

Pulsed Ultra-Wideband Transceivers

Patrick P. Mercier*, Denis C. Daly*, Fred S. Lee, David D. Wentzloff,
and Anantha P. Chandrakasan

Abstract Ultra-wideband (UWB) radios offer tremendous promise in terms of achievable data rates due to the large capacity afforded by their inherently large occupied bandwidth. While achieving ultra-high data rates may have been one of the original intents of UWB radios, pulsed-UWB radios have another potential advantage over their narrowband counterparts: energy. By exploiting the large available bandwidth in conjunction with non-coherent signaling, low-complexity and ultra-energy-efficient transmitters can be designed using all-digital architectures that do not require the use of a PLL. Similarly, energy-detecting receivers can receive pulses with low energy-per-bit at high data rates and can be rapidly duty-cycled to minimize overall power consumption. This chapter outlines the main challenges in UWB design, while discussing several representative receiver and transmitter implementations in detail.

Keywords Ultra-wide band • UWB • IR-UWB • Pulsed radios • Low-power radios

* Author contributed equally.

P.P. Mercier (✉)

University of California San Diego, La Jolla, CA 92093, USA

e-mail: pmercier@ucsd.edu

D.C. Daly

Maxim Integrated, North Chelmsford, MA 01863, USA

e-mail: denis.daly@maximintegrated.com

F.S. Lee

Google[x], Mountain View, CA 94043, USA

D.D. Wentzloff

University of Michigan, Ann Arbor, MI 48109, USA

A.P. Chandrakasan

Massachusetts Institute of Technology, Cambridge, MA 02139, USA

© Springer International Publishing Switzerland 2015

P.P. Mercier, A.P. Chandrakasan (eds.), *Ultra-Low-Power Short-Range Radios*,
Integrated Circuits and Systems, DOI 10.1007/978-3-319-14714-7_8

1 Introduction

1.1 Background

The world-wide use of portable electronics has never been more prevalent than in today's society. To stay competitive, it is increasingly important to design portable electronics to have high performance, small form factors, and long battery life. As a result, much research and development efforts have been spent maximizing the performance and integration of portable electronics in power-constrained environments. However, for applications such as wireless sensor networks, medical monitoring, and asset tracking, the ultimate goal of maximizing performance is superseded by minimizing energy consumption, area, and/or cost [8, 41]. In these types of energy-starved applications, the radio-frequency (RF) circuits typically dominate the overall energy budget. Thus, in order to maximize battery lifetime or minimize the required amount of energy harvesting, new and innovative design techniques are required to reduce the RF circuitry energy burden. This chapter describes ultra-wideband (UWB) circuits and their applicability for low power RF applications, focusing primarily on impulse radio ultra-wideband (IR-UWB).

UWB communication was first demonstrated at the end of the nineteenth century by Marconi with spark gap transmitters, but by the end of the twentieth century, UWB communication was primarily used for only niche military and radar applications; instead, narrowband communication was the dominant wireless communication scheme. However, in 2002 the United States Federal Communications Commission (FCC) issued a First Order and Report permitting the development and operation of UWB systems for communication, measurement, imaging and vehicular radar, which reinvigorated academic and industrial UWB development efforts [19, 20]. The FCC established emission limits of -41.3 dBm/MHz in three different frequency bands: below 960 MHz, 3.1-to-10.6 GHz and 22-to-29 GHz. The 22-to-29 GHz band is intended for vehicular radar systems whereas the other two bands can also be used for communication, measurement and imaging systems.

Naturally, these bands overlap with other FCC band allocations, resulting in direct interference between UWB devices and competing narrowband devices. The FCC only permits this by limiting the *average* emissions of UWB radios to be less than the Part 15 radiation limit for consumer electronic devices (-41.3 dBm/MHz). In other words, UWB devices operate below ambient noise levels. This brings up two important questions. First, if the average signal power is below ambient noise, how does a receiver correctly receive and demodulate UWB communications? To answer this, keep in mind that the FCC regulates the *average* radiated output power; since an IR-UWB signal outputs narrow pulses followed by periods of zero radiated output power, the *instantaneous*, or *peak* radiated output power can be large—often well above the noise floor.¹ In other words, UWB signals are,

¹The FCC also limits the peak power of a UWB signal to be less than 0 dBm in a 50 MHz resolution bandwidth, though this is often well above the noise floor.

instantaneously, not necessarily completely buried in noise (though to be fair, signal-to-noise ratios (SNRs) can often be low). The second question is: why bother using UWB in the first place? The answer to this question lies with Shannon: a channel's theoretical capacity (assuming additive white Gaussian noise (AWGN)) is equal to $C = B \times \log_2(1 + SNR)$, and thus a larger bandwidth will give much larger capacity. UWB circuit and system designers can thus leverage this additional capacity in order to achieve ultra-high data rates; alternatively, they can trade-off the available capacity to employ modulation schemes that may be spectrally inefficient, yet result in circuits or architectures that reduce power dramatically. In addition, the fine time resolution enabled by short pulses can enable accurate measurement of distances between objects (i.e., ranging).

According to the FCC, a signal is considered ultra-wideband if it has a -10 dB bandwidth that exceeds the lesser of 20% of its own center frequency or 500 MHz. Ultra-wideband is commissioned to be an overlay technology, such that it does not disrupt the operation of narrowband devices operating in the same frequency span. Due to concerns about interference with low SNR devices such as the global positioning system (GPS), the average power spectral density (PSD) limit is further reduced in other frequency ranges, as shown in Fig. 1.

The resulting power limit constrains high data rate communication to a range of approximately 1-to-10 m, which is appropriate for wireless personal area network (WPAN) and body-area network (BAN) applications. It is possible, however, to trade-off data rate and/or spectral efficiency for increased transmit distance and/or energy efficiency, which opens up UWB to other low power applications. For instance, applications such as miniaturized flying vehicles require communication distances upwards of 100 m, while minimizing both energy consumption and weight due to limited payload carrying capacities [6]. As discussed throughout this chapter, leveraging the wide available bandwidth of UWB signaling can lead to the possibility of achieving small, energy efficient radios.

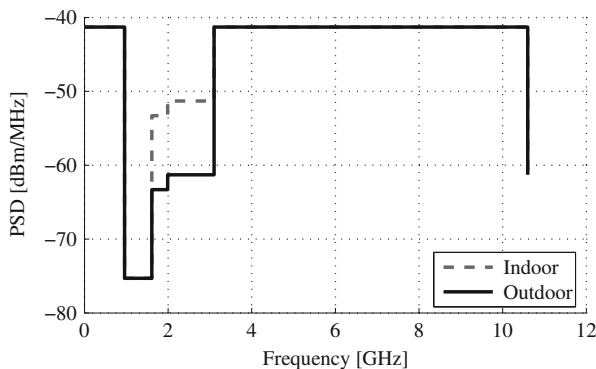


Fig. 1 FCC mask restricting power spectral densities from 0-to-10.6 GHz

1.2 UWB Standards, Proposals, and Communication Schemes

Since the 2002 FCC report did not restrict UWB signaling to any particular scheme, circuit and systems designers have the freedom to choose any type of implementation, provided the spectral masks are met. As a result, several very different techniques were proposed for standardization. One of the early standardization efforts was initiated by the Institute for Electrical and Electronics Engineers (IEEE) 802.15.3a task group, which attempted to add a UWB physical (PHY) layer to the 802.15.3 high-rate WPAN standard. After much deliberation, the task group consolidated the many submitted proposals into two separate proposals: one relying on orthogonal frequency-division multiplexing (OFDM), and the other relying on a form of IR-UWB communication called direct sequence UWB (DS-UWB). Unfortunately the parties could not agree to further consolidation and the 802.15.3a task group was disbanded in 2006 and each technology sought standardization and development elsewhere. The OFDM proposal was adopted by the WiMedia Alliance for high-rate communication and eventually standardized as ECMA-368 and ECMA-369 in 2008. Several industry products compatible with this standard were released but ultimately little commercial success was achieved. The DS-UWB proposal was adapted for low-rate communication as part of the IEEE 802.15.4a amendment. 802.15.4a compliant UWB radios have been released by industry, but long term commercial success of the standard is too early to be determined.

Despite the limited commercial success to-date of products compliant with the EMCA and IEEE UWB standards and amendments, there has been ongoing research and development into using UWB technology in applications that require either energy efficiency, high data rates, localization, difficult-to-intercept communication, or some combination therein. For example, UWB systems have found utility in applications ranging from automotive radar [42], respiration monitoring [63], medical implants, RFID tags [1], and secure military communications. Given the relatively low output power limits, UWB technology appears most differentiated in short range links where either precision ranging is required or unique features of UWB signaling can be exploited in the circuit domain to reduce power consumption, cost, or area. The majority of these applications employ impulse-based communication (impulse radio ultra-wideband (IR-UWB)).

1.3 IR-UWB

A promising approach for implementing low-power UWB communication involves a time-domain IR-UWB approach [56]. With this technique, pulses of very short duration (on the order of 200 ps to 2 ns) are used to create inherently wideband signals capable of both transmitting digital data and providing ranging and localization information [26]. These wideband signals can be generated to lie directly in the

band of interest, or can be generated at baseband and subsequently mixed-up to RF frequencies. Since the radiated pulse power is relatively low due to FCC regulations, IR-UWB receivers must operate at very low signal-to-noise ratios. Correlation and comparison operations are typically required to separate signal information from noise, even at low-to-medium transmission distances.

Compared to narrowband signals, IR-UWB signals are more amenable to be processed in the time domain rather than the frequency domain, which allows for different transceiver architectures with the potential for reduction in cost, area, or power. Due to the wide bandwidth of UWB signals, they can be efficiently amplified and processed with wide-bandwidth, low Q resonant or non-resonant circuits, which can be easily integrated on-chip with minimal area [55]. IR-UWB signaling is highly compatible with digital architectures, and very simple digital pulse transmitters consisting of only digital logic and delay elements have been successfully demonstrated [54].

This chapter describes several IR-UWB transceiver implementations in detail while also highlighting other implementations to provide an overview of the current state-of-the art. Section 2 focuses on IR-UWB receiver implementations, describing a 3-to-5 GHz noncoherent receiver for insect motion control applications and a 9.8 GHz noncoherent receiver for ultra-low power cubic-mm sensor nodes. Section 3 focuses on IR-UWB transmitter implementations, beginning with an overview and classification of architectures, followed by detailed descriptions of an all-digital 3-to-5 GHz transmitter with pulse shaping [29, 31].

1.4 Coherency

Before diving into architectural details, it is first necessary to make an important note about modulation schemes. There are two fundamentally different ways to demodulate data in carrier-based communication systems: coherent versus non-coherent demodulation. Coherent receivers typically lock the incoming carrier phase with a locally generated (and very accurate) carrier or pilot tone, whereas non-coherent receivers discard phase information. For example, consider the non-coherent receiver shown in Fig. 2. The incoming signal is amplified, squared, then integrated over a set window of time. The squaring and integrating operation does

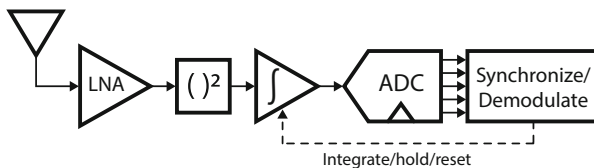


Fig. 2 Block diagram of a non-coherent energy-detecting receiver

not consider phase, and is in fact equivalent to finding the *energy* of a signal in a given window of time. For this reason, this type of non-coherent receiver is called an energy-detecting receiver.

Non-coherent systems have a lower effective data rate for a given bit error rate (BER) compared to the coherent case, since the loss of phase information reduces the number of potential signaling dimensions by one. This limits the types of modulation that can be used, and may result in decreased symbol distances in the constellation diagram.

However, since phase information is discarded, non-coherent systems do not require phase alignment between the transmitter and receiver. Thus, non-coherent receivers are only sensitive to variations in the transmitted *frequency*. If the fractional bandwidth of a system is large (as is the case for UWB), then the *absolute* transmitted frequency accuracy can be relatively low compared to narrowband systems. For example, the IEEE 802.15.4a standard specifies an RF frequency accuracy requirement of ± 20 ppm for coherent signaling, whereas noncoherent UWB signaling can tolerate RF frequency accuracies over $\pm 1,000$ ppm. For this reason, non-coherent systems can employ simple architectures with relaxed frequency requirements, and often do not require the use of phase-locked loops (PLLs) or cordic blocks. Thus, noncoherent signaling is frequently used in systems where minimizing power consumption is the main priority over spectral efficiency or wireless range.

2 IR-UWB Receiver Design

2.1 System and Architecture Level Considerations

Due to the differences between narrowband and UWB signals, UWB receivers frequently are implemented with different architectures and circuits than traditional narrowband receivers. For instance, to achieve ultra-low power operation, it is useful for a UWB receiver to be able to quickly turn on and off multiple times during a packet between individual pulses or bits. This behavior differs compared to narrowband receivers which typically remain on at all times while receiving a packet.

Synchronization is also a key challenge of UWB receiver design because pulses are often transmitted with large gaps in between them, multi-path must be carefully considered, and extremely precise synchronization is required for ranging. Included in 802.15.4a is a packet structure and frame format for the UWB PHY. The frame consists of a synchronization header, a start frame delimiter (SFD), a packet header and a data field. The synchronization header provides time for the receiver to detect a signal, realize automatic gain control (AGC), synchronize with the transmitter, and implement frequency tracking and several other functions. Embedded in the synchronization header are length 31 or length 127 ternary codes

which are repeatedly sent by the transmitter. The UWB PHY specifies forward error correction to be implemented with an outer Reed-Solomon systematic block code and an inner half-rate systematic convolutional code [40]. An interesting characteristic of the UWB PHY is that both coherent and noncoherent signaling are supported. With noncoherent signaling, the receiver can only demodulate the pulse-position modulation (PPM) modulated data and not the binary phase-shift keying (BPSK) modulated data. Thus, the overall data rate is lowered, but simpler, energy-detection receiver architectures are supported.

2.2 Receiver Performance Metrics

To evaluate the performance of receivers, several performance metrics are used including power consumption, data rate and sensitivity. For UWB radios, which frequently operate at fast instantaneous data rates but low duty cycles, it is important to differentiate between peak power consumption and average consumption when duty cycled. A key metric used to compare the energy efficiency of radios is energy per bit, corresponding to the energy required to send or receive a bit of information. Low power radios typically consume less than 5 nJ/bit. The energy/bit metric, while useful, must be evaluated in parallel with receiver sensitivity as well as average and peak power consumption, as generally lower energy/bit values are achieved at higher data rates and worse sensitivity. Table 1 and Fig. 3 present key performance metrics of recently published low power receivers, both narrowband and wideband as well as coherent and noncoherent.

2.3 Design Example: Implementation of a 3-to-5 GHz IR-UWB Receiver

Based on Fig. 3 one can see that a clear trade-off exists between receiver sensitivity and energy per bit. Daly et al. [14] and Mercier et al. [33] achieves a good balance between receiver sensitivity and energy per bit and its implementation is described in detail in this section. The IR-UWB receiver is designed for insect flight control system, with the goal to be able to wirelessly receive commands that control the flight direction of an insect. This system has extremely stringent weight, volume and power consumption requirements, due to the limited carrying capacity of insects. These requirements are similar to distributed sensor network applications.

Figure 4 shows a block diagram of the wireless receiver. The receiver is a noncoherent, energy detection based IR-UWB receiver designed for the 802.15.4a wireless standard. The receiver operates at a peak data rate of 16 Mbps in the

Table 1 Key performance metrics of recently published low power receivers

Author	Data rate (kbps)	Power (mW)	E/bit (nJ/bit)	Sens. at data rate (dBm)	Sens. scaled to 100 kbps (dBm)
Porret [39]	24	1	41.6	-95	-89
Choi [12]	200	21	105	-82	-85
Emira [18]	11,000	114	10.3	-86	-106
Otis [36]	5	0.4	80	-101	-88
Darabi [15]	11,000	360	32.7	-88	-108
Chen [10]	500	2.8	5.6	-80	-87
Lee [25]	16,700	42	2.5	-77	-99
Marholev [28]	3,000	43	14.3	-83	-98
Pletcher [38]	100	0.052	0.5	-72	-72
Zheng [61]	15,600	102	6.51	-75	-97
Weber [50]	2,000	36	17.8	-90	-103
Bohorquez [5]	120	0.4	3.3	-93	-94
Retz [43]	250	30.25	121	-96	-100
Verhelst [49]	20,000	3.1	0.159	-65	-88
Daly [14, 33]	16,000	22.5	1.4	-76	-98
		11	0.7	-50	-72
Brown [7]	30	0.037	1.2	-67	-62

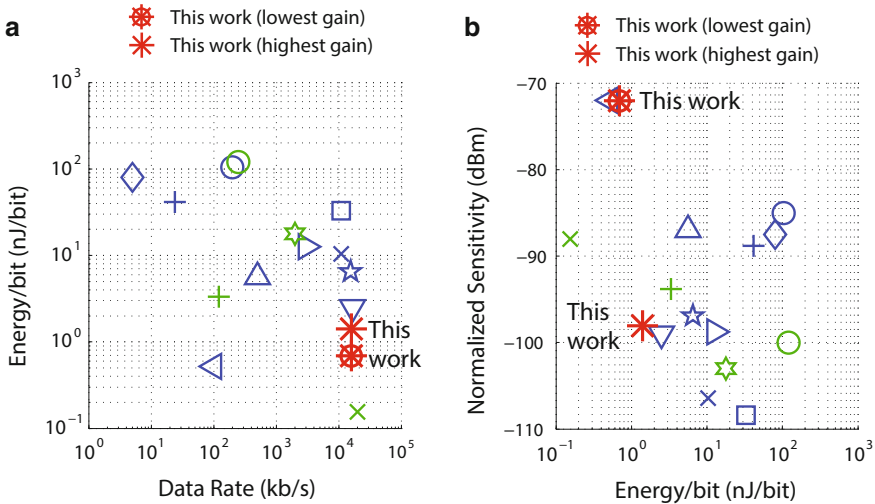


Fig. 3 Two comparison plots of receiver with previously published work: (a) energy/bit versus data rate, and (b) normalized sensitivity versus energy/bit. In both plots, a point is shown for the receiver at its highest and its lowest gain setting. Data for these plots are found in Table 2

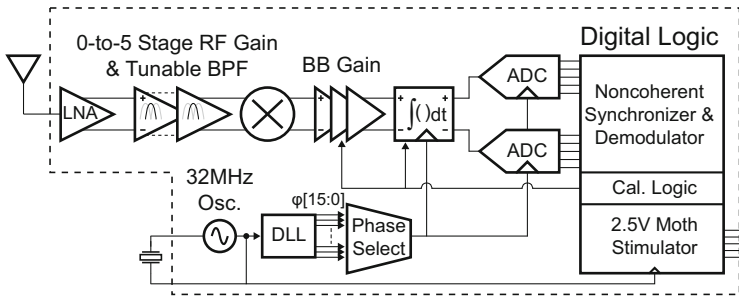


Fig. 4 Detailed block diagram of receiver SoC

3-to-5 GHz UWB band, communicating in one of three 500 MHz channels at 3.5, 4.0, and 4.5 GHz. Through duty cycling, the receiver can operate at lower data rates, thereby reducing average power consumption.

Non-coherent signaling is employed to reduce power consumption on the receiver as it allows for a simple, energy detection architecture without any high frequency clocks. The receiver mixes the received signal with itself at RF, and a windowed integrator and analog-to-digital converter (ADC) at baseband generate a digital signal representing the total energy received in a given time window. This architecture allows for demodulation of both on-off keying (OOK) and PPM signals.

The first stage of the receiver signal chain is an RF front end that amplifies the received signal by up to 40 dB while attenuating out-of-band interferers. This amplified RF signal is then squared, resulting in the RF signal being mixed to baseband. Following the squarer is a baseband amplifier, and then the amplified signal is integrated and quantized by an ADC. The ADC values are passed to a digital backend, which performs packet detection, synchronization and decoding. Also included in the receiver system-on-chip (SoC) is a crystal oscillator and a delay-locked loop (DLL). The entire receiver is clocked by a fixed, 32 MHz clock. After synchronization, the appropriate DLL phase is selected and is used by the windowed integrator and ADC. Each of the specific components of the receiver SoC is described in the following subsections.

2.3.1 RF Front End

For noncoherent receivers, significant gain is required prior to the squarer to obtain a sufficient signal swing such that semiconductor device nonlinearity can be exploited in the squaring element. Passive and active squarers require input voltages on the order of milli-Volts whereas low noise amplifier (LNA) input voltages can be on the order of tens of micro-Volts, thus requiring voltage gain of approximately 40 dB. To achieve such large gain, noncoherent receivers typically employ one of two methods: a super-regenerative architecture [48] or a multi-stage linear amplifier [25]. Although a multi-stage linear amplifier requires more power than

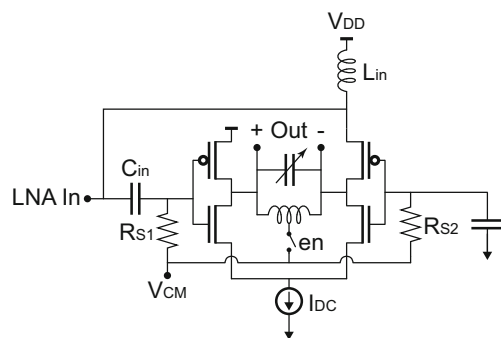
a super-regenerative amplifier, it allows for simple support of any arbitrary squaring and integration interval. Moreover, a multi-stage linear amplifier is less subject to RF leakage out of the antenna, which can potentially result in FCC spectrum violations or require the use of an RF isolation amplifier. Based on these advantages, a multi-stage linear amplifier topology is selected, with a per-stage gain of approximately 8 dB.

A key design choice is whether to implement the multi-stage amplifier with single-ended or differential circuits. As the RF front end is integrated on the same chip as digital logic and baseband analog circuits, a differential architecture offers significant advantages in terms of substrate noise and power supply immunity. In addition, reduced decoupling capacitance is required, and a differential structure allows for higher quality factor inductors and virtual ground ‘center-tap’ nodes. Thus, a differential RF architecture is selected; however, as all commercially available UWB antennas are single ended, and thus the LNA has a single ended input. Single-ended to differential conversion is realized by the LNA and all later stages are differential. Resonant LC loads are used instead of non-resonant loads as they offer superior gain in the 3-to-5 GHz frequency band at the same power consumption and also have a second order bandpass characteristic which rejects out-of-band interferers [25].

The schematic of the LNA is shown in Fig. 5. When the LNA is enabled, the switch en is closed, connecting the dc output of the differential inverters with the dc input of the inverters. Through negative feedback, the dc voltages at all of the nodes normalize to the same value, V_{CM} . To allow the LNA to turn on rapidly, switches are placed in parallel with R_{S1} and R_{S2} and these switches are briefly enabled while the LNA turns on. In normal operation, R_{S1} and R_{S2} are sufficiently large that the negative feedback does not degrade gain. When the LNA is disabled, the switch en is opened, I_{DC} is set to 0 A, and V_{CM} is actively driven to V_{DD} . This allows the output dc voltage to freely float, which is necessary for proper calibration of the receiver.

Following the LNA are five stages of RF gain. Figure 6 presents the schematic of the multi-stage RF amplifier, including the LNA. To dc bias the RF gain stages, the center tap of each stage’s inductor is connected to the center taps of adjacent stages’ inductors. Due to the differential voltage across each inductor, these center tap

Fig. 5 Schematic of low noise amplifier



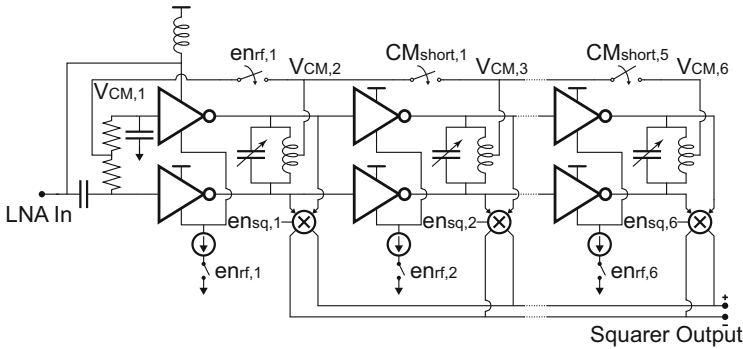


Fig. 6 Schematic of 6-stage RF amplifier, including the LNA. A variable number of stages can be enabled depending on the gain required

nodes are virtual grounds. Moreover, as all RF amplifiers are biased with the same current density, these nodes are nominally at the same dc voltage. By connecting these nodes together with a low impedance connection, the common-mode rejection ratio (CMRR) is superior to what is achieved with more traditional common-mode feedback (CMFB) techniques like resistive feedback. The Monte Carlo simulated common-mode gain of the five stages of RF gain after the LNA has a mean of 7.7 dB and a standard deviation of 7.5 dB at the RF resonant frequency and a mean of 2 dB and a standard deviation of 0.1 dB at low frequencies.

Each gain stage has a squarer at its output, although at any time only one squarer is enabled. Depending on how much RF gain is needed, a variable number of RF gain stages are enabled, as well as the appropriate squarer.

2.3.2 Squarer

A squarer serves two functions in the receiver: to frequency shift (or mix) the received RF signal to baseband and to square its amplitude. It is possible to design an entirely passive squarer that consumes no dc bias current; however, these passive squaring circuits are traditionally single ended [25] or pseudo-differential [37]. In this work, a passive, differential squarer is employed that uses transistors biased in the triode region (Fig. 7). The differential squarer is made possible by the inverter-based RF amplifier, as the output voltage of the RF amplifier is nominally mid-range, thereby allowing both NMOS and PMOS devices to have sufficient gate overdrive. The squarer consumes no static bias currents or active power and has near zero dc output voltage offsets. A key advantage of this structure is that fairly well matched differential outputs are generated. Due to its nonlinear transfer function, the squarer requires RF inputs with amplitudes above approximately 10 mV. At a 10 mV RF input, the single-ended output voltage amplitude is ~ 0.7 mV.

Fig. 7 Schematic of differential, passive squarer

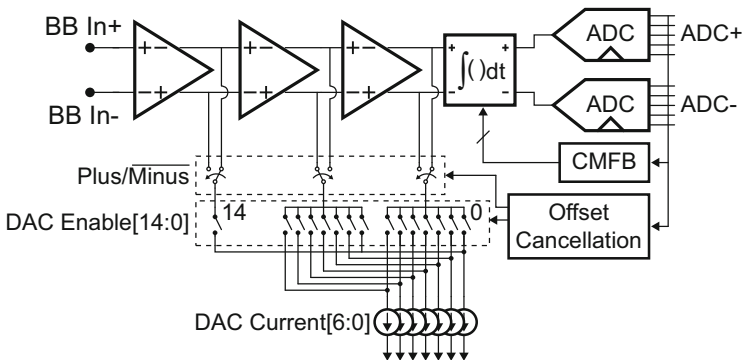
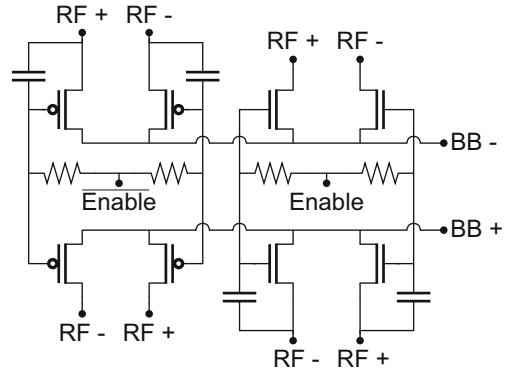


Fig. 8 Baseband signal chain, consisting of a baseband amplifier, an integrator, an ADC, a current-mode DAC that is used to cancel baseband amplifier offsets, and digital calibration logic

2.3.3 Baseband Amplifier

Following the squarer is a baseband signal chain consisting of a three-stage amplifier followed by an integrator and ADC (Fig. 8). The baseband amplifiers are simple differential pairs with resistive loads. The cumulative differential gain of the baseband amplifier chain is simulated to be 83 V/V and the 3 dB bandwidth is 230 MHz. The large baseband gain is required to amplify the squarer output from amplitudes as low as 0.5 mV. Each differential pair operates off a 1 V supply, is supplied 320 μ A of current, and has resistive and capacitive loads of 2.5 k Ω and 150 fF, respectively. A multi-stage amplifier is used rather than an op amp due to the wide signal bandwidths and because a high Q filter is not required.

Due to the small input levels and high gain, offset compensation is a critical component of the baseband amplifier. An input referred offset of merely 10 mV would saturate the baseband amplifier. Traditionally, the goal of offset compensation is to establish a 0 V differential output voltage given a 0 V differential input voltage; however, in this system a fixed offset at the output needs to be established to

maximize dynamic range. This fixed output offset is required because the baseband signal generated by the squarer is monopolar, meaning that the positive squarer output only increases from its 'zero-input' level and the negative squarer output only decreases. Thus, the positive baseband amplifier output should nominally be biased near the bottom of the amplifier's dynamic range.

Offset compensation is implemented digitally with a current-mode digital-to-analog converter (DAC) in a discrete time process. Rather than a traditional architecture of a binary-weighted DAC connected to the output of the first baseband amplifier stage, the DAC consists of current sources that can connect to any of the three baseband amplifier stages. This allows for fine offset control without requiring very small current sources. To ensure monotonicity as the DAC code increases, the current sources transition from being unconnected, to being connected to the final amplifier stage, to eventually being connected to earlier amplifier stages. Depending on whether a positive or negative offset needs to be cancelled, the current sources can connect to the positive or negative output nodes.

During calibration, the LNA is disabled and the baseband inputs are shorted to the same dc value. Next, the integrator and ADC convert the baseband output to a digital value. The ADC output code is processed by a slope tracking state machine to adjust the DAC until the ADC output code approaches the desired ADC value.

2.3.4 Integrator and ADC

Following the baseband amplifier is an integrator and ADC. Both the integrator and ADC are clocked at 32 MHz, resulting in an integration period of 31.25 ns. The output of the ADC is a digital representation of the total RF energy received within the 31.25 ns integration period. This absolute measurement of energy is preferred to a relative measurement of energy, because it allows for demodulation of both PPM and OOK data.

The ADC consists of two single ended ADCs, operating on the positive and negative integrator outputs and each generating 5 bits of information. The difference between these ADC values generates a 6 bit output code, although if perfect matching is assumed, only 5 bits of useful information is generated. Despite this limitation, the pseudo-differential structure offers improved power supply rejection and common-mode rejection compared to a single ended 5 bit structure, while also allowing for a simpler implementation than a fully differential structure.

Having the integration output quantized to multiple bits is useful for gain control and for accurate timing synchronization. Due to the 5 bits of ADC information combined with coding on the transmitter, the receiver is able to synchronize with an accuracy of ± 1 ns while being clocked with a period of 31.25 ns [32].

The integrator and ADC are jointly designed to not require any high frequency clocks, as well as to allow for a simple integrator that does not need op amps, loads with high output impedance, or positive feedback. A detailed block diagram of the integrator and ADC are shown in Fig. 9. Together, the integrator and ADC are similar to a single-slope integrating ADC, but with some key differences.

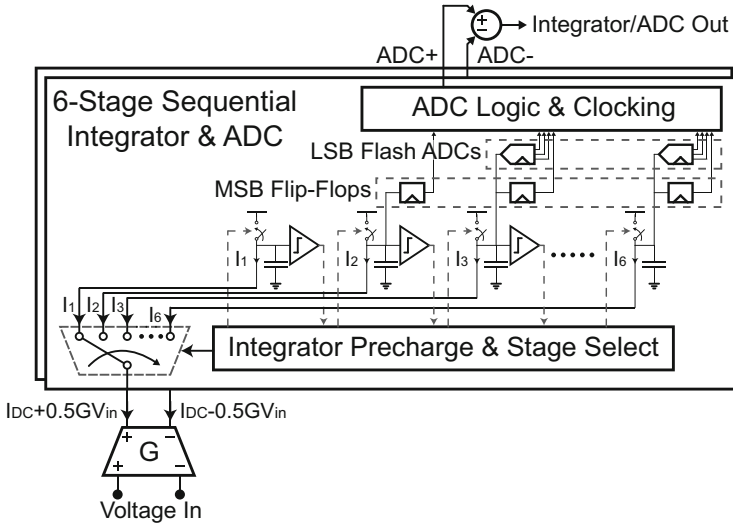


Fig. 9 Six stage sequential integrator and ADC

The differential inputs are first passed through a differential transconductor to convert the input voltage to a current. This current discharges up to six stages from V_{DD} in succession, similar to that of a dynamic inverter. The differential rate of discharge between the positive and negative ADCs is based on the differential input voltage, and thus an integration function is realized. Based on the number of stages that are discharged in the integration period, 2 bits of coarse quantization are generated. Only 2 bits of information are generated from the six stages because the first two stages are not considered in the coarse quantization. The first two stages should ideally always be discharged by the end of an integration period and thus do not contribute information. These first two stages serve to cancel out the static, zero-input dc current of the differential transconductor that is required to appropriately bias the transconductor in a linear region. Additionally, the time while these first two stages are being discharged is leveraged by the final four stages to evaluate the previous integration value.

The ADC generates an additional 3 bits of fine quantization that are combined with the 2 bits of coarse quantization. These 3 bits are generated by quantizing the capacitor voltage of the stage that was being discharged at the end of the integration period with a flash ADC. The capacitor voltages on stages three through six are temporarily held constant while the appropriate flash ADC resolves. During this time period, the next integration period has already begun by discharging stage one. A simple flash ADC with a resistive ladder DAC is used to generate these 3 bits. Thus, 5 bits of data are generated by the integrator and ADC. Both positive and negative outputs of the transconductor are independently processed by this integrator and ADC structure, and thus a pseudo-differential output is generated. The integrator and ADC architecture would only need slight modifications to allow for the use of a differential ADC.

2.3.5 Clocking

The SoC is designed to be clocked off a fixed 32 MHz oscillator that is always enabled. Due to the noncoherent signaling, clock frequency and timing synchronization accuracy requirements between transmitter and receiver are dramatically reduced. Through the use of a Pierce oscillator stabilized with a quartz crystal, it is possible to achieve frequency accuracies on the order of ± 20 ppm [2], allowing the transmitter and receiver to require only one synchronization per packet, without any phase tracking during the packet payload of up to 1,600 bits.

For the receiver to successfully decode data, the integrator and ADC must be phase aligned with the received data. This phase alignment is achieved with a digital synchronization algorithm and a DLL. Based on the result of the digital synchronization, an appropriate phase from the DLL is used to clock the integrator and ADC. During synchronization, the DLL is bypassed and the integrator and ADC are provided the same clock phase as the rest of the digital logic. As the DLL is not being used, the DLL can be calibrated during this time by a successive approximation register (SAR) state machine.

The digital baseband achieves synchronization accuracy of ± 1 ns in an integration window of 31.25 ns, and the DLL is designed to match these specifications. The DLL has 16 outputs, each nominally spaced 1.95 ns apart from one another. Due to the noncoherent signaling, the DLL does not need to have good linearity, and thus it is possible to use very simple delay elements and simple calibration logic. The core delay element consists of a current starved inverter, and a simple DAC is used to control the bias current of the inverter. All outputs of the DLL are passed to a digital, synchronous state machine.

As the integrator and ADC operate from a different clock phase than the rest of the digital logic, there is a potential for timing violations or clock offsets at the interface. To address this problem, the ADC outputs are retimed with registers. These retiming registers can be either positive or negative-edge triggered to ensure sufficient setup and hold time.

2.3.6 Digital State Machine and Duty Cycling

Since the receiver peak data rate of 16 Mbps is much larger than the required data rate in the system, the receiver is designed to be duty cycled. Duty cycling is implemented through the use of a programmable digital state machine. Between packets, the radio and modem are disabled and all digital logic is clock gated except for a sleep counter. This low power sleep mode continues until the sleep counter reaches a programmable count value. At this point, the receiver state machine is triggered, and the receiver attempts to receive a packet.

To receive a packet, the digital state machine first enables the RF and analog circuits, which turn on within one clock cycle. Before the receiver modem performs packet detection, the receiver state machine performs calibration of the DLL, baseband amplifier and integrator. This calibration only takes a few microseconds,

and is performed before every packet reception to account for any change in temperature or supply voltage since the last packet reception attempt.

2.3.7 Digital Baseband Synchronizer

Since the transmitter and receiver are not normally phase synchronized, there is no guarantee that the integration window of the receiver front end lines up perfectly with the pulses generated by the transmitter. Thus, the primary purpose of the digital baseband is to perform this synchronization in order to maximize the SNR seen at the receiver ADC. In addition, the digital baseband should understand where the start and end of a packet is in order to properly demodulate the received payload.

In conventional RF systems, synchronization between the transmitter and receiver is typically achieved by transmitting a known preamble code, and having the receiver compare all possible time-shifts of this known code to the signal it is receiving; in doing so the receiver will acquire the precise phase of the incoming signal. Comparison between all time shifts of the code under the presence of noise is typically achieved using a correlator structure—specifically, a matched filter. While this is indeed the optimal solution under a linear, AWGN channel assumption, the squaring element employed by the energy detector in this architecture is inherently non-linear. As a result, a matched filter is not the optimal solution. As discussed in [4], the optimal maximum likelihood solution involves the computation of Bessel functions, which are computationally inconvenient in a low-power implementation. To overcome this, the receiver employs a quadratic correlation technique that simplifies the maximum likelihood expression into one that is amenable to a low-complexity implementation while offering improved performance compared to a simple matched filter.

The digital baseband is comprised of 512 parallel quadratic correlators that computes all 2,048 possible preamble code shifts in a minimum of 14 μ s. A detailed description of the digital baseband design, including considerations regarding code choice and circuit-level optimizations, is discussed in [33].

2.3.8 Measurement Results

The receiver is implemented in a 90 nm CMOS process and a die photo of the chip is shown in Fig. 10. The die area is 2.6 mm by 2.1 mm, and the area is dominated by digital logic, which occupies the right side of the die. Due to the significant amount of digital logic integrated on the same die as the RF front end, there is significant potential for digital supply and substrate noise to result in degraded analog and RF performance. This motivated the use of a differential receiver architecture. Additionally, substrate contact rings are used to isolate the digital and analog blocks, as well as reduce the potential for feedback coupling in the high gain RF front end. The receiver is packaged in a 40-lead quad flat no-leads (QFN) package and mounted on an FR4 printed circuit board (PCB).

Fig. 10 Die photograph of pulsed UWB receiver SoC

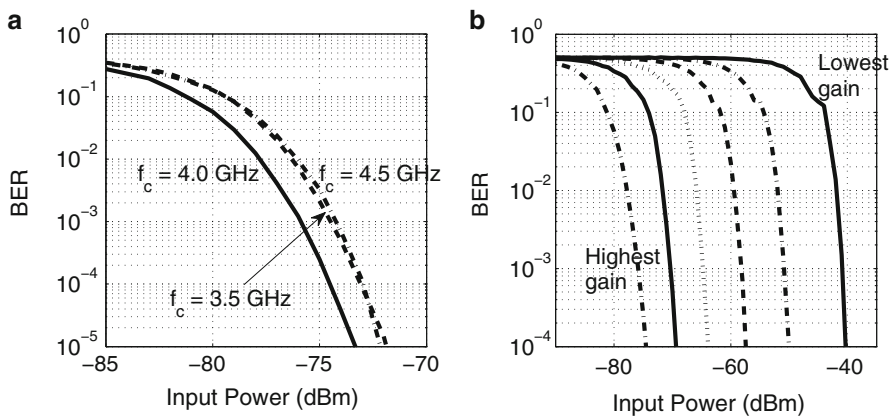
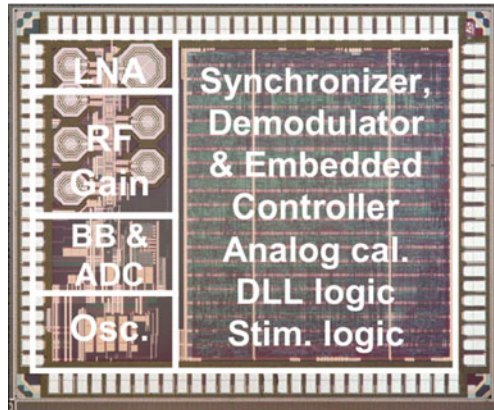


Fig. 11 BER of receiver (a) at its highest gain setting at the three center frequencies, and (b) at the different gain settings with $f_c = 4.0$ GHz

Figure 11 presents the BER of the receiver in different frequency bands at its highest gain setting and at different gain settings with $f_c = 4.0$ GHz. The receiver achieves a maximum sensitivity of -76 dBm at a data rate of 16 Mbps and a BER of 10^{-3} . The sensitivity scales by 35 dB from the lowest to highest gain setting, allowing for a trade-off of power consumption for sensitivity.

As the receiver SoC is targeted for low power, highly energy constrained applications, significant effort was spent to minimize overall power consumption and energy/bit. A breakdown of power consumption is shown in Table 2. Due to the extensive digital logic and the absence of power gating switches, the total leakage power is 0.64 mW. The always-on crystal oscillator consumes 0.15 mW. When the receiver is in idle mode, the majority of the clock tree is gated; however, an additional 0.13 mW of power is still consumed. The overall receiver power consumption is dominated by the LNA and the RF amplifiers that follow the LNA. Each individual RF amplifier consumes approximately 2.85 mW of power

Table 2 Receiver instantaneous power consumption breakdown

Receiver component	Power consumption (mW)
Leakage	0.64
Crystal oscillator	0.15
Clock tree (idle)	0.13
Delay locked loop	0.05
Baseband amplifier & ADC	1.51
LNA	5.90
RF Amplifier	0–14.30
Total idle power	0.92
Total active power	8.38–22.69

consumption, and the five-stage RF amplifier consumes a total of 14.30 mW of power when all five stages are enabled. At a data rate of 16 Mbps at the lowest gain setting, the entire receiver consumes 8.38 mW of power and at the highest gain setting, the receiver consumes 22.69 mW of power. When the receiver is duty cycled to low, kb/s data rates, the average power consumption is reduced to the order of a few milli-Watts, ultimately limited by leakage power. By adding power gating switches, the average receiver power consumption could approach the micro-Watt level at kb/s data rates. The receiver power consumption is constant regardless of the RF center frequency and includes the power of the digital backend when decoding data; however, these power measurements do not account for the energy required for synchronization at the start of a packet. As the receiver operates at an instantaneous data rate of 16 Mbps, the energy/bit of the receiver is 0.5-to-1.4 nJ/bit depending on the gain setting.

2.3.9 Receiver System Implementation

For the insect flight control system, some additional electronic components are required alongside the SoC. Figure 12 shows a block diagram of the electronics that are used. The key components include the receiver SoC, a microcontroller, 2.5 V DC-DC converter, 1 V low-dropout regulator (LDO) regulator, miniature coin cell battery, on-off switch, crystal resonator, LED, antenna, and discrete inductors, resistors and capacitors. The electronic components are soldered to a flexible, 4-layer PCB. A flexible PCB allows for a 60–70 % reduction in weight and thickness compared to a rigid PCB. Photos of the PCB are shown in Fig. 13. The entire system consumes an average power of 2.5 mW when the receiver attempts to receive a 68 bit synchronization packet every 1 ms.

The electronics are powered by a 1.4-to-1.6 V Silver-Oxide, size 362 coin cell battery that is capable of sourcing the 2.5 mW consumed by the electronics. The battery has a typical capacity of 27 mAh, weighs 0.32 g, and has an impedance at 40 Hz of 10-to-20 Ω . As the receiver SoC requires 1.0 and 2.5 V supply voltages, dc-dc converters are used to generate the required voltages from the battery.

Fig. 12 Block diagram of electronics mounted on a flexible PCB and attached to a moth

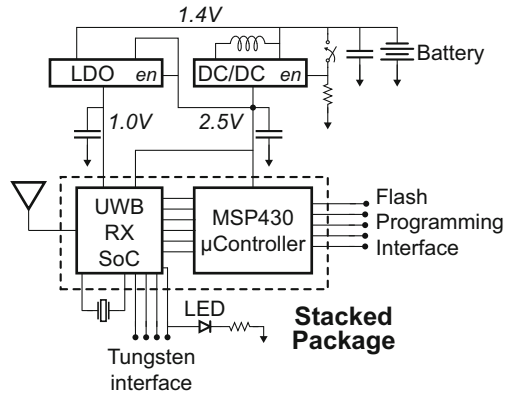


Fig. 13 Flexible PCB (a) top and (b) side

To further reduce form factor and weight, only a single decoupling capacitor is used for each supply voltage. A miniature on-off power switch is used to enable the dc-dc converters, so that the receiver does not consume any static current when turned off. Additional details on the flight control system implementation and measurement results are presented in [14].

2.4 Design Example: Implementation of a 9.8 GHz IR-UWB Receiver

Future biomedical and internet-of-things applications are driving the volume of wireless sensors into the cubic-mm regime, with power and volume requirements significantly more stringent than those demonstrated by the 3-to-5 GHz receiver described in the previous section. At the mm-scale, complete integration is necessary, and operation within the limits of a micro-battery becomes a primary challenge [11]. With CMOS scaling and ultra-low-power circuits reducing battery volume, the antenna and crystal quickly become the largest components in a cubic-mm node. In [7], a UWB transceiver for cubic-mm sensor nodes is demonstrated that achieves average receiver power levels of $37 \mu\text{W}$ at 30 kbps. Such a low average power consumption is achieved through extensive circuit and system optimizations,

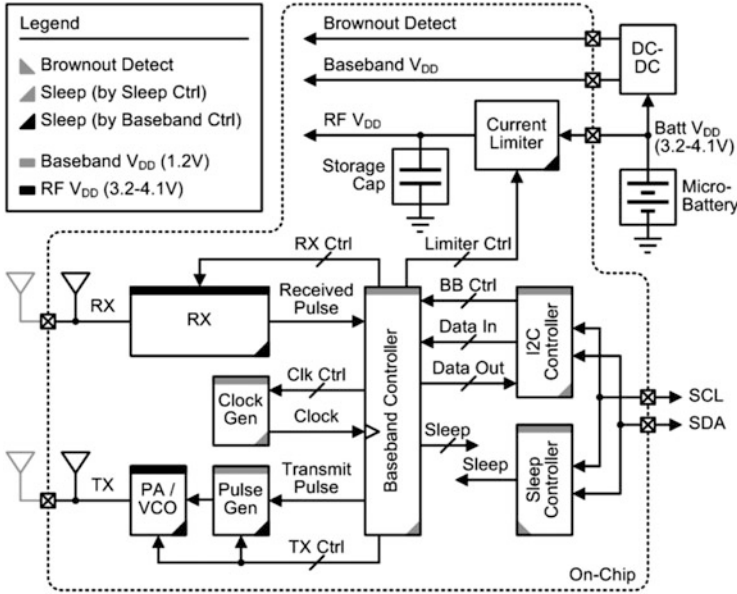


Fig. 14 System block diagram of the entire crystal-less UWB radio

including removing the need for a crystal oscillator, minimizing leakage current, fast duty cycling, operating circuits directly off a battery, and allowing for degraded sensitivity and range. This section briefly summarizes the receiver implementation and measurement results.

A block diagram of the receiver and its associated transmitter is shown in Fig. 14. The transmitter and receiver operate at the battery voltage, through a current limiter (CL) to protect the micro-battery from over-current and under-voltage. An internal storage capacitor allows higher current draws from the transmitter (TX) and receiver (RX) during duty-cycled operation. Digital baseband blocks operate from a 1.2 V V_{DD} to reduce power consumption. To survive on the limited resources of the micro-battery, all blocks on the radio have a low-power sleep state. RF and other analog blocks are duty-cycled at the bit-level by the baseband controller, while baseband blocks are duty-cycled at the packet-level by a separate sleep controller. The sleep controller remains on-continuously unless an under-voltage condition occurs. The sleep controller begins and ends the wake-up procedure for each packet via I2C communication with modified I/Os to eliminate pull-up resistors. The I2C controller provides bidirectional communication with other stacked die in a sensor node.

The receiver uses the non-coherent, energy-detection architecture shown in Fig. 15, similar to the 3-to-5 GHz receiver described in the previous section. Four RF gain stages amplify the 9.8 GHz UWB pulses before down-converting with a squaring mixer. The signal then passes through a baseband gain stage before the

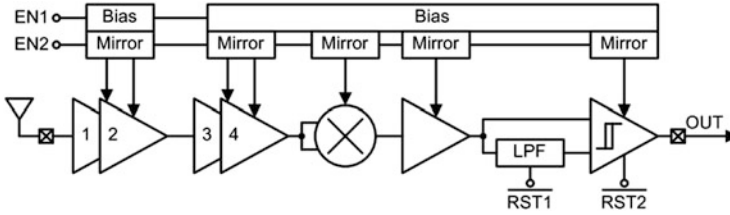


Fig. 15 Block diagram of the crystal-less UWB RX

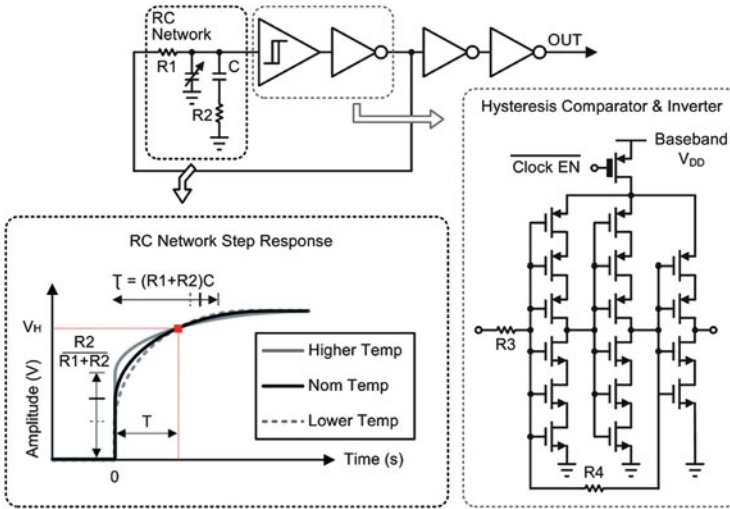


Fig. 16 Schematic of the temperature-compensated relaxation oscillator

signal path is split. Along one path, the pulses are passed directly to a comparator. The other path low-pass filters (LPFs) the signal to provide an auto-zeroed, DC-compensated reference level for comparison. A reset signal enables fast settling of the LPF for fast RX turn-on. Finally, a continuous-time latching comparator with controllable hysteresis digitizes the incoming pulses. BJTs are used for higher RF gain efficiency (gm/I), while the RF gain stages are stacked in order to reuse current and better utilize the supply voltage. The RF center frequency is tunable via 4 binary-weighted control bits. After RF amplification, the signal is self-mixed to dc using a common emitter amplifier with resistive load as a squaring mixer.

To reduce both power and area, the radio includes a relaxation oscillator with a modified RC network and a single-ended hysteretic comparator for on-chip clocking (Fig. 16). The RC network adds an additional zero in the transfer function from R2 over conventional relaxation oscillators, providing an additional degree-of-freedom for temperature compensation. As temperature increases, the initial step at $t = 0$ from the zero increases, but the time constant of the exponential decay also increases, offsetting the step and resulting in a constant time, T, to trigger

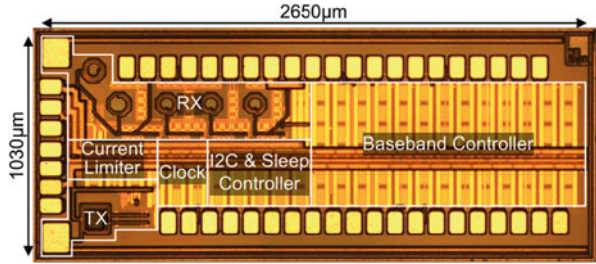
General	Process	0.18 μ m BiCMOS	RX Frontend	Average Power	37 μ W @3.6V
	Modulation	PPM		Sleep Power	347pW @3.6V
	Total Area	2.73mm ²		Center Frequency	9.7-10.2GHz
	Data Rate	30kb/s		Sensitivity	-67dBm @10 ⁻³ BER
	RF Voltage	3.2-4.1V			
	Baseband Voltage	1.2V			
	Sleep Power	1.0nW @3.6V 1.8nW @1.2V			
Baseband	Active Power	269 μ W	TX Frontend	Average Power	22.4 μ W @3.6V
	Clock	13 μ W		Sleep Power	170pW @3.6V
	I2C & Sleep Ctrl	10 μ W		Center Frequency	9-12GHz
	Baseband Ctrl	246 μ W		Output Power	0.1dBm
				Pulse Width	1.2-6.0ns
Clock	Frequency	3MHz	Current Limiter	Active Power	223nW @3.6V
	Supply Sensitivity	\pm 1.5% @ \pm 4% V _{DD}		Sleep Power	550pW @3.6V
	Temp Sensitivity	\pm 1.0% @0 to 50°C		Output Current	6-38 μ A
	Temp Coefficient	-584ppm/°C		Temp Sensitivity	-34nA/°C
				Line Sensitivity	0.21 μ A/V
				Efficiency	94% @3.6V

Fig. 17 Summary of the Radio Performance

the switching threshold, V_H , so that the overall period remains unchanged. The comparator consists of two stacked inverters with hysteresis levels set by R3 and R4. Stacking the FETs reduces leakage power while the oscillator is asleep, and a 5-bit capacitor bank is added to the oscillator for one-time process calibration of frequency. The oscillator has a measured variation of 1% over a range of 0–50°C that allows the TX and RX to be heavily duty-cycled between pulses in order to give the on-chip storage capacitor time to fully recharge and also sufficient accuracy to maintain network synchronization.

The radio was fabricated in 0.18 μ m BiCMOS with MIM capacitors. At a 10⁻³ BER, the RX has a sensitivity of -67 dBm and a 30 kb/s data rate while consuming an average of 37 μ W from a 3.6 V supply with 6% duty-cycling. The modem uses PPM and includes early/late tracking of pulses for each PPM window to maintain synchronization. At a 3 MHz oscillation frequency, the entire baseband system consumes 269 μ W, of which the clock consumes 12.7 μ W. The CL has a 6–38 μ A tuning range, which is sufficient for sustained operation of the TX and RX. The CL consumes only 223 nW, yielding a 94% efficiency. Each block consumes <1 nW while asleep by carefully including thick-oxide headers on all blocks, making this system ideal for heavily duty-cycled cubic-mm sensor nodes. A complete performance summary is provided in Fig. 17. The die occupies approximately 2.73 mm², dominated by the modem (Fig. 18). The entire radio is designed to operate from just the seven pads on the left edge to enable die stacking; the remaining pads are for debugging and may be left open.

Fig. 18 Die photo of the radio



3 IR-UWB Transmitter Design

This section will first review general classes of IR-UWB transmitters, followed by a detailed description of an example architecture [29, 31].

3.1 IR-UWB Transmitter Architectures

At the simplest level, an IR-UWB transmitter must generate narrow RF pulses and interface these pulses with an antenna, often through a power amplifier. Generation and radiation can typically be distinguished using two sets of criteria:

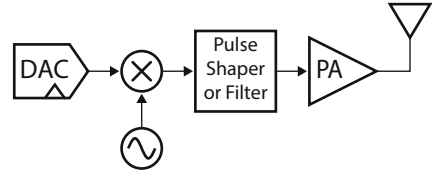
1. *RF generation:* There are two primary techniques used to synthesize UWB pulses at RF frequencies. The first technique involves mixing a baseband pulse with a local oscillator (LO) running at the desired RF center frequency. The second technique involves generating UWB pulses to lie directly at the desired RF center frequency. In other words, the second technique does not use a local oscillator.
2. *Power amplification:* There are two different techniques used to amplify and interface pulsed signals with an antenna. The first technique involves using analog circuits biased in their active regions for small-signal amplification and balanced conversions. The second technique uses digital circuits to buffer pulses at the interface to the antenna.

This criteria will be used to classify various pulse-generation techniques into four different categories. As a forewarning, it should be mentioned that it is sometimes difficult to make clear classifications, as some pulse generators use a combination of different techniques.

3.1.1 Traditional Small-Signal, Mixer-Based Transmitters

In these types of architectures, baseband data is typically converted from the digital to analog domain and subsequently mixed with an LO. The output of the LO is then amplified by an analog power amplifier (PA), often biased as class A or class AB in

Fig. 19 A traditional small-signal, mixer-based pulse generator architecture



order to meet linearity requirements. A simplified example architecture can be seen in Fig. 19. The initial popularity of this technique stemmed mainly from the fact that similar techniques are well established in traditional narrowband radio design.

From a signaling point of view, this type of architecture is the most robust, as both phase and amplitude modulation are possible.² Pulse shaping, used to attenuate RF sidelobes in order to meet FCC spectral masks, is also easily achieved in these types of architectures by either shaping the baseband data, or the RF data before the PA. For instance, the transmitter considered in [53] employed approximate Gaussian pulse shaping in the mixer by utilizing the exponential response of bipolar transistors.

The transmitter considered in [60] operates in dual bands by simultaneously up-converting two data streams onto two separate RF carriers. This is made possible by a wide bandwidth power amplifier employing shunt peaking and inductive feedback. A similar design is shown in [61], however only one band is operated in at a time. The transmitter supports all of the 802.15.4a specifications and reduces the power consumption over [60] by aggressively duty-cycling the class A power amplifier.

3.1.2 LC-Based Transmitters

LC-based transmitters use an LO to generate RF content, yet an explicit mixer is not necessarily required. For instance, a simple switch can either pass or block the LO output, thus effectively mixing the RF signal with a rectangular baseband pulse. As an example, the transmitter considered in [37] operates in a similar fashion to a superregenerative receiver; that is, the output of an LC oscillator is the transmitter output itself. A schematic is shown in Fig. 20. The oscillator can directly connect to the antenna (as a “power oscillator”), or an explicit PA can be employed.

In this example, a rectangular quenching pulse train acts as the baseband mixing signal. Like most oscillators, this circuit can be modeled as a second order system with poles in the right half portion of the s -plane. It is well known that the oscillatory output of such systems grow exponentially in time until circuit non-linearities limit the output swing. This oscillation growth can be leveraged to employ simple, low-overhead pulse shaping.

²Note that in-phase and quadrature paths are often used to enable quadrature amplitude modulation (QAM) for high data-rate communication.

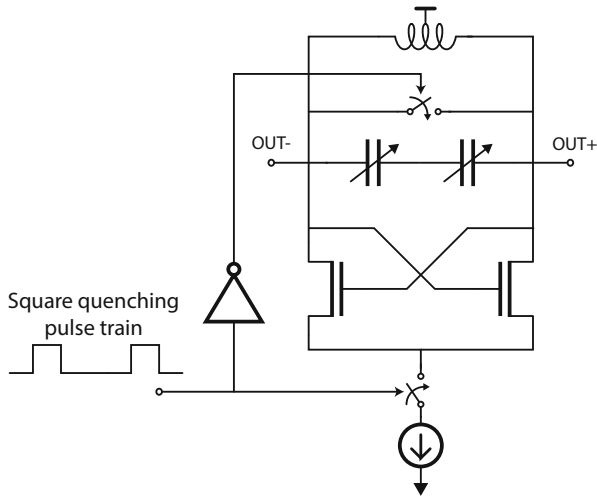


Fig. 20 An LC-based transmitter

3.1.3 Carrier-Less Transmitters

Carrier-less transmitters do not have an explicit local oscillator to mix baseband data up to RF. Instead, baseband data typically triggers a pulse generator to synthesize a pulse directly in the RF band of interest. One advantage 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 downconverting receiver typically cannot share the RF generation circuits and therefore must have a separate LO.

Implementation strategies typically involve generating pulses by combining edges of various delay elements, then amplifying the result using a power amplifier [24, 59]. Architectures that use delay lines to synthesize RF frequencies typically have half-RF cycles available at the output of each delay cell. By exploiting the fact that each half-RF cycle can be manipulated, simple pulse shaping schemes and differential-to-single-ended conversions are possible [17, 34, 62]. A popular architecture involves feeding half-RF cycles to alternating sides of a wideband balun, as shown in Fig. 21. This architecture ensures there is close to zero DC content at the transmitter output, thus enabling clean BPSK modulation.

3.1.4 All-Digital Transmitters

All-digital pulse generators attempt to reduce the power consumption over their analog counterparts by eliminating large static currents required to bias transistors in their active regions. Instead, digital static CMOS gates are used to generate high frequency rail-to-rail voltage swings. These digital architecture dissipate only $CV^2 f$

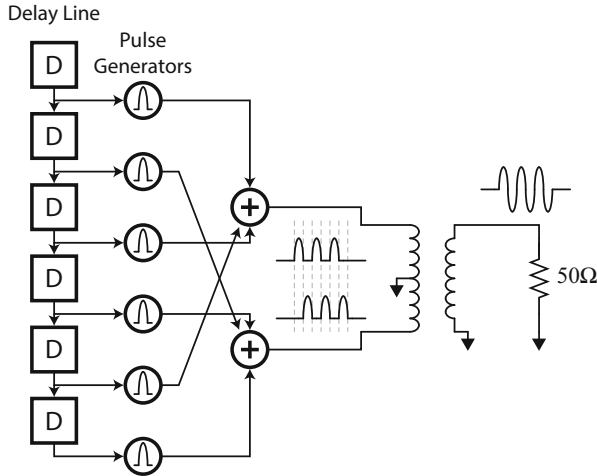


Fig. 21 A carrier-less architecture employing a balun for zero-DC voltage pulse generation

switching power and subthreshold leakage power. Since digital edges have harmonic content, pulse shaping and filtering may become necessary to reduce RF sidelobes.

A similar all-digital technique is popular in narrowband radio design, where linear power amplifiers are replaced by switched-mode power amplifiers. A major drawback of this approach is that constant-envelope modulation schemes must often be used, as switched-mode PAs have poor linearity and thus cannot support amplitude modulation techniques. Similarly, all-digital UWB transmitters are often restricted to phase and position modulations schemes only, unless clever pulse-shaping techniques are introduced.

Pulses in all-digital transmitters can either be synthesized using carrier-less techniques, or by modulating the output of a digitally-controlled oscillator (DCO). Examples of carrier-less techniques include the transmitters presented in [46, 47]. In these examples, UWB pulses are generated directly in the band of interest by NOR-ing two delayed edges together, converting from single-to-differential, and applying the differential signal to a dipole antenna. Similarly, the transmitter considered in [23] generates pulses by combining inverter gate delays using NOR and NAND structures. Relying on uncalibrated gate delays, however, leads to significant deviations in frequency and bandwidth targets over process voltage and temperature (PVT) variation.

The transmitter considered in [45] generates pulses in a carrier-less fashion by combining output edges from a delay line. BPSK modulation is achieved by applying full-swing pulses to either input of a balun, as illustrated in Fig. 22. However, a bandpass filter is required to reduce the low-frequency content typically associated with digital pulse generation driving a single-ended antenna.

The other technique to generate UWB pulses digitally is to modulate the output of a DCO. For instance, the transmitter considered in [44] generates pulses by

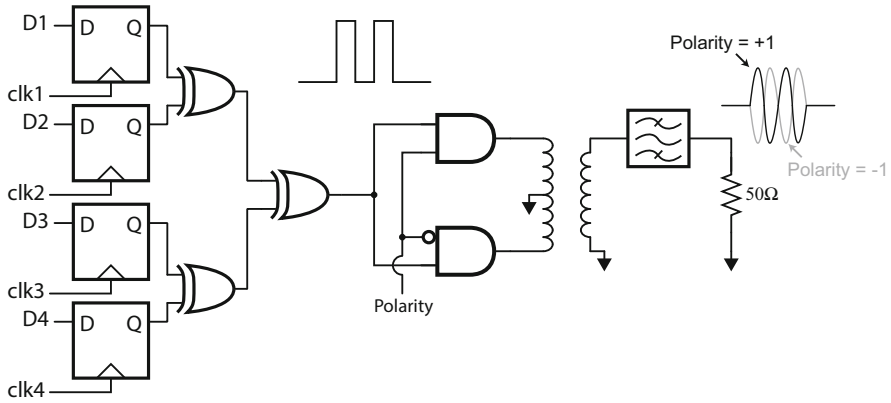


Fig. 22 An all-digital architecture employing a balun for BPSK modulation

directly modulating the output of a three-stage inverter-based ring oscillator. By utilizing the phases of an on-chip frequency divider, discrete two-level pulse shaping is employed. Since digital circuits are used in this architecture, reconfigurability and calibration are easily implemented.

3.2 Design Example: An All-Digital Non-coherent IR-UWB Transmitter Meeting FCC Spectral Masks Without Off-Chip Filters

3.2.1 Motivation for Non-coherent Transmitter Architecture

Coherent modulation schemes (e.g., BPSK or QAM) are generally more spectrally efficient than non-coherent modulation schemes (e.g., OOK or PPM), and thus in theory should be preferred. However, coherent modulation requires phase synchronization between the transmitter and receiver, resulting in more complex implementations that may not feature superior energy per bit. In addition, coherent IR-UWB systems suffer from significant multi-path fading, requiring high-complexity and power-hungry RAKE-based techniques for path consolidation [3].

While less spectrally efficient, non-coherent architectures often feature lower-complexity architectures, resulting in circuits that may consume lower power. Importantly, since precise phase information is not required and the RF bandwidth is large, a precise oscillator derived from a PLL is not required; instead, low-complexity, low-power RF generation techniques are available for use. In addition, non-coherent architectures can have inherent robustness to multi-path effects [51]. In an energy-detecting architecture, for example, the incoming signal is squared, then integrated over a set window of time. If the integration time window is set to be larger than the width of the pulse, the energy of several propagation paths will

be collected. Furthermore, the shape of the received pulse is no longer of concern to the receiver.³ Thus, non-coherent IR-UWB architectures have considerable promise in terms of energy efficiency, and thus a non-coherent architecture is employed in this design example.

3.2.2 Digital Pulse Generation

In principal, it is actually quite simple to design an all-digital IR-UWB pulse generator. For example, the transmitter shown in Fig. 23 uses an LO, a switch, and an inverter-based PA to generate and radiated pulsed-RF waveforms. Data in this transmitter can be modulated using OOK, or PPM, as illustrated in Fig. 24.

Naturally, this overly simple architecture suffers from several drawbacks including lack of programmability and calibration. Additionally, it is difficult to control the spectra, and as a result it is nearly impossible to meet the FCC spectra mask without dramatically reducing the average output power. As shown in Fig. 25, the resulting power spectral density of a square pulse train with non-zero DC content centered at 4 GHz clearly surpasses the FCC indoor mask.

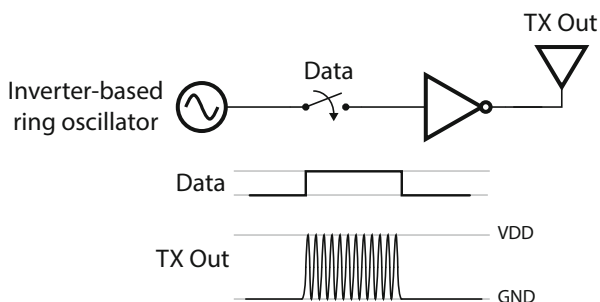


Fig. 23 A simple way to generate UWB pulses using all-digital circuits



Fig. 24 Pulse position modulation represents data by the presence of a pulse in a particular window in time

³The transmitter must still adhere to any pulse shape regulations to be standards-compliant.

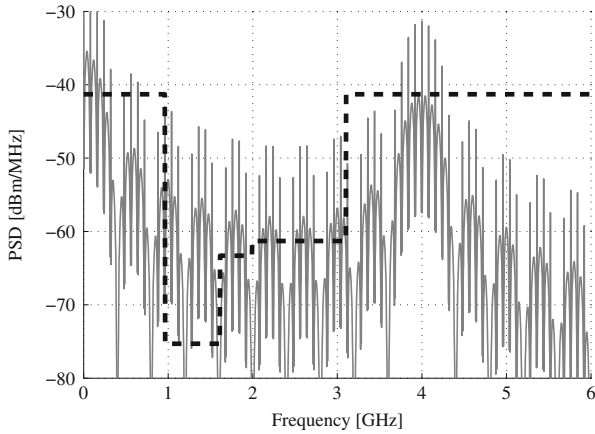


Fig. 25 Power spectral density of a train of PPM-modulated square UWB pulses

3.2.3 Achieving Spectral Compliance

There are three main problems with the spectrum shown in Fig. 25, all three of which must be addressed in order to meet the FCC spectral mask:

1. Large spectral lines spaced at integer multiples of the pulse repetition frequency.
2. Sidelobes centered at the carrier frequency.
3. Sidelobes centered at DC.

Spectral Lines

The problem of spectral lines is conceptually easy to fix. If the UWB pulses were phase modulated with random (or pseudo-random) data during transmissions, the tones would be scrambled out. This effect is most easily achieved by implementing a BPSK scrambler or modulator. The resulting power spectral density is illustrated in Fig. 26.

RF Sidelobes

This spectrum of Fig. 26 still suffers from drawbacks two and three: namely, it contains undesired sidelobes centered at both RF and DC. Although these sidelobes can easily be eliminated by bandpass or highpass filters, the area penalty of this approach is significant. For instance, this particular example would require at least a fourth order passive filter to ensure the necessary roll-off of roughly 20 dB in 0.29 decades (from the 1.61 to 3.1 GHz mask boundaries). A fourth order passive

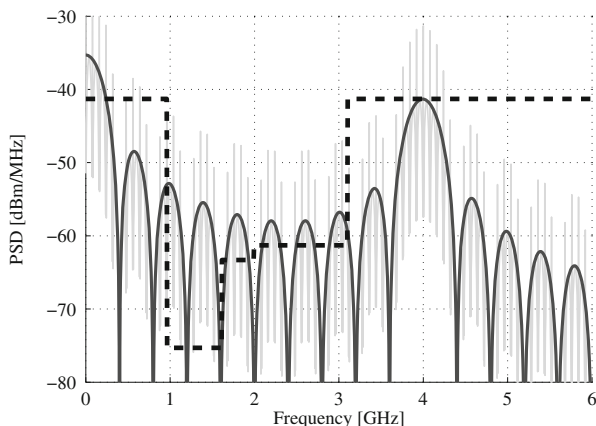


Fig. 26 Power spectral density of a train of PPM-modulated, BPSK-scrambled, square UWB pulses overlaid on top of the non-BPSK-scrambled case

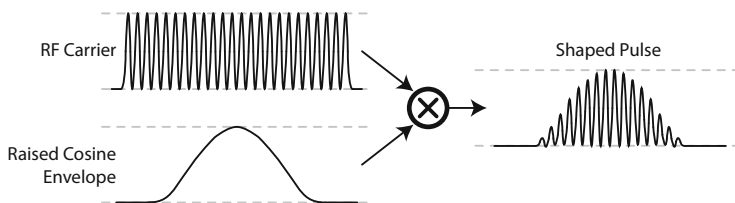


Fig. 27 Time domain view illustrating how to generate a raised-cosine pulse

filter requires several inductors and capacitors, which are not only lossy in modern semiconductor technologies, but also consume significant area.

An alternative to filtering is to employ pulse shaping to reduce the sidelobes centered at RF. As demonstrated in [9, 16, 34, 35, 62], and many other designs, pulse shaping is a very good method for obtaining high-order roll-off without the use of large passive filters. An excellent overview of several popular pulse shapes can be found in [52].

To illustrate the virtues of pulse shaping, consider shaping a pulse with a raised-cosine envelope, as illustrated in Fig. 27. The resulting spectrum achieves up to 17 dB of sidelobe rejection, as shown in Fig. 28.

Low Frequency Sidelobes

The raised cosine envelope greatly suppresses the sidelobes centered around the carrier frequency. However, the sidelobes centered at DC remain. This problem does not depend on pulse shape, but is rather fundamentally related to the method in which the digital pulses are synthesized. The issue stems from the fact that single

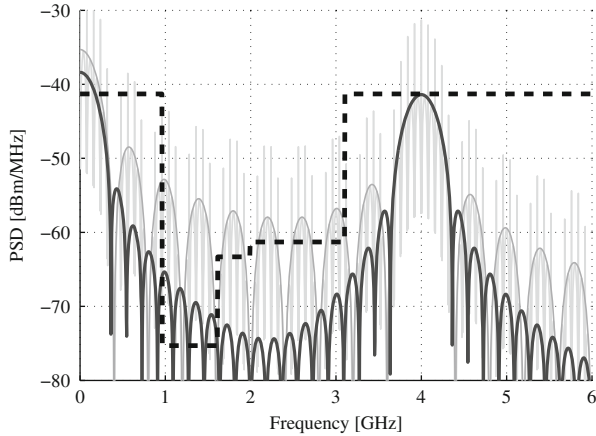


Fig. 28 Power spectral density of a train of PPM-modulated, BPSK-scrambled, raised-cosine UWB pulses overlaid on top of the spectrum in Fig. 26

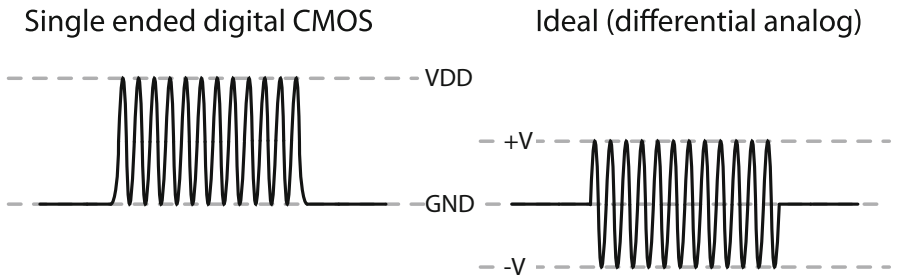


Fig. 29 Digital CMOS circuits can only generate one of two different reference levels. On the other hand, differential analog circuits can generate multiple reference levels at different bias voltages

ended digital circuits have only two stable operating points: the lowest and highest potentials in the circuits (typically GND and VDD). To eliminate the DC content and its associated sidelobes, the generated pulses must have *three* effective levels: GND, +V, and -V, as illustrated in Fig. 29.

For continuous wave systems, this is a relatively easy problem: simply insert an AC-coupling capacitor before the antenna, as illustrated in Fig. 30. This solution is unfortunately not ideal for pulses of short duration, the reason for which will become clear momentarily. Digitally generated pulses with two reference voltage levels (e.g. GND and VDD), can be decomposed into an RF carrier and a baseband pulse, as illustrated in Fig. 27. The baseband pulse will require a finite amount of time to charge and discharge the voltage across the capacitor, as shown in Fig. 31. The time required to charge and discharge, given by t_{charge} and $t_{discharge}$ respectively, is proportional to the RC time constant of the circuit.

Fig. 30 A coupling capacitor providing a DC block to a wideband antenna. A large resistor may be connected from the output to ground in order to provide a stable DC voltage to the antenna if necessary

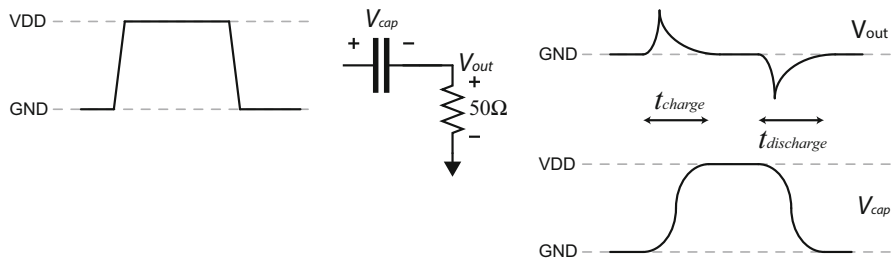
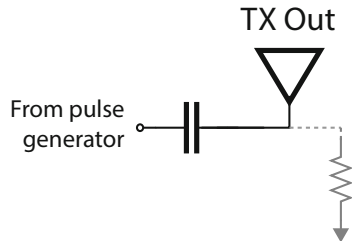


Fig. 31 A baseband pulse requires a finite amount of time to charge and discharge a coupling capacitor

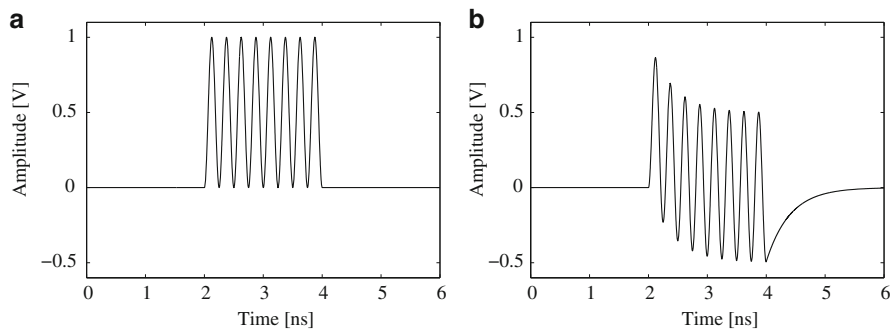


Fig. 32 The effect of passing a UWB pulse through an AC-coupling capacitor. (a) Digitally generated pulse. (b) After AC-coupling filter

The effect of finite low frequency capacitor charging and discharging times when AC-coupling UWB pulses is illustrated through a time-domain simulation of a square pulse in Fig. 32. It can be noted here that the AC-coupled pulse has a non-zero DC value, as well as some low-frequency turn-on and turn-off transients. The power spectral density of a train of BPSK-scrambled raised-cosine UWB pulses before and after the AC-coupling filter is shown in Fig. 33. The AC-coupled spectrum does not comply with the FCC mask, since the first order roll off is not sufficient to eliminate all of the low frequency sidelobes.

There are several techniques to reduce or even eliminate the low frequency-content of digitally generated UWB pulses. The most common technique relies on

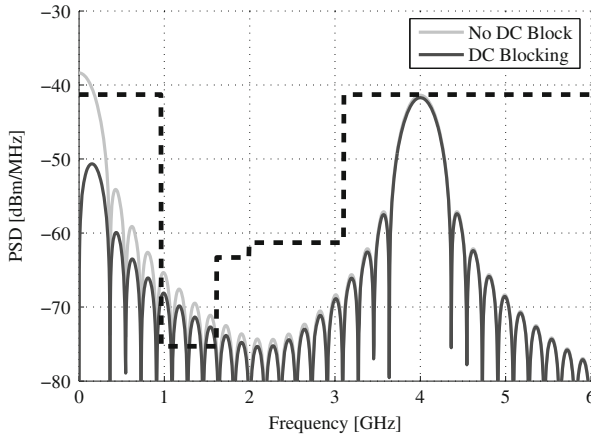


Fig. 33 Power spectral densities of raised-cosine pulses before and after an AC-coupling filter

generating individual half-RF cycles and applying them differentially to a wideband balun [34]. This technique produces excellent spectral results, however it requires the use of inductors which consume more-than-desired chip area. Another potential drawback of this type of architecture is that the half-RF cycles are generated from a delay line instead of a free-running ring oscillator. This can be seen as a benefit if the designed system is only a transmitter which generates a single pulse, then immediately turns off for a period of time. If, instead the designed system is a transceiver that transmits multiple pulses back-to-back, it is beneficial to design a single oscillator which is shared between the receiver and transmitter.

The proposed solution to attenuating the low frequency content using scalable digital structures involves capacitively coupling two paths which have differential baseband signals, yet contain in-phase RF tones. To elaborate, consider the network shown in Fig. 34a, where the two capacitors nominally have opposite DC voltages across them (GND and VDD, generated from digital logic). If they are driven with a differential baseband pulse, the upper capacitor will ideally charge at the same rate that the lower capacitor is discharging, thereby inducing zero voltage at the output.

If the low frequency baseband pulses are multiplied with in-phase RF tones as illustrated in Fig. 34b, then the low frequency common-modes will cancel, and the in-phase RF components will propagate to the output.

Since the two inputs into the capacitive combination network start off with opposite common modes, there is an inherent half RF cycle delay between the start of the effective baseband pulses shown in Fig. 34b. This, combined with circuits mismatches, will create non-idealities including turn-on and turn-off transients leading to spectral impurities. Ideally, the output spectrum will contain zero low frequency content, as illustrated by the spectrum of ideal raised cosine pulses in Fig. 35. In practice, the output spectrum will have a small amount of low-frequency content.

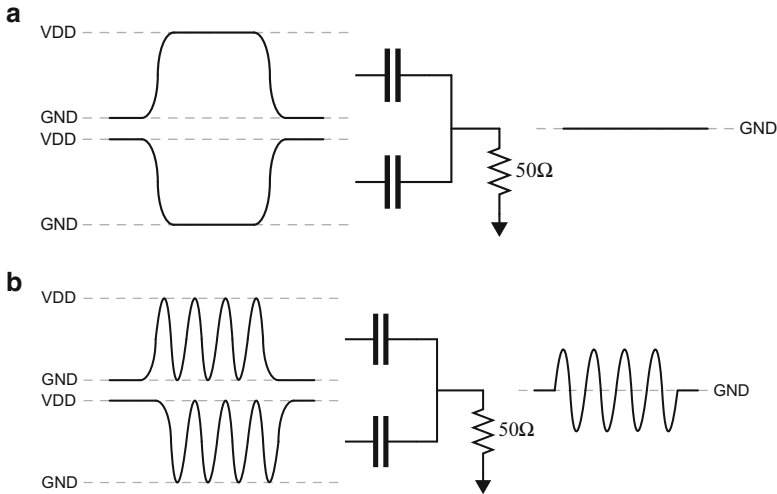


Fig. 34 Differential baseband pulses cancel as shown in (a), while in-phase RF signals propagate relatively undisturbed to the output, as shown in (b)

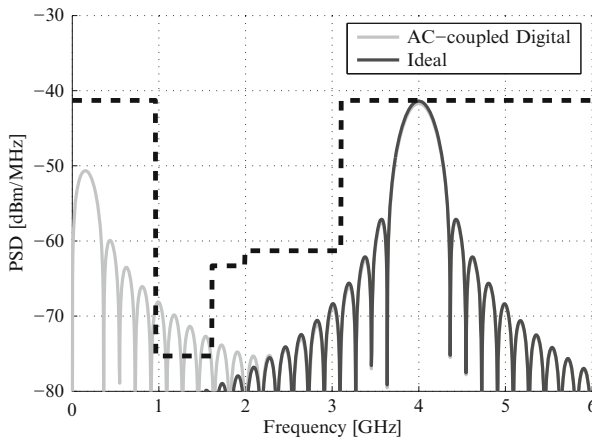


Fig. 35 Power spectral densities of ideal raised-cosine pulses and AC-coupled digitally generated raised cosine pulses

3.2.4 Transmitter Architecture

Given the dual capacitively-coupled paths, a block diagram of the presented transmitter is shown in Fig.36 [29, 31]. The transmitter is designed to operate in all three channels of the low-band group of the 802.15.4a standard. As per the 802.15.4a specifications, payload data is modulated using time-hopped (TH)-PPM, where a PPM symbol is represented by a burst of several back-to-back pulses contained in a fixed window of time [27]. In idle mode between bursts, all transmitter circuits are off and the transmitter consumes only leakage power.

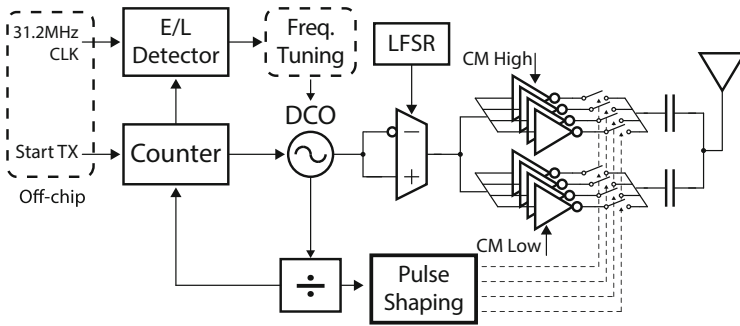


Fig. 36 Transmitter block diagram

Pulse bursts are generated on the rising edge of the off-chip *Start-TX* signal. This edge enables a DCO, whose output is BPSK-scrambled via an linear feedback shift register (LFSR) and subsequently buffered through dual single-ended digital PAs employing capacitive combination.

The DCO output frequency is calibrated and dynamically adjusted using an early-late detector in a digital frequency locked loop (FLL). The DCO output is also synchronously divided to a 499.2 MHz clock as specified by the 802.15.4a standard [22]. Several phases of the divided clock are used by pulse shaping circuitry to dynamically shape the PA envelope to one of four discrete levels. The 499.2 MHz clock sets the pulse repetition frequency (PRF) within a burst, and is also used in conjunction with a counter to program the number of pulses transmitted per burst.

3.2.5 Dual Digital Power Amplifiers

The circuit shown in Fig. 37 implements the dual capacitor technique in order to generate low-DC content RF pulses by driving two 2 pF coupling capacitors with two separate digital PAs.

Each PA consists of 30 tri-state inverters. A single oscillator signal is fed as an input to all 60 tri-state inverters, thus ensuring both paths receive in-phase RF signals. Each tri-state inverter is sized such that all 60 inverters operating in parallel can drive the antenna and associated parasitics up to 800 mV when switching at 4 GHz. Output power control can be configured by programming the number of tri-state inverters enabled at a given time. Furthermore, by dynamically adjusting the number of enabled tri-state inverters during pulse transmission, pulse shaping can be realized. Section 3.2.7 discusses the implementation details of the pulse shaping logic.

The differential baseband pulses (i.e. opposite common mode low-frequency pulses) are generated by ensuring that the outputs of the two PAs are at opposite supply rails immediately before and after pulse generation. Thus, during pulse

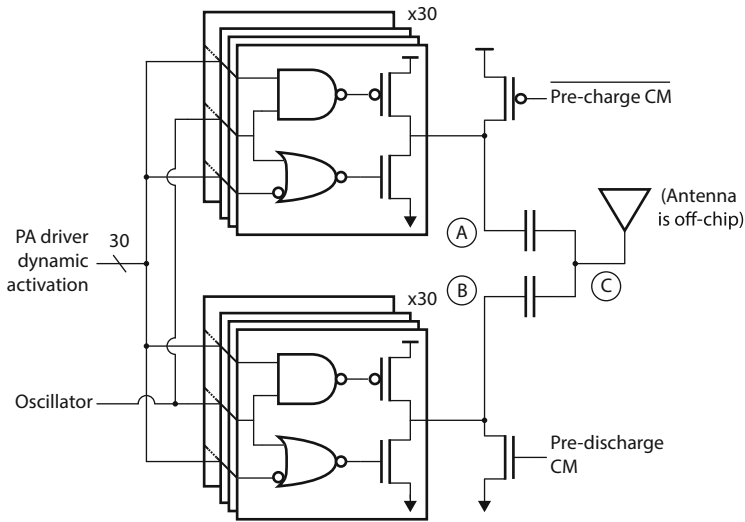


Fig. 37 Dual digital power amplifiers

generation the two capacitively coupled paths begin to charge and discharge low-frequency content at the same rate, resulting in close to zero low-frequency content on the output.

The opposite common modes for the two PA outputs are set by pre-charge and pre-discharge transistors during the idle mode between pulses (i.e. when the PA outputs are tri-stated). The dynamic PA control logic should ensure that the pre-charge and pre-discharge transistors are never turned on during pulse generation in order to avoid static power dissipation. It should be mentioned that since the PA outputs can be tri-stated, the transmitter can easily share the antenna with an integrated receiver without requiring an explicit transmit/receive switch. If required, the DC voltage of node C can be set with a large resistance or inductor to GND in order to eliminate any potential build up of charge.

Figure 38 shows a representative timing diagram of the dual digital power amplifiers with pulse shaping applied. Since the coupling capacitors are charging and discharging at roughly the same rate, the average voltage of nodes A and B approach the same value (ideally $V_{DD}/2$) during pulse generation. For this reason, a very visible low-frequency transient is seen on nodes A and B at the end of pulse generation when the pre-charge/discharge devices are turned on. If the two paths are matched and the pre-charging and pre-discharging begin at the same voltage on nodes A and B, this low-frequency transient will not be seen at the output (node C).

However, if the two paths are not matched, nodes A and B will discharge with different initial conditions, thus leading to some low-frequency content on output node C. For example, Fig. 39 shows the simulated output spectrum of the dual PAs with ideally-matched paths overlaid on top of Monte Carlo process variation, showing up to 4 dB degradation at both DC and 1.2 GHz. Since there are relatively

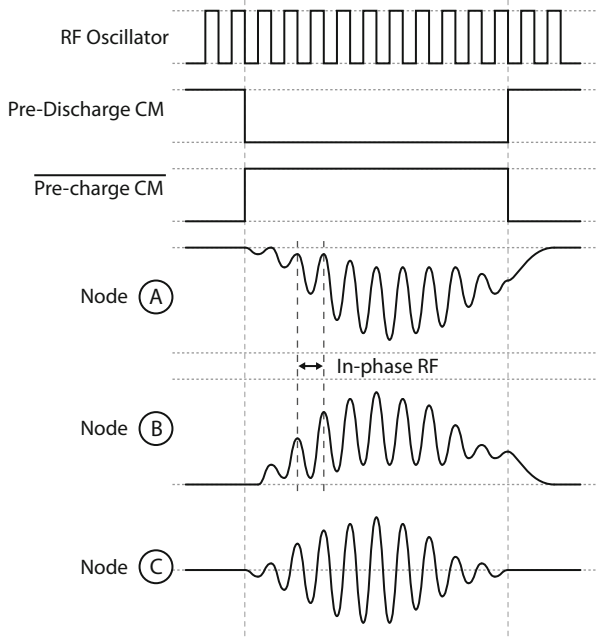


Fig. 38 Timing diagram for the dual digital power amplifiers

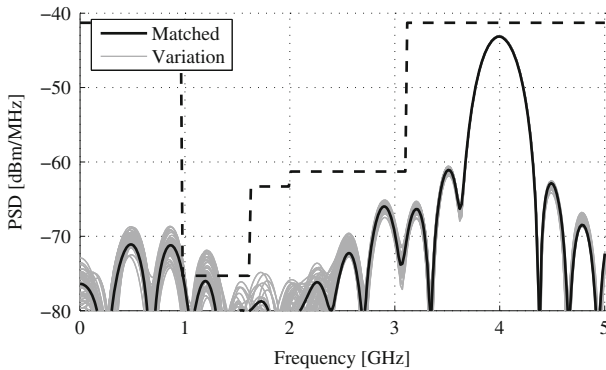


Fig. 39 Simulated output spectra with Monte Carlo process variation

small spectral differences across process variation and device mismatches, a large number of pulse shape configurations should be able to guarantee, with a reasonable degree of confidence, that the FCC mask will be met. This idea of implementing *redundancy* in order to guarantee desired operation is almost necessary in high density memory design, and is becoming more popular for other types of circuits such as ADCs [13, 21].

3.2.6 Clocking

Since non-coherent pulsed-UWB receivers have large input bandwidths *and* discard phase information, precise transmitter frequency tolerances are not necessarily required. To quantify this claim, consider a UWB transmitter with a 6,000 ppm accuracy. At a 4 GHz center frequency, this corresponds to a maximum frequency error of ± 12 MHz. If an ideal receiver with a 500 MHz brick wall input filter received a 500 MHz input signal offset by 12 MHz, a loss of only 0.1 dB would be incurred. The situation typically improves when dealing with non-ideal signal bandwidths and filters. As an example, the system presented in [55] has a transmitter center frequency accuracy of 6,000 ppm. While unacceptably large for coherent and/or narrowband systems, the receiver still achieves a sensitivity of -99 dBm at a BER of 10^{-3} and a data rate of 100 kbps. As a result, the transmitter oscillator design requirements can be relaxed considerably with the ultimate goal of improving energy efficiency.

A DCO that meets the needs of this design is shown in Fig. 40. The current-starved inverter-based three-stage ring structure is designed to have a fast turn-on time on the order of 2 ns to reduce energy consumption in duty-cycled operation. Furthermore, the delay elements are all single-ended to further reduce energy consumption over a differential structure. Although single-ended structures are more susceptible to power supply noise compared to their differential counterparts, the resulting increase in phase noise is of negligible concern to a non-coherent energy-detecting receiver.

Coarse frequency tuning is provided by switchable load capacitors, while fine frequency tuning is provided with NMOS and PMOS current starving DACs. To simplify the frequency locking algorithm, all three current starving DACs are set to the same digital value, except that the second and third stage DACs can be individually incremented by one for increased resolution. This technique results in a resolution of 7.5 bits from the DACs and 2 bits from the three thermometer encoded capacitors, totaling 9.5 bits.

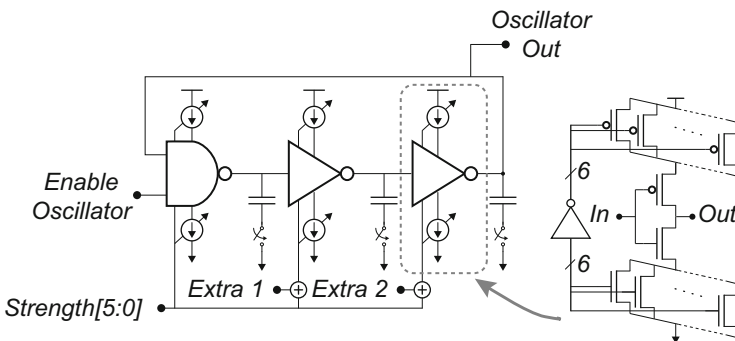


Fig. 40 Digitally controlled oscillator

The output of the DCO is fed to a programmable synchronous frequency divider. The divider is realized using true single-phase clock (TSPC) logic [57] in order to accommodate inputs up to 6 GHz. The divider consists of fourteen half-transparent latches (HTLs) which can be individually bypassed, thus allowing a programmable divide ratio of up to fourteen. The design is based on the work presented in [44] and [58].

The output of the divider drives the pulse shaping circuitry, which in turn determines the effective transmitted pulse width (and thus the bandwidth). To comply with the 802.15.4a standard, the transmitted pulse width is maximally set to one over the 499.2 MHz PRF within a burst (i.e. 2 ns). If the DCO frequency is set to one of the 802.15.4a channels, an integer division will always yield the required 499.2 MHz clock.

The transmitter contains an early/late detector which can be used for periodic frequency calibration in a frequency locked loop. In addition, the transmitter employs BPSK scrambling in order to smooth out spectral lines associated with non-phase modulated PPM signaling while maximizing peak power. More details of both of these blocks and their design considerations can be found in [30, 31].

3.2.7 Pulse Shaping Logic

The output phases of the frequency divider do not necessarily have 50 % duty cycles. In fact, the duty cycles of the internal phases vary approximately linearly from 10 % to 90 %, depending on the divide ratio. Since the period of each phase is set to be equal to the pulse width of 2 ns, it is possible to combine several of these phases in order to generate the timing required for pulse shaping. This is illustrated in Fig. 41,

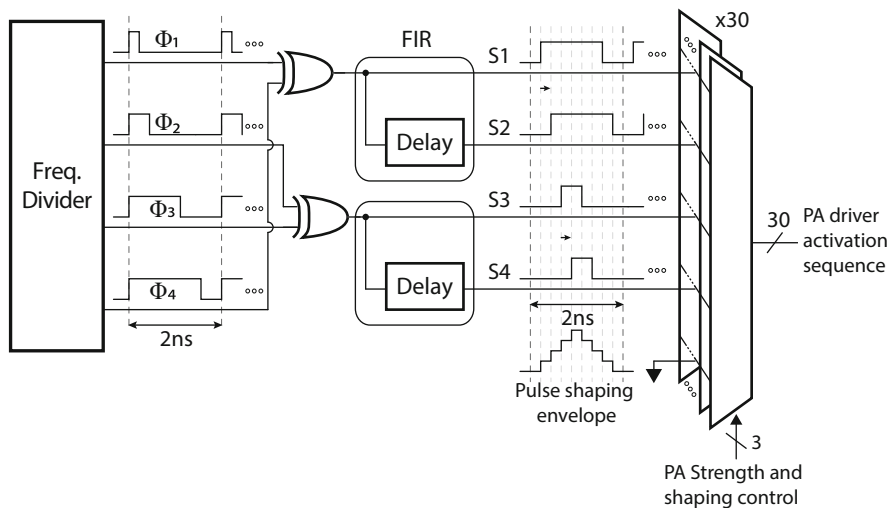


Fig. 41 Pulse shaping logic

where signals Φ_{1-4} are appropriately chosen to have duty cycles of approximately 20 %, 40 %, 60 %, and 80 %, based on which HTLs are enabled or disabled.

By XOR-ing Φ_1 with Φ_4 and Φ_2 with Φ_3 , two pulse shaping signals are generated. These two pulse shaping signals are each passed through simple one-tap finite impulse response (FIR) filters to increase the number of pulse shaping signals to four (signal S1-S4). The delay elements of the FIR filters are simply comprised of a programmable number of inverter-based buffers.

Each tri-state inverter of the dual PAs is individually programmed through a five-input multiplexer network to receive one of the four pulse shaping signals as a dynamic activation input. The fifth multiplexer input is grounded in order to allow statically disabled tri-state inverters, as the inverters are typically disabled to perform gain control. The four pulse shaping signals can be thought of as the output of FIR filter taps which are added together at the input of the coupling capacitors via the parallel combination of PA tri-state inverters. Maximum PA output swing is achieved when all four pulse shaping signals are high, i.e. the maximum number of PA inverters are enabled in parallel simultaneously. This pulse shaping configuration also ensures that the output signal amplitude is zero during BPSK phase transitions in order to avoid common-mode glitching and inter-pulse interference.

The actual pulse shape can be modified through two different methods. One method involves changing the relative times at which the shaping signals arrive. This can be accomplished by selecting different frequency divider phases or changing the FIR filter delay element. This does not typically produce a desired pulse shape; an example is shown in Fig. 42a. The FIR filter delays have $2^3 = 8$ possible permutations. The second method involves changing how many tri-state inverters receive a particular pulse shaping signal. This is similar to choosing the weights of the FIR tap coefficients, and can be used to more closely approximate a raised-cosine shape [35]. An example pulse shape is shown in Fig. 42b. There are approximately $2^{30} \approx 10^9$ total pulse shape strength permutations. Many of these permutations

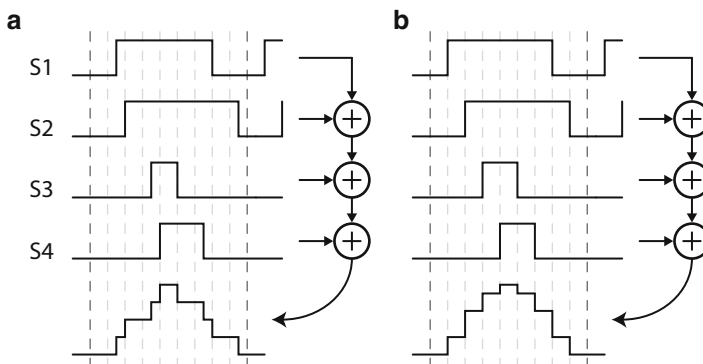


Fig. 42 Modifying pulse shapes by changing (a) delays and (b) weights

Fig. 43 Die photo of fabricated transmitter

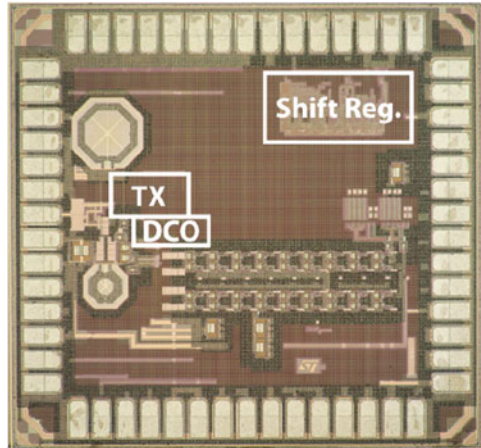
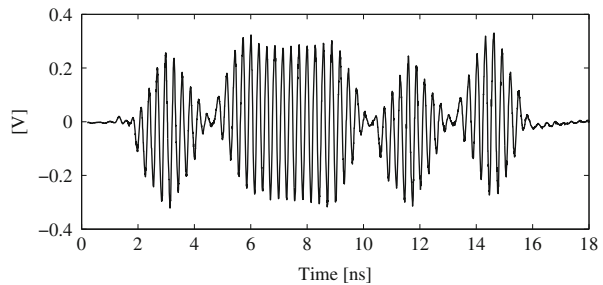


Fig. 44 Measured transient waveform of a burst of five individually BPSK-modulated pulses



are not practically useful, however as discussed in Sect. 3.2.5, there are more than enough possible configurations to guarantee FCC spectral compliance with a reasonable degree of confidence.

3.2.8 Measurement Results

The transmitter was fabricated in a 90 nm bulk CMOS process and packaged in a 40-lead, wirebonded QFN package; all measurement results were taken from the packaged chip. A die photograph is shown in Fig. 43b. The transmitter core and DCO consume an area of 0.07 mm².

The transmitter operates at data rates from 0-to-15.6 Mbps on a 1 V power supply. It has a turn-on time of 7.2 ns, measured as the time it takes pulses to appear at the output of the dual PAs after the rising edge of *Start-TX* has arrived.

Figure 44 shows the output when the transmitter is configured to generate a burst of five individually BPSK-modulated pulses at a time as measured by a Tektronix TDS 8000 Sampling Oscilloscope with an 80E04 sampling module. The dual PAs were measured to have an output voltage swing range from 160-to-710 mV, resulting in an output power control range of 13 dB.

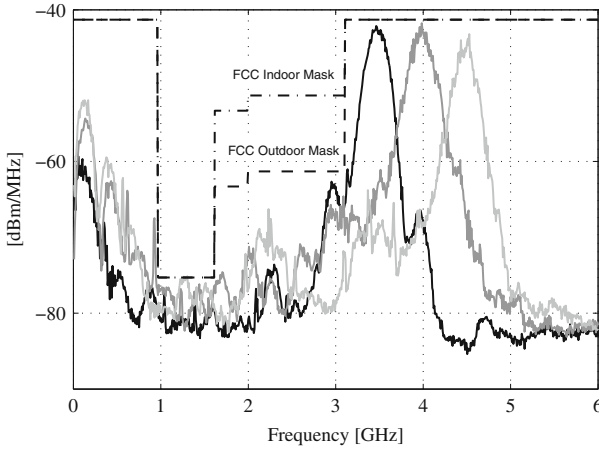


Fig. 45 Overlaid power spectral densities of the three channels low 802.15.4a bands

Frequency domain measurements were taken at the output of the capacitive combination network using an Agilent MXA N9020A Spectrum Analyzer. Unless otherwise noted, the resolution bandwidth was set to 1 MHz. During normal operation, the proposed transmitter achieves both indoor and outdoor FCC spectral compliance in all three low-band 802.15.4a channels. The resulting spectra for sixteen-pulse bursts are superimposed together over the FCC spectral mask in Fig. 45. Note that no off-chip filters were used to make this measurement.

To determine how effective capacitive combining and pulse shaping are at reducing both low frequency content as well as RF sidelobes, Fig. 46 shows the output power spectral density of two-pulse bursts with: capacitive combining and no pulse shaping, pulse shaping and no capacitive combining, and capacitive combining and pulse shaping. Here it can be seen that capacitive combination achieves up to 12 dB of low-frequency attenuation. Furthermore, pulse shaping achieves 15-to-20 dB of sidelobe rejection.

Operating on a 1 V supply, the transmitter draws 4.36 mW when generating 16-pulse bursts at an symbol repetition frequency (SRF) of 15.6 MHz. The total output in this configuration is -16.4 dBm. This results in an energy efficiency of 280 pJ/burst, or 17.5 pJ/pulse. Since all transmitter circuits are inherently off between pulse transmissions, the power consumption scales with data rate. However, the impact of leakage power becomes significant at symbol rates below 1 MHz. The standby (i.e. idle-mode) leakage power is 123 μ W. Table 3 summarizes the transmitter's performance.

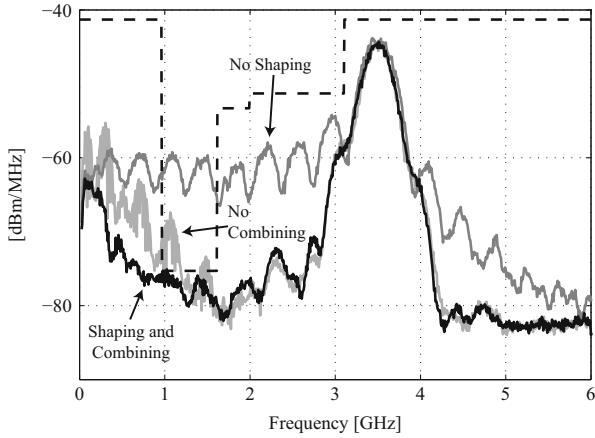


Fig. 46 Overlaid power spectral densities with shaping disabled, combining disabled, and normal operation

Table 3 Transmitter performance summary

Specification	Value	
Process	90 nm CMOS	
Active die area	0.07 mm ²	
Modulation	PPM+BPSK	
SRF range	0-to-15.6 MHz	
Supply	1 V	
	Standby power	Active E/pulse
Power amplifier	22 μ W	2 pJ
DCO/clock/control	83 μ W	15 pJ
Shift register	11 μ W	–
I/O and ESD	7 μ W	<1 pJ
Total	123 μ W	17 pJ
Output voltage swing	165 mV-to-710 mV	
DCO frequency range	2.1 GHz-to-5.7 GHz	
Turn-on time	7.2 ns	
Symbol rate (SRF)	100 kHz	15.6 MHz
Energy/16-pulse-burst	1.6 nJ	280 pJ
Energy/pulse	103 pJ	17.5 pJ
Total power	164 μ W	4.36 mW
Output power	–25 dBm	–16.4 dBm

4 Summary

In this chapter, we presented two IR-UWB receivers and one IR-UWB transmitter that consume low power for low data rate and short range radio applications. We have shown that IR-UWB is a viable technology choice that provides competitively

low energy consumption per bit (17.5 pJ/bit TX, 700 pJ/bit RX) when compared to narrowband radio counterparts with similar operating requirements. Furthermore, we have shown that IR-UWB systems require fewer off-chip components for a full transceiver implementation. IR-UWB transmitters do not require much analog design effort since they can be built with all-digital logic and digitally synthesized. They do not require high Q filters; instead, simple digitally-configurable pulse-shaping conforms the transmit power to the spectral mask requirements. Though the average power density for UWB signaling is restricted to be below the noise floor (-41.3 dBm/MHz), it is more energy efficient to transmit the largest allowable instantaneous power in each pulse (less than 0 dBm in a 50 MHz resolution bandwidth) rather than a series of lower power pulses. Finally, since the transmitters are built with all-digital logic, they consume no static power and are inherently duty-cycled between pulse transmissions.

IR-UWB receivers can leverage non-coherent signaling and large signal bandwidth to duty cycle the receiver between pulses, much like the transmitter. However, since RF/analog circuits must be used in the LNA, RF, and baseband gain stages, it requires some effort to ensure these blocks can power up/down as fast as possible for maximum energy savings. Because the signal bandwidth is >500 MHz wide, the LNA and RF bandpass filters are much more robust to any changes in absolute frequency than their narrowband radio counterparts, thereby easing the design effort. The pulses can be converted from RF to baseband without an LO by utilizing a self-mixer. However, since the self-mixer requires at least 10 mV of input amplitude to produce a reasonable output (0.7 mV), RF gain stages are required between the LNA and the self-mixer. Both radios that were presented leverage these aspects of IR-UWB. In addition, we have shown two methods of demodulating data after the self-mixer. One method uses a simple DC-tracking analog threshold voltage and hysteretic comparator. These two blocks operate together to function as a single-bit data converter. This method is the simplest if sensitivity and flexibility are not as important as power consumption. The other method involves an integrator and a 5-bit digitizer, both operating at a 31.5 ns period long enough to capture a typical UWB channel impulse response. The advantage of this more complex demodulator is the increased level of visibility into timing and voltage amplitudes of the RF channel. With this level of detail, the digital baseband is able to synchronize with an accuracy of ± 1 ns, and have the flexibility to demodulate both PPM and OOK IR-UWB signals. The main outstanding challenge for UWB receivers lies in the fact that since the UWB signal bandwidth overlays with many other narrowband standards with much larger allocated power levels, immunity to these narrowband interferers must be developed at the signaling, circuit, and even network levels.

At the system level, we have shown that IR-UWB transceivers do not require accurate RF synthesizers nor external crystal references, because the large signal bandwidth along with non-coherent signaling allows the system to tolerate larger RF frequency differences ($>1,000$ ppm) with a negligible (approximately 0.1 dB @ 6,000 ppm) impact on sensitivity; however, it should be noted that if the bit-period clock accuracy between TX and RX is improved, the system could exhibit improved performance and further power savings. At the network layer, the order of magnitude

difference in energy/bit between TX and RX may also spur development of new networks of wireless sensors where the overall system can be energy-optimized by leveraging the cheaper cost of data transmission using IR-UWB signaling.

From the work presented here, it is clear that IR-UWB radio systems are a viable technology choice in applications that require energy efficient, low rate wireless communication. Additionally, we have shown IR-UWB systems can be reliably implemented in silicon, and can operate with fewer off-chip components than narrowband radio counterparts because they do not need a transmit filter nor an external crystal oscillator. Finally, if the application requires more transmitters than receivers, wireless communication systems built on IR-UWB technology would be a natural fit to allow an overall lower network energy consumption. If these features are important in the application that is under consideration, IR-UWB radios could be the right technology choice.

References

1. M. Baghaei-Nejad, D.S. Mendoza, Z. Zou, S. Radiom, G. Gielen, L.-R. Zheng, H. Tenhunen, A remote-powered RFID tag with 10Mb/s UWB uplink and -18.5 dBm sensitivity UHF downlink in $0.18 \mu\text{m}$ CMOS, in *Proceeding of the IEEE ISSCC Digest Technical Papers* (February, 2009), pp. 198–199
2. ABM10: Ceramic SMD Ultra Miniature Quartz Crystal, Abracon, Technical Report (2008)
3. M.-G.D. Benedetto, T. Kaiser, A.F. Molisch, I. Oppermann, C. Politano, D. Porcino, *UWB Communication Systems: A Comprehensive Overview* (Hindawi Publishing Corporation, New York, 2006)
4. M. Bhardwaj, Communications in the observation limited regime, Ph.D. thesis, Massachusetts Institute of Technology, 2009
5. J.L. Bohorquez, A.P. Chandrakasan, J.L. Dawson, A 350 W CMOS MSK transmitter and 400 W OOK super-regenerative receiver for medical implant communications. *IEEE J. Solid State Circuits* **44**(4), 1248–1259.
6. A. Bozkurt, R. Gilmour, D. Stern, A. Lal, MEMS based bioelectronic neuromuscular interfaces for insect cyborg flight control, in *Proceedings of the IEEE International Conference on MEMS Systems* (January, 2008), pp. 160–163
7. J. Brown, K.-K. Huang, E. Ansari, R. Rogel, Y. Lee, D. Wentzloff, An ultra-low-power 9.8ghz crystal-less uwb transceiver with digital baseband integrated in $0.18 \mu\text{m}$ bimos, in *IEEE ISSCC Digest of Technical Papers* (February, 2013), pp. 442–443
8. B.H. Calhoun, D.C. Daly, N. Verma, D. Finchelstein, D.D. Wentzloff, A. Wang, S.-H. Cho, A.P. Chandrakasan, Design considerations for ultra-low energy wireless microsensor nodes. *IEEE Trans. Comput.* **54**(6), 727–749 (2005)
9. X. Chen, S. Kiaei, Pulse generation scheme for low-power low-complexity impulse ultra-wideband. *IEE Electron. Lett.* **43**(1), 44–45 (2007)
10. J.-Y. Chen, M. Flynn, J. Hayes, A fully integrated auto-calibrated super-regenerative receiver, in *Proceeding of the IEEE ISSCC Digest Technical Papers* (2006), pp. 1490–1499
11. G. Chen, S. Hanson, D. Blaauw, D. Sylvester, Circuit design advances for wireless sensing applications. *Proc. IEEE* **98**(11), 1808–1827 (November, 2010)
12. P. Choi, H. Park, I. Nam, K. Kang, Y. Ku, S. Shin, S. Park, T. Kim, H. Choi, S. Kim, S.M. Park, M. Kim, S. Park, K. Lee, An experimental coin-sized radio for extremely low power WPAN (IEEE802.15.4) application at 2.4GHz, in *Proceeding of the IEEE ISSCC Digest Technical Papers* (February, 2003), pp. 92–480

13. D.C. Daly, A.P. Chandrakasan, A 6b 0.2-to-0.9V highly digital flash ADC with comparator redundancy, in *IEEE ISSCC Digest Technical Papers* (February, 2008), pp. 554–555
14. D. Daly, P. Mercier, M. Bhardwaj, A. Stone, Z. Aldworth, T. Daniel, J. Voldman, J. Hildebrand, A. Chandrakasan, A pulsed UWB receiver SoC for insect motion control. *IEEE J. Solid-State Circuits* **45**(1), 153–166 (2010)
15. H. Darabi, S. Khorram, Z. Zhou, T. Li, B. Marholev, J. Chiu, J. Castaneda, E. Chien, S. Anand, S. Wu, M. Pan, R. Roufoogaran, H. Kim, P. Lettieri, B. Ibrahim, J. Rael, L. Tran, E. Geronaga, H. Yeh, T. Frost, J. Trachewsky, A. Rotougaran, A fully integrated SoC for 802.11b in 0.18 μm CMOS, in *Proceeding of the IEEE ISSCC Digest Technical Papers* (February, 2005), pp. 96–586
16. M. Demirkan, R.R. Spencer, Antenna characterization method for front-end design of pulse-based ultrawideband transceivers. *IEEE Trans. Antennas Propag.* **55**, 2888–2899 (2007)
17. M. Demirkan, R.R. Spencer, A 1.8Gpulse/s UWB transmitter in 90nm CMOS, in *IEEE ISSCC Digest Technical Papers* (February, 2008), pp. 116–117
18. A. Emira, A. Valdes-Garcia, B. Xia, A. Mohieldin, A. Valero-Lopez, S. Moon, C. Xin, E. Sanchez-Sinencio, A dual-mode 802.11b/Bluetooth receiver in 0.25 μm BiCMOS, in *Proceeding of the IEEE ISSCC Digest Technical Papers* (February, 2004), pp. 270–527
19. FCC, First report and order, FCC 02–48 (2002)
20. FCC, Second report and order and second memorandum opinion and order, FCC 04–285 (2004)
21. B.P. Ginsburg, A.P. Chandrakasan, Highly interleaved 5b 250MS/s ADC with redundant channels in 65nm CMOS, in *IEEE ISSCC Digest Technical Papers* (February, 2008), pp. 240–241
22. IEEE 802.15.4a Wireless MAC and PHY specifications for LR-WPANs (2007) [Online]. Available <http://www.ieee802.org/15/pub/TG4a.html>
23. H. Kim, D. Park, Y. Joo, All-digital low-power CMOS pulse generator for UWB system. *IEE Electron. Lett.* **40**(24), 1534–1535 (2004)
24. V. Kulkarni, M. Muqsith, H. Ishikuro, T. Kuroda, A 750Mb/s 12pJ/b 6-to-10GHz digital UWB transmitter, in *Proceeding of IEEE Custom Integrated Circuits Conference* (September, 2007), pp. 647–650
25. F.S. Lee, A.P. Chandrakasan, A 2.5nJ/b 0.65V 3-to-5GHz subbanded UWB receiver in 90nm CMOS, in *IEEE ISSCC Digest of Technical Papers* (February, 2007), pp. 116–117
26. J.-Y. Lee, R. Scholtz, Ranging in a dense multipath environment using an UWB radio link. *IEEE J. Sel. Areas Commun.* **20**, 1677–1683 (2007)
27. D. Marchaland, F. Badets, M. Villegas, D. Belot, 65nm CMOS burst generator for ultrawideband low data rate systems, in *Proceedings of IEEE Radio Frequency Integrated Circuits Symposium* (June, 2007), pp. 43–46
28. B. Marholev, M. Pan, E. Chien, L. Zhang, R. Roufoogaran, S. Wu, I. Bhatti, T.-H. Lin, M. Kappes, S. Khorram, S. Anand, A. Zolfaghari, J. Castaneda, C. Chien, B. Ibrahim, H. Jensen, H. Kim, P. Lettieri, S. Mak, J. Lin, Y. Wong, R. Lee, M. Syed, M. Rofougaran, A. Rofougaran, A single-chip Bluetooth EDR device in 0.13 μm CMOS, in *Proceeding of the IEEE ISSCC Digest Technical Papers* (February, 2007), pp. 558–759
29. P.P. Mercier, An all-digital transmitter for pulsed ultra-wideband communication, Master's thesis, Massachusetts Institute of Technology, Cambridge, MA, 2008
30. P.P. Mercier, D.C. Daly, M. Bhardwaj, D.D. Wentzloff, F.S. Lee, A.P. Chandrakasan, Ultra-low-power UWB for sensor network applications, in *Proceedings of IEEE ISCAS* (2008), pp. 2562–2565
31. P.P. Mercier, D.C. Daly, A.P. Chandrakasan, An energy-efficient all-digital UWB transmitter employing dual capacitively-coupled pulse-shaping drivers. *IEEE J. Solid-State Circuits* **44**, 1679–1688 (2009)
32. P.P. Mercier, M. Bhardwaj, D.C. Daly, A.P. Chandrakasan, A 0.55V 16Mb/s 1.6mW non-coherent IR-UWB digital baseband with $\pm 1\text{ns}$ synchronization accuracy, in *IEEE ISSCC Digest of Technical Papers* (February, 2009), pp. 252–253
33. P.P. Mercier, M. Bhardwaj, D.C. Daly, A.P. Chandrakasan, A low-voltage energy-sampling IR-UWB digital baseband employing quadratic correlation. *IEEE J. Solid-State Circuits* **45**, 1209–1219 (2010)

34. T. Norimatsu, R. Fujiwara, M. Kokubo, M. Miyazaki, A. Maeki, Y. Ogata, S. Kobayashi, N. Koshizuka, K. Sakamura, A UWB-IR transmitter with digitally controlled pulse generator. *IEEE J. Solid-State Circuits* **42**, 1300–1309 (2007)
35. A. Oncu, B.B.M.W. Badalawa, M. Fujishima, 22–9 GHz ultra-wideband CMOS pulse generator for short-range radar applications. *IEEE J. Solid-State Circuits* **42**, 1464–1471 (2007)
36. B. Otis, Y. Chee, J. Rabaey, A 400 μ W-RX, 1.6mW-TX super-regenerative transceiver for wireless sensor networks, in *Proceeding of the IEEE ISSCC Digest Technical Papers* (February, 2005), pp. 396–606
37. T.-A. Phan, V. Krizhanovskii, S.-G. Lee, Low-power CMOS energy detection transceiver for UWB impulse radio system, in *Proceedings of IEEE Custom Integrated Circuits Conference* (September, 2007), pp. 675–678
38. N. Pletcher, S. Gambini, J. Rabaey, A 2GHz 52 μ W wake-up receiver with -72 dBm sensitivity using uncertain-IF architecture, in *Proceeding of the IEEE ISSCC Digest Technical Papers* (February, 2008), pp. 524–633
39. A.-S. Porret, T. Melly, D. Python, C. Enz, E. Vittoz, An ultralow-power UHF transceiver integrated in a standard digital CMOS process: architecture and receiver. *IEEE J. Solid-State Circuits* **36**(3), 452–466 (2001)
40. J.G. Proakis, *Digital Communications*, 4th edn. (McGraw-Hill, New York, 2001)
41. J.M. Rabaey, M.J. Ammer, J.L. da Silva Jr., D. Patel, S. Roundy, PicoRadio supports ad hoc ultra-low power wireless networking. *Computer* **33**, 42–48 (2000)
42. E. Ragonese, A. Scuderi, V. Giannello, E. Messina, G. Palmisano, A fully integrated 24GHz UWB radar sensor for automotive applications, in *IEEE ISSCC Digest of Technical Papers* (February, 2009), pp. 306–307,307a
43. G. Retz, H. Shanan, K. Mulvaney, S. O'Mahony, M. Chanca, P. Corowley, C. Billon, K. Khan, P. Quinlan, A highly integrated low-power 2.4GHz transceiver using a direct-conversion diversity receiver in 0.18 μ m CMOS for IEEE802.15.4 WPAN, Feb 2009, pp. 414–415
44. J. Ryckaert, G. Van der Plas, V. De Heyn, C. Desset, B. Van Poucke, J. Craninckx, A 0.65-to-1.4 nJ/burst 3-to-10 GHz UWB all-digital TX in 90 nm CMOS for IEEE 802.15.4a. *IEEE J. Solid-State Circuits* **42**, 2860–2869 (2007)
45. L. Smaini, C. Tinella, D. Helal, C. Stoecklin, L. Chabert, C. Devaucelle, R. Cattenoz, N. Rinaldi, D. Belot, Single-chip CMOS pulse generator for UWB systems. *IEEE J. Solid-State Circuits* **41**, 1551–1561 (2006)
46. T. Terada, S. Yoshizumi, Y. Sanada, T. Kuroda, Transceiver circuits for pulse-based ultra-wideband, in *Proceeding of IEEE ISCAS* (May, 2004), pp. 349–352
47. T. Terada, S. Yoshizumi, M. Muqith, Y. Sanada, T. Kuroda, A CMOS ultra-wideband impulse radio transceiver for 1Mb/s data communications and ± 2.5 cm range findings. *IEEE J. Solid-State Circuits* **41**, 891–898 (2006)
48. P. Thoppay, C. Dehollain, M. Declercq, A 7.5mA 500 MHz UWB receiver based on super-regenerative principle, in *Proceeding of the IEEE European Solid-State Circuits Conference* (September, 2008), pp. 382–385
49. M. Verhelst, N. Van Helleputte, G. Gielen, W. Dehaene, A reconfigurable, 0.13 μ m CMOS 110pJ/pulse, fully integrated IR-UWB receiver for communication and sub-cm ranging, in *IEEE ISSCC Digest of Technical Papers* (February, 2009), pp. 250–251
50. D. Weber, W. Si, S. Abdollahi-Alibeik, M. Lee, R. Chang, H. Dogan, S. Luschas, P. Husted, A single-chip CMOS radio SoC for v2.1 Bluetooth applications, in *Proceeding of the IEEE ISSCC Digest Technical Papers* (February, 2008), pp. 364–620
51. M. Weisenhom, W. Hirt, Robust noncoherent receiver exploiting UWB channel properties, in *Proceedings of IEEE Joint UWBST & IWUWBS* (May, 2004), pp. 156–160
52. D.D. Wentzloff, Pulse-based ultra-wideband transmitters for digital communication, Ph.D. thesis, Massachusetts Institute of Technology, 2007
53. D.D. Wentzloff, A.P. Chandrakasan, Gaussian pulse generators for subbanded ultra-wideband transmitters. *IEEE Trans. Microw. Theory Tech.* **54**, 1647–1655 (2006)
54. D.D. Wentzloff, A.P. Chandrakasan, A 47pJ/pulse 3.1-to-5GHz all-digital UWB transmitter in 90nm CMOS, in *IEEE ISSCC Digest of Technical Papers* (February, 2007), pp. 118–119

55. D.D. Wentzloff, F.S. Lee, D.C. Daly, M. Bhardwaj, P. Mercier, A.P. Chandrakasan, Energy efficient pulsed-UWB CMOS circuits and systems, in *Proceedings of IEEE International Conference on Ultra-Wideband* (September, 2007), pp. 282–287
56. M.Z. Win, R.A. Scholtz, Impulse radio: how it works. *IEEE Commun. Lett.* **2**, 36–38 (1998)
57. J. Yuan, C. Svensson, High-speed CMOS circuit techniques. *IEEE J. Solid-State Circuits* **24**, 62–70 (1989)
58. J. Yuan, C. Svensson, Fast CMOS nonbinary divider and counter. *IEE Electron. Lett.* **29**, 1222–1223 (1993)
59. Y. Zheng, Y. Tong, C.W. Ang, Y.-P. Xu, W.G. Yeoh, F. Lin, R. Singh, A CMOS carrier-less UWB transceiver for WPAN applications, in *IEEE ISSCC Digest of Technical Papers* (February, 2006), pp. 378–387
60. Y. Zheng, K.-W. Wong, M. Annamalai Asaru, D. Shen, W.H. Zhao, Y.J. The, P. Andrew, F. Lin, W.G. Yeoh, R. Singh, A 0.18 μm CMOS dual-band UWB transceiver, in *IEEE ISSCC Digest of Technical Papers* (February, 2007), pp. 114–115
61. Y. Zheng, M.A. Arasu, K.-W. Wong, Y.J. The, A.P.H. Suan, D.D. Tran, W.G. Yeoh, D.-L. Kwong, A 0.18 μm CMOS 802.15.4a UWB transceiver for communication and localization, in *IEEE ISSCC Digest of Technical Papers* (February, 2008), pp. 118–119
62. Y. Zhu, J.D. Zuegel, J.R. Marciante, H. Wu, A 10 GS/s distributed waveform generator for sub-nanosecond pulse generation and modulation in 0.18 μm standard digital CMOS, in *Proceedings of IEEE Radio Frequency Integrated Circuits Symposium* (June, 2007), pp. 35–38
63. D. Zito, D. Pepe, M. Mincica, F. Zito, A 90nm cmos soc uwb pulse radar for respiratory rate monitoring, in *IEEE ISSCC Digest of Technical Papers* (February, 2011), pp. 40–41