

A Low-Voltage Energy-Sampling IR-UWB Digital Baseband Employing Quadratic Correlation

Patrick P. Mercier, *Student Member, IEEE*, Manish Bhardwaj, *Member, IEEE*, Denis C. Daly, *Member, IEEE*, and Anantha P. Chandrakasan, *Fellow, IEEE*

Abstract—This paper describes a digital baseband designed for use in a non-coherent IR-UWB system. Owing to the nonlinear statistics introduced by the energy-sampling RF front-end, the baseband employs a new quadratic correlation technique that achieves comparable performance to a matched filter classifier, with the added benefit of being robust to SNR estimation errors. Additionally, “alias-free codes” are introduced that allow for pulse-level synchronization accuracy without requiring any increase in front-end complexity. Fabricated in a 90 nm CMOS process, the digital baseband utilizes significant parallelism in addition to clock and data gating to achieve low-power operation, with supply voltages as low as 0.55 V. At a clock frequency of 32 MHz, the baseband requires 14-to-79 μ s to process a preamble during which it consumes an average power of 1.6 mW, while payload demodulation requires 12 pJ/bit.

Index Terms—Correlation, energy detection, maximum likelihood, modem, non-coherent, ultra-wideband, UWB, VLSI.

I. INTRODUCTION

BATTERY-POWERED applications such as medical monitoring and personal telemetry systems often require sophisticated sensing and communication functionality in small, unobtrusive packages. These types of applications share many attributes and requirements. For example, personal telemetry-type systems typically require long battery lifetimes, operate over short distances, and produce data rather infrequently. This implies that resulting system implementations should be energy efficient, and communicate with low average data rates and short payloads.

Impulse-radio ultra-wideband (IR-UWB) systems are finding increasing use in such applications, in part due to their Federal Communications Commission (FCC)-regulated short range [1] and inherently duty-cycled nature [2]. In particular, non-coherent IR-UWB radios, which discard phase information and instead use amplitude or position modulation schemes, simplify RF front-end complexity and thus lower the amount of energy

required to transmit [3]–[6] and/or receive [7]–[10] a single bit of information. Furthermore, due to relaxed frequency tolerances of wideband, non-coherent communication, RF phase-locked loop (PLLs) are not necessarily needed, and the RF front-end circuits may be aggressively duty-cycled between pulses and/or packets, thereby dramatically lowering the average power consumption.

Depending on an application’s data rate and quality of service (QoS) needs, it is often more convenient in IR-UWB systems to relay bursts of information at high instantaneous data rates, rather than continuously relaying information at low data rates. There are three main reasons why bursting a packet of information is beneficial: 1) although IR-UWB circuits can generally duty cycle very quickly, there is less energy spent turning on and off bias currents and/or calibrating delay-lines or digital-to-analog converter (DACs) if duty cycling is at the packet-level rather than at the bit-level; 2) external characteristics such as temperature and multipath objects are unlikely to significantly change during the course of a short packet and thus do not need to be tracked; and 3) clock offset-tracking circuitry may not be required if the packet is sufficiently short and an appropriate modulation scheme is used.

Unfortunately, since RF front-end circuits are typically on and consuming energy during all phases of a packet, a larger proportion of system energy is spent on detection and synchronization tasks when payloads become short. For instance, over 50% of the energy consumed when transmitting a 100-bit packet using Zigbee is spent on the preamble and header, where no useful information is conveyed [11]. As a result, it is not sufficient to compare wireless systems solely by the energy consumed per raw bit. Instead, a more useful metric is the energy per *useful* bit (EPUB), which accounts for time and energy spent during preambles and headers [12].

By discarding phase information and accumulating squared noise, non-coherent square-and-integrate energy samplers offer inherently less performance than their coherent counterparts, typically by as much as 5–10 dB in signal-to-noise ratio (SNR) for the same bit error rate (BER) [13]. However, due to the increased front-end energy efficiency and eased synchronization accuracy requirements, such energy-sampling systems can offer superior EPUB to coherent systems when operating with very short payloads [14].

The main goal of this work is to design an energy-sampling digital baseband modem that maximizes battery lifetimes by reducing the amount of time an RF front-end must be turned on during a preamble. This goal coincides with designing a system

Manuscript received August 12, 2009; revised December 11, 2009; accepted February 25, 2010. Current version published June 09, 2010. This paper was approved by Associate Editor Stefan Rusu. This work was supported in part by DARPA HI-MEMS under Contract FA8650-07-C-7704 and the Natural Sciences and Engineering Council of Canada (NSERC). Chip fabrication was provided by STMicroelectronics.

P. P. Mercier and A. P. Chandrakasan are with the Department of Electrical Engineering and Computer, Science, Massachusetts Institute of Technology (MIT), Cambridge, MA 02139 USA (e-mail: pmercier@mit.edu).

M. Bhardwaj is with Innovators In Health, Cambridge, MA 02138 USA.

D. C. Daly is with Cambridge Analog Technologies, Bedford, MA 01730 USA.

Digital Object Identifier 10.1109/JSSC.2010.2046245

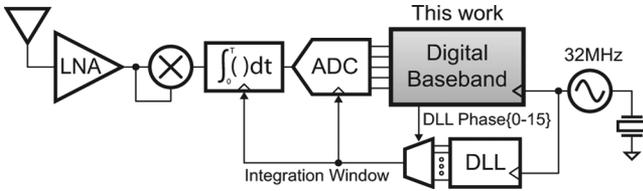


Fig. 1. Top-level energy-sampling receiver block diagram.

with minimum EPUB. Additionally, the baseband should support improved synchronization performance over traditional techniques without requiring additional analog processing, be integrated with a radio front-end in a deep sub-micron CMOS process, and consume minimal energy itself.

This paper is organized as follows. Section II describes the non-coherent receiver system architecture and packet structure. Section III describes previously published synchronization approaches, and presents a new class of codes designed for energy-sampling systems. Section IV introduces the proposed quadratic correlation scheme for improved synchronization performance. Section V presents the digital baseband architecture and discusses circuit implementations. Section VI compares measured chip results with results from other classes of receivers. Finally, Section VII concludes the paper.

II. SYSTEM ARCHITECTURE

The chosen receiver architecture employs a non-coherent energy-sampling front-end, as shown in Fig. 1 [15]. The receiver operates in the 3.5, 4.0, and 4.5 GHz bands set by the 802.15.4a standard [16]. Following initial amplification by a tunable low noise amplifier (LNA), energy is measured by self-mixing (i.e., squaring) the signal to baseband, and integrating the result over a 31.25 ns window. This window length is chosen such that up to 16 back-to-back 1.95 ns pulses (or, more generally, chip periods) fit in a single window. The integrated result is then quantized to 5 bits and fed to the baseband in parallel at 32 MHz.

Since the front-end discards phase information and samples over 31.25 ns windows of time, a high speed clock is not required; all circuits are clocked off of a free-running 32 MHz crystal oscillator.¹ In fact, the large integration window, coupled with a non-coherent modulation scheme, relaxes clock frequency and phase requirements to the point where no PLLs or clock-offset tracking circuits are necessary for payloads under 6000 bits (given a crystal accuracy of 40 ppm).

A packet structure is shown in Fig. 2. The preamble is modulated using on-off keying (OOK) and consists of the following phases:

- **Detection** involves determining if an appropriate packet is being received by the RF front-end. The receiver averages a programmable number of codewords, S_0 , followed by five codeword cycles spent computing correlation results. If it is deemed that no packet is being received, the baseband goes into idle-mode for a programmable length of time before re-entering the detection phase.

¹With the exception of the integrator and analog-to-digital converter (ADC), which take a particular output phase from the integrated delay-locked loop (DLL) only after synchronization has been achieved.

- **Synchronization** involves determining where the beginning of the current codeword S_0 is in relation to the current energy sample. As in the detection phase, the receiver averages a programmable number of codewords, followed by five codeword cycles used for correlation. The integration window phase is shifted to the inferred location immediately following synchronization.
- **Start frame delimiter (SFD) search** involves determining where the preamble ends and where the payload begins. This is accomplished by correlating to a sequence of 1 or 5 codewords.

The payload is modulated using pulse-position modulation (PPM) and consists of a header and actual data. Following the payload, the system enters a low-power idle mode where all circuits, with the exception of the crystal oscillator and a watchdog timer, are clock gated (digital circuits) or bias-current gated (analog circuits) until the next packet is expected to arrive.

III. SYNCHRONIZATION APPROACHES

In an energy-sampling system, non-reliance on phase simplifies front-end circuits, but may significantly increase synchronization complexity. This is particularly true in positioning applications, where synchronization accuracy specifications are much more stringent than those normally required by a PPM-modulated payload.

Since the preamble consists of repetitions of the synchronization sequence (codeword), synchronization is the problem of inferring the chip-level cyclic delay suffered by the sequence. Physically, this can be broken up into two alignments: inferring the cyclic shift of the integrated codeword (codeword alignment), and inferring the integration window location (phase alignment). Phase alignment is illustrated in Fig. 3 with a simple example of a 5-sample code, where the integration window is three chip periods long.

In the literature, non-coherent IR-UWB synchronization codes can be split up into two main categories: repetition codes and pseudo-random codes. What follows is a brief description of synchronization approaches using these two classes of codes.

A. Repetition Codes

The simplest synchronization approach involves transmitting repetitions of the chip-level code,

$$c = [1 \ 1 \ 1 \ \dots \ 1 \ 0 \ 0 \ 0 \ \dots \ 0],$$

where a '1' represents a single pulse, and a '0' represents a chip period of zero output. Typically the number of '1's and '0's are equal, and are set to exactly fill an entire integration window [17]. The receiver then collects energy in multiple windows that are relatively shifted in time. This may be done by either having multiple parallel integrators [7], or by sliding the integration window [18]. To increase reliability, multiple energy measurements are made in each window and averaged. The window with the maximum measured energy is inferred to be the optimum.

While conceptually simple, repetition codes suffer from poor performance and typically require significant averaging in low-SNR environments before reliable results are achieved.

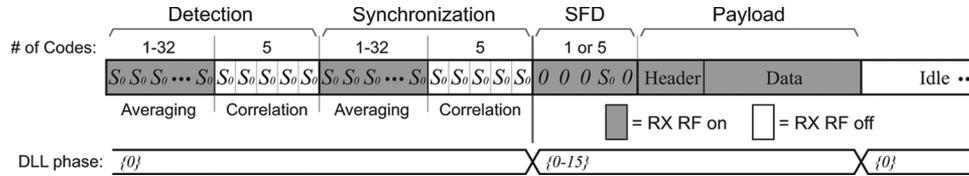


Fig. 2. Packet structure used by the digital baseband.

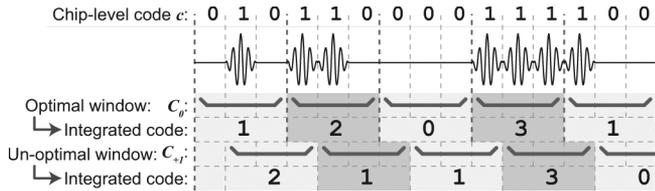


Fig. 3. Illustration of phase alignment required during synchronization with an integration window of three chip periods. Codeword alignment involves determining the correct cyclic shift of the resulting integrated codewords.

B. Pseudo-Random Codes

In an energy-sampling system, every chip-level synchronization code has an integrated counterpart, obtained by passing the sequence through an energy-sampler (squared-integrator). These codes can be classified as either being *alias-free* or aliased. If two distinct chip-level shifts of a sequence yield identical integrated sequences, these shifts are deemed to be aliased. If no such shifts exist, then the sequence is deemed to be *alias-free*.

1) *802.15.4a Codes*: Synchronization in the 802.15.4a standard requires the use of length 31 or 127 codes drawn from a ternary alphabet $\{-1, 0, +1\}$ [16]. Each non-zero ternary-symbol² represents a single pulse followed by either 3, 15, or 63 chip periods of zero output, depending on the channel number and system configuration. The purpose of transmitting a single pulse followed by several chip periods of zero output is to help resolve multipath.

These 15.4a codes are aliased if the integration period is greater than the chip period, which is generally true for low-power energy-samplers. A disadvantage of aliased codes is that achieving chip-level timing resolution requires the integration window to be shifted in time by a single chip period, thus increasing the complexity of the receiver. Also, we have recently shown that 15.4a codes are indeed optimum in coherent systems, but not when energy-sampling is used [19]. In other words, 15.4a codes do not necessarily minimize the preamble length required to achieve a specified probability of synchronization error in energy-sampled systems.

2) *Alias-Free Codes*: We have previously shown that the optimum synchronization sequence in an energy-sampling system depends on the underlying SNR, and have proposed alias-free sequence families for both the low and high SNR regimes [19]. To illustrate alias-free codes, consider an integration window of two chip periods (i.e., pulse periods), and the chip-level code

$$\mathbf{c} = [1 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0].$$

²For the purposes of non-coherent systems, the polarity of the pulse is unimportant.

If this code is integrated two chips at a time, it yields the integrated result³

$$\mathbf{C}_0 = [1 \ 1 \ 0 \ 0 \ 0].$$

A single right cyclic shift of \mathbf{c} when integrated yields

$$\mathbf{C}_{+1} = [1 \ 0 \ 1 \ 0 \ 0].$$

It may be verified that no two shifts of this code yield the same integrated sequence. Hence, this is an alias-free sequence.

Alias-free codes can be constructed for any integration window length, and can achieve chip-level synchronization resolution *regardless* of sampling resolution. In other words, alias-free codes do not need fractional shifts or accurate, high speed clocks, thus simplifying receiver hardware. These factors motivate the choice of alias-free codes in the proposed digital baseband.

Sparse alias-free codes can be constructed that are very similar to 802.15.4a codes and thus share common characteristics (such as performance in multipath environments). For example, an alias-free code can be constructed by shifting each pulse in a 802.15.4a code by its index value (a “walking code”). This does not affect peak or average radiated power limits and can be easily implemented, since many IR-UWB transmitters inherently have sufficient timing resolution.

IV. QUADRATIC CORRELATION

A matched filter (or correlator) is the optimum synchronizer in coherent systems, i.e., it minimizes the probability of error. Energy-sampling systems must also “correlate” the received signal with all possible cyclic shifts of the transmitted sequence, but the optimum correlation function is nonlinear, owing to the squaring operation in the energy sampler. The optimum function may be obtained via the maximum-likelihood (ML) expression and involves Bessel functions, making it computationally cumbersome. In contrast, prior work involving matched filters work well in ideal conditions, but may be very sensitive to SNR estimation errors for certain sequences.

When the period of integration is large compared with the chip period (which is desirable for low-power systems), the central limit theorem permits a Gaussian approximation, and yields a new and robust classifier: the quadratic correlator (QCORR). Hence, the correlation measure is of the form

$$P_j = \sum_k \frac{(y_k - \mu_{kj})^2}{\sigma_{kj}^2} \quad (1)$$

³The notation here is that a lower-case \mathbf{c} refers to a chip-level codeword, while an upper-case \mathbf{C}_x refers to an integrated codeword with the shifted chip-level index, x .

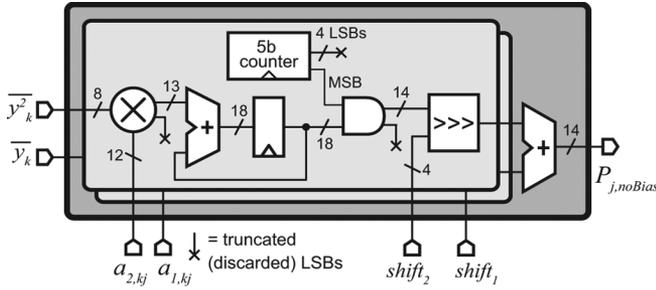


Fig. 4. Circuit diagram of the proposed quadratic correlator (QCORR). Truncation and arithmetic shifts are used to preserve dynamic range.

where μ_{kj} and σ_{kj}^2 are the expected signal mean and variance for the k 'th coordinate (i.e., sample) of the j 'th codeword shift.

Sequences are typically repeated to improve synchronization performance at low SNRs. It is more efficient to average the signal and then correlate, rather than computing correlations for each repetition and then averaging correlator outputs. The quadratic form implies that *two* averages are required: the signal ($\overline{y_k}$), and the signal squared ($\overline{y_k^2}$). Equation (1) can then be expanded for ease of implementation into the following form:

$$P_j = \sum_k \left(a_{2,kj} \overline{y_k^2} + a_{1,kj} \overline{y_k} \right) + b_j \quad (2)$$

where $a_{2,kj} = 1/\sigma_{kj}^2$ and $a_{1,kj} = -2\mu_{kj}/\sigma_{kj}^2$ are multiplier coefficients and $b_j = \sum \mu_{kj}^2/\sigma_{kj}^2$ is a bias term that can be added after the main summation. Fig. 4 illustrates the implemented form of the quadratic correlator.

The proposed quadratic correlation approach has a performance that is nearly indistinguishable from the ML approach at all SNRs (provided, of course, that the integration period is large compared with the chip period), and, as shown in Section VI, it is much more robust to SNR mismatches than a matched filter.

V. DIGITAL BASEBAND ARCHITECTURE

A simplified block diagram of the digital baseband is shown in Fig. 5 [20]. A highly-parallel, low fan-in architecture allows for low voltage operation without sacrificing throughput. Additionally, since storage requirements are relatively small, all memories use flip-flop-based registers (which can also operate down to low voltages). Fine-grained clock gating, where the clock nets to individual blocks are explicitly gated when they are not in use, is implemented to reduce dynamic power dissipation in the clock tree network. Since there are many high fan-out nets in the design that would normally be broadcasted to many different blocks, input vector gating techniques, where the state of parallel nets are driven to constant static values, are implemented to dynamically prevent broadcasting vectors to blocks that are not being used. This further reduces unnecessary dynamic power by limiting bus activity factors.

Baseband circuits corresponding to different phases of a packet are described below. Since the bulk of the circuitry is devoted to synchronization, it is described first, even though detection is always performed prior to synchronization.

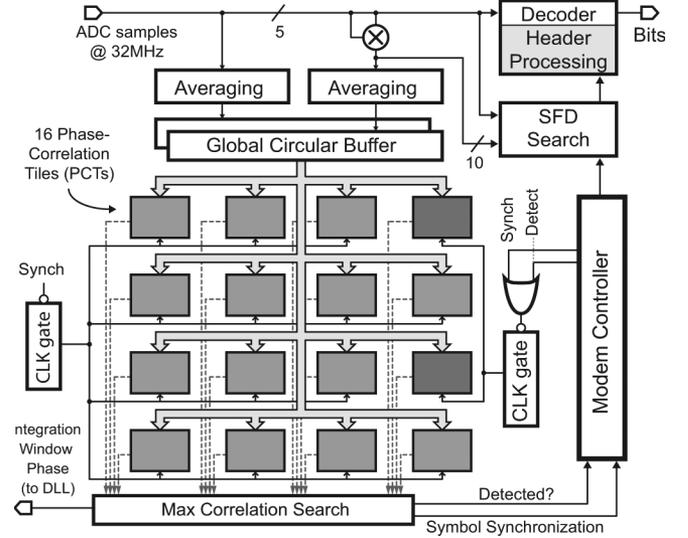


Fig. 5. Top-level block diagram of digital baseband modem.

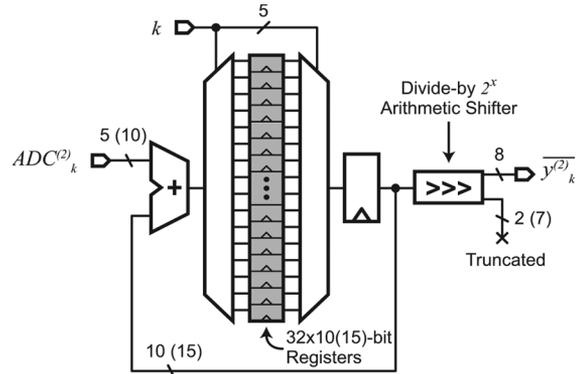


Fig. 6. Circuit that separates and averages linear ADC samples into 32-sample codewords. The circuit that separates the squared (quadratic) samples is the same, except with modified bit-widths (shown in parentheses).

A. Synchronization

During detection and synchronization, the transmitter sends several repetitions of a 32-sample (512 chips) alias-free integrated codeword, S_0 . The receiver front-end squares, integrates, and digitizes the received signal. A 5-bit, 32 MS/s signal is generated, which is passed to the digital baseband. These samples are again squared, then both the 5-bit linear and 10-bit quadratic samples are separated into 32-sample codewords and averaged (Fig. 6). Up to 32 codewords (restricted to powers of two) can be averaged for both detection and synchronization. This amount of averaging can achieve ± 1 ns synchronization accuracy at reasonably low SNRs while tolerating clock offset effects for crystal accuracies of up to 40 ppm.

Once averaging is complete, the receiver front-end is temporarily shut off to save energy (by gating bias currents), and the averaged linear and quadratic samples are distributed to the detection and synchronization QCORR network via a global circular buffer (which is shown in greater detail in Fig. 7). Recall that the goal of synchronization is to determine both the number of sample shifts required to achieve codeword alignment, and the phase shift within a sample required to align the integration

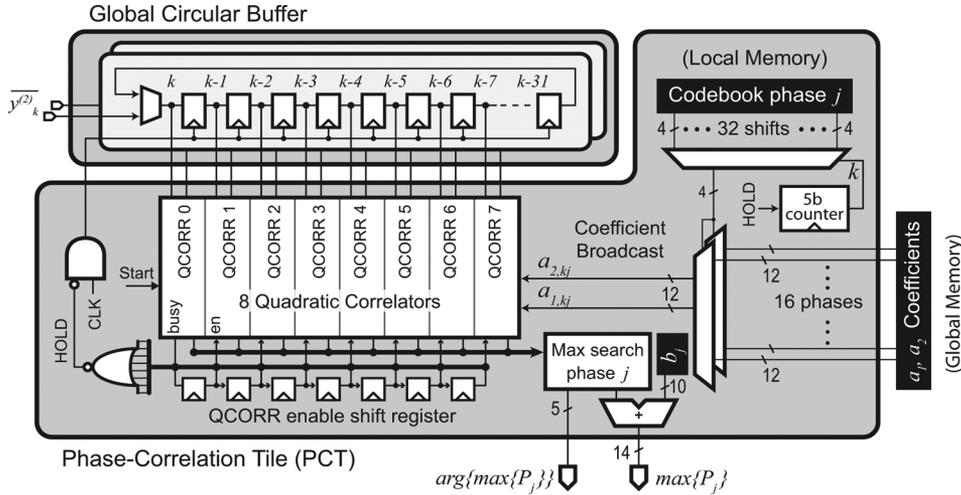


Fig. 7. Block diagram of a phase correlation tile (PCT) that computes all 32 quadratic correlations of a particular integration window phase. Also shown is the global circular buffer that feeds all 16 PCTs on the chip.

window. This is achieved by correlating the received samples with digitally-stored coefficients corresponding to every possible shift of the original codeword. The shift that results in the maximum correlation is deemed to be the most likely possibility, and thus is selected for synchronization purposes.

With an integration window of 31.25 ns, every sample duration corresponds to 16 possible signal starts, or phases, each separated by a single pulsewidth of 1.95 ns. Since the length of the integrated codeword S_0 is 32 samples, a total of $16 \times 32 = 512$ quadratic correlations are required (i.e., the chip-level codeword, s , is of length 512) to attain chip-level synchronization accuracy. These correlations are distributed over 16 phase correlation tile (PCTs) (Fig. 7), which each compute all 32 correlations for a particular integration window phase.

A quadratic correlation involves multiplying averaged 32-sample linear ($\overline{y_k}$) and quadratic ($\overline{y_k^2}$) input data with a group of 32 pairs of coefficients representing one of the 512 distinct shifts of the original codeword (as shown in (2)). To attain desired performance, each coefficient requires 12 bits. Without any optimization, this would involve storing $512 \times 32 \times 2 \times 12 = 393,216$ bits of multiplier coefficients in memory. However, there are only a total of 16 distinct pairs of linear and quadratic multiplier coefficients—one pair for each possible combination of the number of pulses in an integration window (0-to-15).⁴ Thus, it is possible to devote only 384 bits to the distinct coefficients, and instead allocate bits to store sequences of 4-bit pointers that point to the 16 pairs of individual coefficients. This pointer memory bank is called the integrated-codebook.⁵

Without optimization, the integrated-codebook would require $512 \times 32 \times 4 = 65,536$ bits of pointer information. However, since the integration window has only 16 phases, the integrated-codebook begins to repeat itself in a circularly-shifted

⁴The case of 16 input pulses is truncated to 15 such that a 4-bit representation may be used. This has little impact on performance.

⁵The chip-level codebook is of the size 512×512 , where all entries are binary, and each column represents a single chip period. The integrated-codebook is of size 512×32 , since each column contains a 4-bit entry describing the expected number of pulses in an integration window.

fashion after the first 16 pulse-width shifts. Thus, only $16 \times 32 \times 4 = 2048$ bits of memory are required to completely define the codebook.

Since each PCT operates on a single integration window phase, only one row of the integrated-codebook is required per PCT. This row, or integrated-codebook phase, is thus locally stored in a programmable shift-register-based memory.

To reduce the latency of the correlation computation, each PCT consists of 8 parallel QCORRs, thus requiring 4 sets of correlations spaced out in time in order to compute the 32 required correlations. Since all 8 QCORR units within a PCT perform correlations with an identical, but shifted set of coefficients, an efficient pipeline schedule can reduce hardware complexity. As shown in Fig. 7, a common pair of linear and quadratic coefficients are broadcast to all 8 QCORRs, while averaged data is offset between QCORRs using a global circular buffer. Rather than sharing coefficients among QCORRs, the alternative is to share input samples. However, this would increase multiplexing costs of non-local coefficient fetches by 8x. A circular-buffer schedule for the first set of 8 correlations is shown in Fig. 8.

Due to the QCORR schedule, correlation results are produced linearly in time. This allows for a simple one-stage comparator with memory to keep track of the local maximum correlation. After all 32 correlations have been computed per PCT, a phase-dependent bias, b_j , is added to the local maximum, which is then sent to the global maximum correlation search (Fig. 5), along with its sample-shift argument.

Immediately following synchronization calculations, the front-end is turned back on, and the LSBs of the argument of the global maximum correlation, which determine the inferred integration window phase, are relayed to the DLL. This occurs while the baseband is waiting to achieve codeword alignment (set by the MSBs of the global maximum correlation), allowing the RF and analog circuits sufficient time to settle.

B. Detection

The goal of the detection phase is to determine whether noise, or a properly-coded message, is being received. The de-

		Time (clock cycles) →																
QCORR	1	2	3	4	5	6	7	8	9	...	32	33	34	35	36	37	38	39
#0	a_0y_0	a_1y_1	a_2y_2	a_3y_3	a_4y_4	a_5y_5	a_6y_6	a_7y_7	a_8y_8		$a_{31}y_3$							
#1		a_1y_0	a_2y_1	a_3y_2	a_4y_3	a_5y_4	a_6y_5	a_7y_6	a_8y_7		$a_{31}y_3$	a_0y_{31}						
#2			a_2y_0	a_3y_1	a_4y_2	a_5y_3	a_6y_4	a_7y_5	a_8y_6		$a_{31}y_2$	a_0y_{30}	a_1y_{31}					
#3				a_3y_0	a_4y_1	a_5y_2	a_6y_3	a_7y_4	a_8y_5		$a_{31}y_2$	a_0y_{29}	a_1y_{30}	a_2y_{31}				
#4					a_4y_0	a_5y_1	a_6y_2	a_7y_3	a_8y_4	...	$a_{31}y_2$	a_0y_{28}	a_1y_{29}	a_2y_{30}	a_3y_{31}			
#5						a_5y_0	a_6y_1	a_7y_2	a_8y_3		$a_{31}y_2$	a_0y_{27}	a_1y_{28}	a_2y_{29}	a_3y_{30}	a_4y_{31}		
#6							a_6y_0	a_7y_1	a_8y_2		$a_{31}y_2$	a_0y_{26}	a_1y_{27}	a_2y_{28}	a_3y_{29}	a_4y_{30}	a_5y_{31}	
#7								a_7y_0	a_8y_1		$a_{31}y_2$	a_0y_{25}	a_1y_{26}	a_2y_{27}	a_3y_{28}	a_4y_{29}	a_5y_{30}	a_6y_{31}
y_k (#0 QCORR)	y_0	y_1	y_2	y_3	y_4	y_5	y_6	y_7	y_8	...	y_{31}	y_0	y_1	y_2	y_3	y_4	y_5	y_6
a_k (All QCORRs)	a_0	a_1	a_2	a_3	a_4	a_5	a_6	a_7	a_8	...	a_{31}	a_0	a_1	a_2	a_3	a_4	a_5	a_6

Fig. 8. Schedule of eight parallel QCORR units showing data and coefficient choices for each one. Multiplier coefficients (a_k 's) are broadcast to all 8 QCORRs, while averaged input samples (y_k 's) are offset in time via the circular buffer. A blank table entry signifies the given QCORR is not operational (i.e., clock gated). Note that the \bar{y} and $a_{1/2,j}$ notations are dropped to simplify the table (only k 's are shown).

tection operation is nearly identical to synchronization: linear and quadratic input samples are first averaged, then correlated to a subset of the possible shifts of the expected codeword, S_0 . Additionally, an extra QCORR is used to correlate the input samples to an all-zero codeword corresponding to noise. If the maximum correlation to any shift of the expected codeword (S_0) exceeds the correlation to the noise codeword by a programmable threshold,⁶ then it is inferred that an appropriate signal is being received, and the receiver may enter the synchronization phase (recall, detection is always performed prior to synchronization).

Only 2 out of the 16 PCTs are used, since only a subset of all possible shifts are required for the desired detection performance. Thus, a fine-grained clock gating scheme is implemented to reduce clock tree power to the unused PCTs. Overall, the entire clock gating strategy reduces dynamic power by 2.7x during idle mode when only a watchdog timer is active.

Note that the detection phase can also play a role in multi-access differentiation—any orthogonal preamble code will appear as noise to the detection circuitry, thus preventing receivers from locking to unintended transmissions.

C. SFD Search

The goal of the start frame delimiter search phase is to determine exactly when the preamble ends and where the payload begins. This is achieved by searching for a new code that is sufficiently uncorrelated to the preamble code used during detection and synchronization. The task of searching for the start frame delimiter (SFD) is split into two sets of correlations:

- 1) Inner quadratic correlations of input samples with the underlying 32-sample codeword, S_0 , and an all-zero 32-sample codeword.
- 2) Outer linear correlations of inner results with all possible shifts of a length 5 SFD code sequence stored in a codebook (low SNR). In the high SNR regime, a length 1 SFD code (the all-zero codeword) is used instead.

The SFD is deemed to be found when the correlation with the unshifted outer sequence exceeds correlations to its possible

⁶The threshold determines the missed detection/false alarm performance trade-off. A typical setting would result in $P(\text{missed detection}) = 10^{-2}$ and $P(\text{false alarm}) = 10^{-6}$.

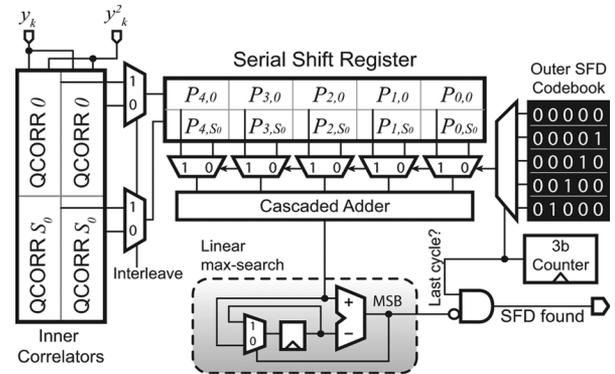


Fig. 9. Block diagram of the SFD search circuitry.

shifts. This inner/outer correlation scheme is similar to what is typically implemented in wireless systems that spread preamble bits over multiple chip periods (such as CDMA).

A block diagram of the circuit that performs computations during the SFD phase is shown in Fig. 9. The SFD codebook is programmable via a local shift-register memory, and contains shifts of the outer code ($\{0 S_0 0 0 0\}$ in this illustration). Note that only codeword shifts in the advancing direction are required, since the baseband either expects to find the SFD, or find no SFD at all.⁷

Recall that immediately after synchronization, the integration window is shifted to the inferred location, and the receiver waits the appropriate number of clock cycles to attain codeword alignment. Thus, the search for the SFD ideally begins exactly as a new codeword is being received. As a result, it is not necessary to perform quadratic correlations to any shifts of the code S_0 ; rather, it is only necessary to correlate to the unshifted code itself.⁸ Consequently, PCTs are not required during the SFD phase; all PCTs are thus clock gated. Instead, individual QCORRs may be used to correlate to the code S_0 and the all-zero codeword. Because it takes a finite number of clock cycles for a QCORR to begin processing data and

⁷In other words, the baseband does not expect to miss the SFD by a single inner-codeword, since immediately following the SFD is the payload, as opposed to a circular shift of the outer SFD code.

⁸Of course, it is never necessary to correlate to shifts of an all-zero codeword.

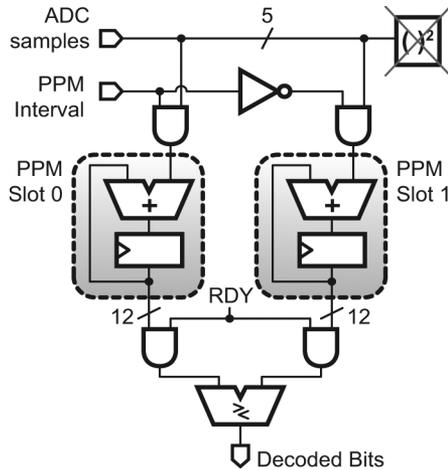


Fig. 10. Payload demodulator circuit that decodes variable-length PPM data.

finalize its computation, two pairs of QCORRs are used in a time-interleaved fashion to compute the required quadratic correlations without dropping input samples or requiring extensive buffering.

The inner (quadratic) correlation results from the interleaved QCORRs fill a pair of 5-entry serial shift registers. After 5 sets of correlations, the shift register is full, and outer (linear) correlations may be computed. Outer correlations are computed in a serial fashion by rotating through the SFD codebook and adding the 5 linear results dictated by the current line in the SFD codebook. A linear max-search circuit is used to keep track of the maximum outer correlation. If the maximum outer correlation value corresponds to the last outer correlation (i.e., the unshifted SFD), the SFD is deemed to be found and the baseband may continue into the payload demodulation phase. If the last outer correlation is not the maximum result, the SFD codebook counter is reset and outer correlations are halted until a new inner result is available. This process repeats until the SFD is found, or a programmable timeout is reached.

D. Payload Demodulator

The payload is found immediately after the SFD, and in this implementation, contains an 8-bit header followed by useful payload information, both of which are modulated using PPM. Since it takes several clock cycles to compute outer SFD correlations and determine if the SFD has been found, payload demodulation must occur at the same time that outer SFD correlations begin in order to avoid dropping data or increasing latency. If the SFD is deemed to have not been found, demodulated data is simply discarded and the process is reset.

Fig. 10 shows the circuit responsible for demodulating payload data. Two accumulators, which can each accept data at 32 MS/s, are used to sum energies in adjacent PPM time slots. The accumulator data paths are sized such that variable-length PPM data can be decoded, with PPM time slots ranging from 31.25 ns, which is the length of a single integration window, to 4 μ s, or 128 adjacent integration windows. As a result, the payload data rate ranges from 125 kbps (spreading factor = 128) to 16 Mbps (spreading factor = 1). Duty-cycling between PPM time slots is not permitted.

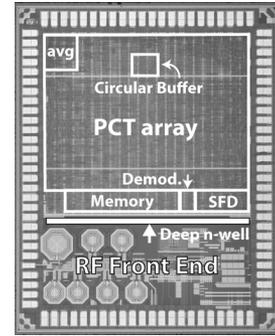
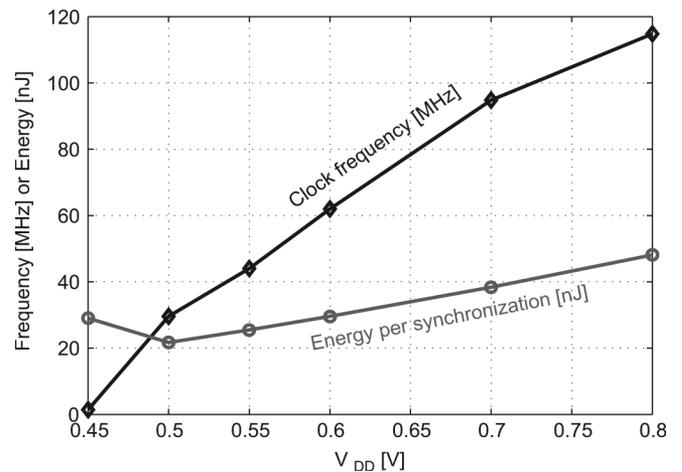


Fig. 11. Die photo of the digital baseband as part of a receiver SoC. A deep n-well isolates digital substrate noise from important analog blocks.

 TABLE I
SUMMARY OF SYSTEM PARAMETERS

Integration window	16 chip periods
ADC sample rate	32MS/s
ADC bit width	5 bits
s chip-level codeword length	512 chips
S_0 integrated codeword length	$512/16 = 32$ symbols
S_0 det./synch. repetitions	2^n , $n \in \{0, \dots, 5\}$
SFD outer code length	1 (high SNR) or 5 (low SNR)
Payload PPM window length	1-to-128 integration windows


 Fig. 12. Measured clock frequency [MHz] and energy per synchronization [nJ] plotted versus the supply voltage, V_{DD} [V].

VI. MEASUREMENT RESULTS

A. Circuits

The digital baseband was fabricated in a 90 nm CMOS process as part of an SoC with an integrated RF front-end [15]. The baseband occupies a core area of 1.5×1.7 mm²; a die photo is shown in Fig. 11. A summary of system parameters is shown in Table I.

The baseband is designed to operate down to low-voltages for improved energy efficiency per operation [21]. As shown in Fig. 12, the baseband is functional at supply voltages as low as 0.45 V. Lowering the supply voltage results in less switching energy and leakage power. However, leakage power begins to integrate over exponentially larger clock periods at lower voltages, thereby comprising an increasing

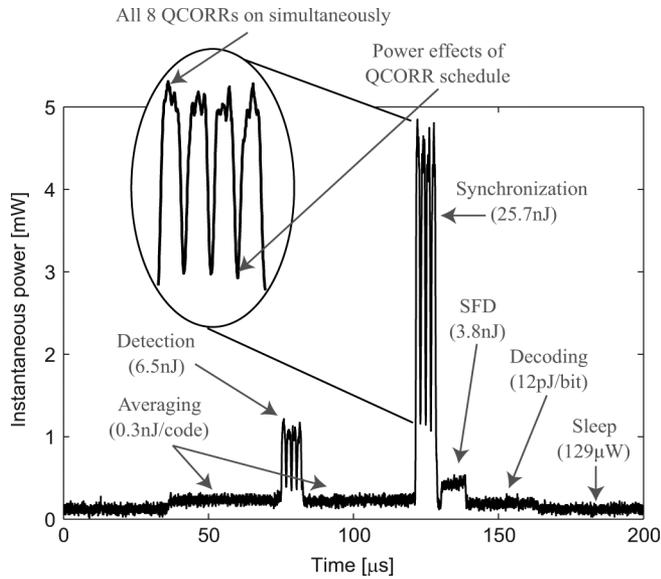


Fig. 13. Power profile of digital baseband operating at 32 MHz with $V_{DD} = 0.55$ V over the course of an entire packet. Note the four spikes in power consumption during synchronization, corresponding to the four sets of correlations that are computed by the PCT network.

portion of the total energy consumed per clock cycle. This switching/leakage energy trade-off results in a minimum energy point at $V_{DD} = 0.5$ V.

Fortunately, the supply voltage which corresponds to a maximum clock frequency of 32 MHz lies very near the minimum energy point. With some margin added for safety, the baseband operates reliably with $V_{DD} = 0.55$ V at a clock frequency of 32 MHz.

Fig. 13 shows the measured power profile of the digital baseband at these operating conditions during the processing of an entire packet. The effects of aggressive clock and input vector gating are clearly demonstrated: for instance, instantaneous synchronization power can vary by as much as 4X between sets of correlations, while consuming a peak instantaneous power that is 22X greater than simply averaging input samples.

B. Synchronization Performance

The synchronization performance of the digital baseband is compared to alternative classifiers in Fig. 14 using synchronization error rate (SER) curves in an AWGN environment. An error event is defined to occur when synchronization results in a phase that is not within ± 1 ns (i.e., one chip period) of the ideal phase. Since data from existing work is difficult to obtain and often uses different codes and preamble sizes, comparisons are instead made to in-house simulations using repetition codes (sliding window, SW) or alias-free codes (all others).

The proposed baseband achieves less than a 1.5 dB SNR loss over an ideal maximum likelihood (ML) classifier, and is within ± 0.5 dB of the optimum matched filter (MF) classifier. Note that all simulations use floating-point arithmetic, and thus would likely suffer an additional 0.5–2 dB implementation loss when quantized in a silicon ASIC. Compared to a sliding window (SW) classifier, the baseband requires at least 10 dB less SNR

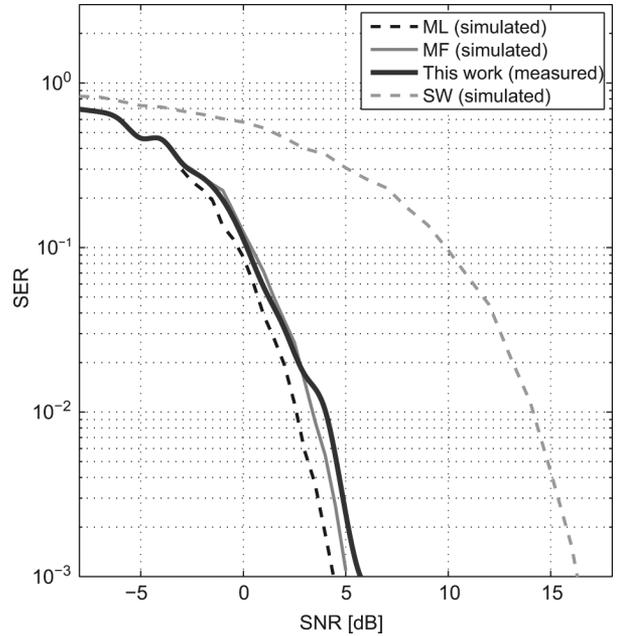


Fig. 14. Measured synchronization error rates (SERs) for the proposed baseband compared to floating-point simulations of maximum likelihood (ML), matched filter (MF), and sliding window (SW) classifiers. All curves are for a ± 1 ns synchronization accuracy.

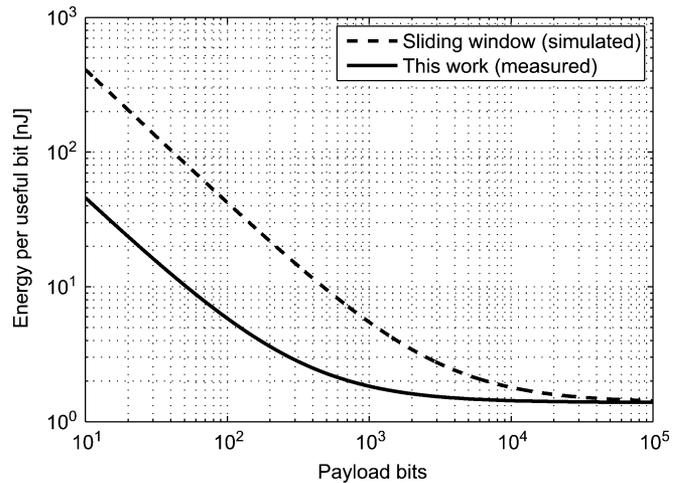


Fig. 15. Energy per useful bit (EPUB) for a synchronization accuracy of ± 1 ns with an SNR of 4 dB and an SER of 10^{-2} .

for an SER of 10^{-2} . In other words, the baseband processes approximately 10X fewer samples than an SW classifier for equivalent performance.

When coupled with the measured energy consumption results of a non-coherent IR-UWB RF front-end [15], the baseband achieves a complete receiver EPUB that is 3-to-9X lower than an SW-based classifier for payload lengths of 1000-to-10 bits. This is illustrated in Fig. 15. Alternatively, for 100-bit payloads and an average data rate of 100 kbps, battery life is extended by 5.8X compared to an SW-based classifier when using a 30 mAh hearing-aid battery.⁹ Note that the SW-based results are deemed

⁹This includes power drawn from the front-end and peripheral circuits on the SoC in [15], but does not include other board components such as an LDO.

TABLE II
COMPARISON OF EXISTING ASIC IMPLEMENTATIONS OF IR-UWB DIGITAL BASEBANDS

	[26]	[27]	[28]	[29]	[22]	This work
CMOS Technology	180nm	180nm	90nm	130nm	180nm	90nm
V_{dd}	1.2V	1.8V	0.4V	0.95V	1.8V	0.55V
Area	1.0mm ²	12.5mm ² †	8.4mm ²	2.0mm ²	0.25mm ²	2.6mm ²
Clock Frequency	62.5MHz	30MHz/300MHz	25MHz/100MHz	80MHz	150MHz	32MHz
Data Rate	62.5Mb/s	193kb/s	100Mb/s	40Mb/s	5Mb/s	16Mb/s
Decoding Type	NR-MF	NR-MF approx.	NR-MF	QAC	NC-QC	NC-ES
ADC resolutions	1bit	4bits	5bits	4bits	1bit	5bits
Payload modulation	BPSK	BPSK	BPSK	BPSK	BFSK	PPM
Preamble power	6.7mW	75mW	7mW	3.5mW	4mW	1.6mW
Payload Efficiency	–	–	17pJ/bit	35pJ/pulse	1300pJ/bit‡	12pJ/bit

abbrev.: NR-MF: Nyquist-rate Coherent Matched-Filter; QAC = Quadrature Analog Correlation; NC-QC = Non-Coherent Quadricorrelation; NC-ES Non-Coherent Energy-Sampling; †includes an RF front-end, ADC, and PLL; ‡includes a quadricorrelator, integrate & dump, and ADC.

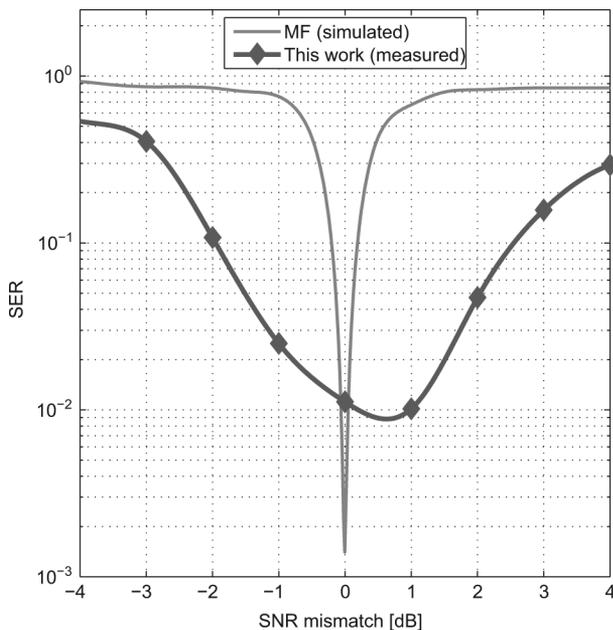


Fig. 16. SER for achieving a ± 1 ns synchronization accuracy in the presence of SNR estimator errors.

to be slightly conservative, as the simulated SW baseband is assumed to consume zero power.

Although the synchronization performance of the proposed baseband is comparable in ideal scenarios to a matched-filter classifier, it has a significant advantage in the presence of SNR estimation errors. Such errors are defined as the difference between the actual SNR seen by the receiver and the estimated SNR used to derive the QCORR multiplier coefficients. Fig. 16 shows the synchronization performance of the baseband and an MF classifier in the presence of SNR estimation mismatch. In this example, the measured eye opening of the baseband is 4.5 dB at an SER of 10^{-1} , which is 11X greater than the 0.4 dB simulated eye opening of the MF classifier. Note that the eye opens up significantly if the absolute SNR is increased or if synchronization accuracy tolerances are relaxed.

C. Summary

The digital baseband consumes 129 μ W of idle-mode power, and can process an entire preamble in a minimum of 14 μ s.

This results in a total energy consumption of 34 nJ, or 2.4 mW average power. The longest possible preamble requires 79 μ s and consumes 55 nJ of energy, equivalent to 0.7 mW of average power. Typical preambles require 25 μ s and consume 1.6 mW of average power. For comparison, the front-end consumes 5-to-14 mW during a typical preamble when duty cycled according to the packet diagram in Fig. 2.

Table II provides a chip summary of the proposed digital baseband while comparing results to previously published work. Note that it is difficult to directly compare existing work in IR-UWB digital basebands, as system parameters such as modulation formats, coherency, codes, preamble lengths, and data rates can vary greatly between different implementations. Additionally, the functionality provided in previous work can greatly differ, such as omitting packet delimiter decoders (e.g., SFDs) [22]. Nevertheless, the proposed baseband achieves the lowest average preamble power and payload demodulation energy efficiency. Note that several important references were omitted from the table due to having only FPGA results [7], [23], [24], synthesized ASIC simulation results [25], or not enough information devoted solely to the digital baseband [10].

VII. CONCLUSION

This paper presents a digital baseband that proposes the use of alias-free codes and a quadratic correlation technique to improve synchronization performance over existing non-coherent IR-UWB work. Alias-free codes allow for highly-accurate ± 1 ns synchronization tolerances while operating all circuits off a 32 MHz clock. Quadratic correlation is introduced as a synchronization classifier that has nearly optimal performance in ideal conditions, while having performance that is more robust than an MF classifier in the presence of SNR estimation errors. The baseband exploits parallelism and clock gating to achieve low-voltage, low-power operation, with supply voltages as low as 0.55 V while still meeting throughput constraints. The baseband can process an entire preamble in a minimum of 14 μ s, and consumes an average power of 1.6 mW.

ACKNOWLEDGMENT

The authors thank Nathan Ickes for testing support.

REFERENCES

- [1] First Report and Order FCC, 2002, FCC 02-48.
- [2] M. Z. Win and R. A. Scholtz, "Impulse radio: How it works," *IEEE Commun. Lett.*, vol. 2, pp. 36–38, Feb. 1998.
- [3] L. Smaini, C. Tinella, D. Helal, C. Stoecklin, L. Chabert, C. Devaucelle, R. Cattenoz, N. Rinaldi, and D. Belot, "Single-chip CMOS pulse generator for UWB systems," *IEEE J. Solid-State Circuits*, vol. 41, pp. 1551–1561, Jul. 2006.
- [4] D. D. Wentzloff and A. P. Chandrakasan, "A 47 pJ/pulse 3.1-to-5 GHz all-digital UWB transmitter in 90 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 118–119.
- [5] J. Ryckaert, G. Van der Plas, V. De Heyn, C. Desset, B. Van Poucke, and 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*, vol. 42, pp. 2860–2869, Dec. 2007.
- [6] P. P. Mercier, D. C. Daly, and A. P. Chandrakasan, "An energy-efficient all-digital UWB transmitter employing dual capacitively-coupled pulse-shaping drivers," *IEEE J. Solid-State Circuits*, vol. 44, pp. 1679–1688, Jun. 2009.
- [7] L. Stoica, A. Rabbachin, H. Repo, S. Tiuraniemi, and I. Oppermann, "An ultrawideband system architecture for tag based wireless sensor networks," *IEEE Trans. Veh. Technol.*, vol. 54, pp. 1632–1645, Sep. 2005.
- [8] F. S. Lee and A. P. Chandrakasan, "A 2.5 nJ/bit 0.65 V pulsed UWB receiver in 90 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 42, pp. 2851–2859, Dec. 2007.
- [9] T.-A. Phan, V. Krizhanovskii, and S.-G. Lee, "Low-power CMOS energy detection transceiver for UWB impulse radio system," in *Proc. IEEE Custom Integr. Circuits Conf.*, Sep. 2007, pp. 675–678.
- [10] D. Lachartre, B. Denis, D. Morche, L. Ouvry, M. Pezzin, B. Piaget, J. Prouvee, and P. Vincent, "A 1.1 nJ/b 802.15.4a-compliant fully integrated UWB transceiver in 0.13 μm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2009, pp. 312–313.
- [11] *IEEE std. 802.15.4, Part 15.4: Wireless MAC and PHY specifications for Low-Rate Wireless Personal Area Networks*, , 2003.
- [12] J. Ammer and J. Rabaey, "The energy-per-useful-bit metric for evaluating and optimizing sensor network physical layers," in *Proc. IEEE Conf. Sensor, Mesh and Ad Hoc Comm. and Networks*, Sep. 2006, pp. 695–700.
- [13] J. G. Proakis, *Digital Communications*, 4th ed. New York, NY: McGraw-Hill, 2001.
- [14] M. Verhelst and W. Dehaene, "Analysis of the QAC IR-UWB receiver for low energy, low data-rate communication," *IEEE Trans. Circuits Syst.*, vol. 55, pp. 2423–2432, Sep. 2008.
- [15] D. C. Daly, P. P. Mercier, M. Bhardwaj, A. L. Stone, J. Voldman, R. B. Levine, J. G. Hildebrand, and A. P. Chandrakasan, "A pulsed UWB receiver SoC for insect motion control," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2009, pp. 200–201.
- [16] IEEE 802.15.4a Wireless MAC and PHY specifications for LR-WPANS 2007 [Online]. Available: <http://www.ieee802.org/15/pub/TG4a.html>
- [17] S.-Y. Jung and D.-J. Park, "Design of preamble signal for synchronization with UWB non-coherent energy detection receiver," in *Proc. IEEE Int. Conf. Ultra-Wideband*, Sep. 2005, pp. 464–468.
- [18] S. Vitavasiri, "A Non-Coherent Ultra-Wideband Receiver: Algorithms and Digital Implementation," M. Eng, Massachusetts Institute of Technology, , 2007.
- [19] M. Bhardwaj, "Communications in the Observation Limited Regime," Ph.D, Massachusetts Institute of Technology, , 2009.
- [20] P. P. Mercier, M. Bhardwaj, D. C. Daly, and A. P. Chandrakasan, "A 0.55 V 16 Mb/s 1.6 mW non-coherent IR-UWB digital baseband with ± 1 ns synchronization accuracy," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2009, pp. 252–253.
- [21] A. Wang and A. Chandrakasan, "A 180-mV subthreshold FFT processor using a minimum energy design methodology," *IEEE J. Solid-State Circuits*, vol. 40, pp. 310–319, Jan. 2005.
- [22] D. Barras, R. Meyer-Piening, G. von Bueren, W. Hirt, and H. Jaekel, "A low-power baseband ASIC for an energy-collection IR-UWB receiver," *IEEE J. Solid-State Circuits*, vol. 44, pp. 1721–1733, Jun. 2009.
- [23] S. Zhang, J. Zhang, M. Liu, S. Wang, L. Heng, and R. Zhou, "An MBOK-UWB SoC transceiver in 0.181 μm CMOS technology," in *Proc. Int. Conf. ASIC*, Oct. 2007, pp. 830–833.
- [24] C. Duan, P. Orlik, Z. Sahinoglu, and A. Molisch, "A non-coherent 802.15.4a UWB impulse radio," in *Proc. IEEE Int. Conf. Ultra-Wideband*, Sep. 2007, pp. 146–151.
- [25] M. Pezzin and D. Lachartre, "A fully integrated LDR IR-UWB CMOS transceiver based on 1.5-bit direct sampling," in *Proc. IEEE Int. Conf. Ultra-Wideband*, Sep. 2007, pp. 642–647.
- [26] C.-H. Yang, K.-H. Chen, and T.-D. Chiueh, "A 1.2 V 6.7 mW impulse-radio UWB baseband transceiver," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2005, pp. 442–608.
- [27] R. Blazquez, P. P. Newaskar, F. S. Lee, and A. P. Chandrakasan, "A baseband processor for impulse ultra-wideband communications," *IEEE J. Solid-State Circuits*, vol. 40, pp. 1821–1828, Sep. 2005.
- [28] V. Sze and A. Chandrakasan, "A 0.4-V UWB baseband processor," in *Proc. IEEE Int. Symp. on Low Power Electron. and Design*, Aug. 2007, pp. 262–267.
- [29] M. Verhelst and W. Dehaene, "A flexible, ultra-low-energy 35 pJ/pulse digital back-end for a QAC IR-UWB receiver," *IEEE J. Solid-State Circuits*, vol. 43, pp. 1677–1687, Jul. 2008.



Patrick P. Mercier (S'04) received the B.Sc. degree in electrical and computer engineering from the University of Alberta, Edmonton, AB, Canada, in 2006, and the S.M. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, in 2008, where he is currently pursuing his doctoral degree.

From May 2008 to August 2008, he worked in the Microprocessor Technology Lab at Intel Corporation, Hillsboro, OR, designing spatial encoding and low-swing circuits for multi-core on-chip communication networks. His research interests include the design of energy-efficient RF circuits and digital systems for both wired and wireless communication applications.

Mr. Mercier was a co-recipient of the 2009 ISSCC Jack Kilby Award for Outstanding Student Paper at ISSCC 2010. Mr. Mercier received a Natural Sciences and Engineering Council of Canada (NSERC) Julie Payette fellowship in 2006, NSERC Postgraduate Scholarships in 2007 and 2009, and an Intel Ph.D. Fellowship in 2009.



Manish Bhardwaj received the Bachelor of Applied Science (Honors) from Nanyang Technological University, Singapore, in 1997, and the Ph.D. in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, in 2009.

He designed application specific digital signal processors for Siemens Semiconductors from 1997 to 1999. From 2001 to 2005, he was a co-founder and Vice President at Engim Inc., a wireless semiconductor startup in Acton, MA, that designed chipsets for wideband 802.11 access points operating at 162 Mbps.



Denis C. Daly (S'02–M'09) received the B.A.Sc. degree in engineering science from the University of Toronto, Toronto, ON, Canada, in 2003, and the S.M. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, in 2005 and 2009, respectively.

Since January 2010, he has been a Senior Member of Technical Staff at Cambridge Analog Technologies, Bedford, MA, designing low power, high performance analog circuits including ADCs and PLLs. From May 2005 to August 2005 he worked in the Wireless Analog Technology Center at Texas Instruments, Dallas, TX, designing transceiver modules. From May 2003 to August 2003, he worked on high-speed signaling systems at Intel Laboratories, Hillsboro, OR.

Dr. Daly is a co-recipient of the 2009 ISSCC Jack Kilby Award for Outstanding Student Paper, the 2006 DAC/ISSCC Student Design Contest (Operational System Category), and was awarded a Student Paper Prize at the 2006 RFIC Symposium. He received Natural Sciences and Engineering Research Council of Canada (NSERC) Postgraduate Scholarships in 2003 and 2005.



Anantha P. Chandrakasan (F'04) received the B.S., M.S., and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1989, 1990, and 1994, respectively.

Since September 1994, he has been with the Massachusetts Institute of Technology, Cambridge, where he is currently the Joseph F. and Nancy P. Keithley Professor of Electrical Engineering. He is the Director of the MIT Microsystems Technology Laboratories. His research interests include

low-power digital integrated circuit design, wireless microsensors, ultra-wideband radios, and emerging technologies. He is a coauthor of *Low Power Digital CMOS Design* (Kluwer Academic Publishers, 1995), *Digital Integrated Circuits* (Pearson Prentice-Hall, 2003, 2nd edition), and *Subthreshold Design for Ultra-Low Power Systems* (Springer, 2006). He is also a co-editor of *Low Power CMOS Design* (IEEE Press, 1998), *Design of High-Performance Microprocessor Circuits* (IEEE Press, 2000), and *Leakage in Nanometer CMOS Technologies* (Springer, 2005).

Dr. Chandrakasan was a co-recipient of several awards including the 1993 IEEE Communications Society's Best Tutorial Paper Award, the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997, the 1999 DAC Design Contest Award, the 2004 DAC/ISSCC Student Design Contest Award, the 2007 ISSCC Beatrice Winner Award for Editorial Excellence and the 2007, 2008, and 2009 ISSCC Jack Kilby Award for Outstanding Student Paper. He has served as a technical program co-chair for the 1997 International Symposium on Low Power Electronics and Design (ISLPED), VLSI Design '98, and the 1998 IEEE Workshop on Signal Processing Systems. He was the Signal Processing Sub-committee Chair for ISSCC 1999–2001, the Program Vice-Chair for ISSCC 2002, the Program Chair for ISSCC 2003, and the Technology Directions Sub-committee Chair for ISSCC 2004–2009. He was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001. He served on SSCS AdCom from 2000 to 2007 and he was the meetings committee chair from 2004 to 2007. He is the Conference Chair for ISSCC 2010.