-less timing references and RF chips based on MEMS resonators. From 2007 to 2008, he was with Rambus, Inc., Los Altos, CA, focusing on mixed-signal multigigahertz wireline and 60 GHz wireless transceiver circuits and systems.

While with MIT, Dr. Lee received many awards with his colleagues, including the ISLPED Low Power Design Contest in 2002, the DAC/ISSCC Student Design Contest Award in 2004, and the Jack Kilby Best Student Paper Award at ISSCC 2007.

Brian P. Ginsburg (Member, IEEE) received the S.B. and M.Eng. degrees in electrical engineering and computer science and the Ph.D. degree from the Massachusetts Institute of Technology (MIT), Cambridge, in 2003, 2003, and 2007, respectively.

He is currently a Mixed-Signal Designer with the wireless unit of Texas Instruments, Dallas. His research interests include low-power mixed-signal and RF circuit design, analog-to-digital converters, and system energy optimization.

Mr. Ginsburg was named a Siebel Scholar in 2003, received the NDSEG Fellowship in 2004, and received the 2006 ISLPED International Low Power Design Contest.

Patrick P. Mercier (Student Member, IEEE) received the B.Sc. degree in electrical and computer engineering from the University of Alberta, Edmonton, AB, Canada, in 2006 and the S.M. degree from the Massachusetts Institute of Technology, Cambridge, in 2008, where he is currently pursuing the Ph.D. degree.

From May to August 2008, he was with the Microprocessor Technology Lab, Intel Corporation, Hillsboro, OR, designing spatial encoding and low-swing circuits for multicore communication networks. His research interests include the design of energy-efficient RF circuits and digital systems for wireless communication applications.

Mr. Mercier received a Natural Sciences and Engineering Council of Canada Julie Payette fellowship in 2006.

Denis C. Daly (Student Member, IEEE) received the B.A.Sc. degree in engineering science from the University of Toronto, Toronto, ON, Canada, in 2003 and the S.M. degree from the Massachusetts Institute of Technology, Cambridge, in 2005, where he is currently pursuing the doctoral degree.

His research interests include low-power mixed-signal circuits and wireless transceivers. From May to August 2005, he was with the Wireless Analog Technology Center, Texas Instruments, Dallas, designing transceiver modules. From May to August 2003, he worked on high-speed signaling systems with Intel Laboratories, Hillsboro, OR.

Mr. Daly received a Student Paper Prize at the 2006 RFIC Symposium and third place in the 2006 DAC/ISSCC Student Design Contest (Operational System Category). He received Natural Sciences and Engineering Research Council of Canada postgraduate scholarships in 2003 and 2005.

Raúl Blázquez (Member, IEEE) received the ingeniero de telecomunicacion degree from the Polytechnical University of Madrid, Spain, in 1999 and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 2006.

Since June 2006, he has been a Systems Engineer with Digital Signal Processing Systems Research and Development, Application Specific Products Division, Texas Instruments, Dallas. His current research interests include ultra-low-power digital systems for both wireless communications and biomedical applications.

Dr. Blázquez was a corecipient of the DAC/ISCCC Student Design Contest Award in 2004. He received the Intel Fellowship in 2004 and La Caixa Fellowship in 1999.