

An Ultra-Wideband Transceiver Architecture for Low Power, Low Rate, Wireless Systems

Ian D. O'Donnell, *Member, IEEE*, and Robert W. Brodersen, *Fellow, IEEE*

(Invited Paper)

Abstract—This paper presents the system architecture, modeling, and design constraints for a baseband, integrated, CMOS, impulse ultra-wideband transceiver targeting very low power consumption on the order of 1 mW. Intended for a sensor network application, the radio supports low communication rates (~ 100 kpbs) and ranging capabilities over short distances (~ 10 m). Based on a “mostly digital” architecture, the analog complexity is reduced by moving the A/D convertor as close to the antenna as is reasonable. Pulses are generated from simple digital switches, overlaying the signal energy on the lower FCC UWB band (0–960 MHz). Reception is achieved using baseband gain blocks feeding a time-interleaved bank of low resolution A/D converters. A window of energy is captured in time and fed to the digital backend for processing. To save power and area, the digital backend implements only a pulse template correlation filter block overlaid with an additional spreading code. As a pulse template is used, no specific channel estimation or interference cancellation is assumed. The system performance is quantified for this case and implementation tradeoffs are explored with a strong focus on reducing power consumption. In particular, the issues of modulation choice, clock generation, gain and noise figure, ADC resolution, and digital signal processing requirements will be discussed.

Index Terms—Digital radio, impulse radio, low power, transceiver, ultra-wideband.

I. INTRODUCTION

THE APPROVAL of approximately 8 GHz of unlicensed spectrum in the U.S.A. (from 0 to 960 MHz and from 3.1 to 10.6 GHz) for ultra-wideband deployment presents an interesting research opportunity. The lack of specified physical layer signaling and low transmit power levels for UWB similar to Part 15 [1] have opened a very wide design space with a large possibility for innovation. While the majority of attention is focused on high-speed communication applications in the 3.1 to 10.6 GHz band over very short distances (~ 1 m), owing to the transmit power density constraints, there is also interest in power efficient ranging, imaging, and distance measurement with communication at relatively low data rates allowed below 960 MHz as well as in the upper band. One attractive method of

ultra-wideband signaling suitable for low power operation uses short pulses, on the order of nanoseconds, to spread energy over at least 500 MHz of bandwidth. The baseband-like nature of this signaling promises a low cost, low power architecture because of the simplified analog front-end design. Further power savings are possible through circuit operation duty-cycling between pulse reception windows. This architecture has the potential for much lower power consumption and higher integration than conventional approaches due to the wideband, i.e., low-Q, nature of the radio.

The anticipated power savings, over a traditional sinusoidal-based transceiver, come from the elimination of frequency translation and synthesis, removal of filtering and reduction of external components, and the duty-cycled nature of pulse generation and reception. However, the use of impulse signaling with a “mostly digital” approach, while easing some problems, moves the design challenge to different dimensions. In particular, the ADC speed and resolution become of utmost importance. Baseband Nyquist sampling of the lower UWB band requires approximately 2 GHz ADC clocking which has the potential to consume enormous amounts of power relative to our 1 mW target. The effect of large, in-band, sinusoidal interferers also must be determined. Additionally, the sub-nanosecond timing for ADC sampling may imply severe limitations on oscillator matching or jitter requirements. The power consumed in the wideband front-end gain stages, and the necessary sensitivity and gain requirements for those blocks, are also important. Finally, the area and power burden required for digital signal processing and demodulation must be considered as well.

Several notable ultra-wideband system architectures have been published. Earlier approaches were based on different analog architectures where the pulse correlation is performed in the analog domain before A/D conversion [2]–[7]. More recent digital architecture publications are split between a channelized, or frequency-based, approach [8]–[10], or direct time-based sampling of the UWB signal [11], [12]. While [12], intended as a baseband system for a 3.1 to 10.6 GHz communication link, is similar in architecture to what will be proposed here, the reported power consumption is too large for our target application (primarily due to the 4-bit ADC and clock generation.) To date, the focus of these architectures has not been on low power consumption or its trade-off with system performance. This paper addresses the issues inherent in a low power design, identifying the time-based, baseband, “mostly digital” architecture as a viable candidate, and concentrating on the mapping of circuit constraints to a low power, low cost implementation.

Manuscript received February 23, 2005; revised June 23, 2005. This research was supported by the Office of Naval Research (Award No. N00014-00-1-0223), an Army Research Office MURI Grant (#065861), and the industrial members of the Berkeley Wireless Research Center. The review of this paper was coordinated by Prof. R. Qiu.

I. D. O'Donnell is with the Berkeley Wireless Research Center (BWRC), Berkeley, CA 94704-1302 USA (e-mail: ian@eecs.berkeley.edu).

R. W. Brodersen is with the Department of Electrical Engineering and Computer Science (EECS), University of California, Berkeley, Berkeley, CA 94720 USA and also with the Berkeley Wireless Research Center (BWRC), Berkeley, CA 94704-1302 USA (e-mail: rwb@bwrc.eecs.berkeley.edu).

Digital Object Identifier 10.1109/TVT.2005.854021

Section II starts with a description of the proposed transceiver architecture. Next, a linear model for the system is developed in Section III to aid in analysis of design tradeoffs, and modulation schemes are analyzed. Section IV groups together the conclusions from this analysis, focusing on ADC resolution, sampling clock generation, front-end gain requirements, and the digital signal processing demands. In addition to calculations, time-domain simulations are run to verify the modeling and analyses. Section V summarizes the system specification and concludes the paper.

II. TRANSCEIVER ARCHITECTURE

In previous publications, the analog-based transceiver typically places A/D conversion after wideband gain, filtering, and high-speed analog correlation. This results in a slower, on the order of the symbol rate, A/D converter with higher resolution requirements. While the ADC power consumption is manageable, this requires more analog circuitry to operate at the full signal bandwidth. In particular, the correlation operation and template generation must be very high speed, which implies a tradeoff between power consumption and template generation accuracy. For low power operation a simple template is desired, such as a rectangular pulse, but this template is inflexible, reducing system performance out of proportion to the power savings gained. In addition, scaling up such an architecture for RAKE reception or faster acquisition places a large load at the critical, high-speed input to the correlators, thus increasing power consumption beyond linear scaling of the number of correlators. Due to these issues and to take advantage of digital circuitry's flexibility, scalability, and ability to trade area for power consumption, the partitioning between analog and digital sections was chosen as close to the antenna as is feasible. A direct, time-based sampling approach was taken to avoid the need to design an integrated well matched bank of filters in the analog frontend and to keep the digital backend simple. This has the cost of increasing the burden upon the A/D converter; possibly to the point where the ADC block consumes the most power in the system. However, we will show that in an interference dominated environment very low resolution ADC's may be utilized, thereby mitigating the power consumption penalty. A byproduct of the flexibility of this architecture is that it also provides a platform for further experimentation. By moving the signal processing into the digital domain, it is easier to prototype different receiver approaches.

Shown in Fig. 1 is a block diagram for the proposed "mostly digital" architecture [13]. As the received energy is localized in time around the channel delay spread, the receiver only needs to operate during that relatively narrow time window. To meet the Nyquist criterion, this window must be sampled at a high rate, on the order of 2 GSamples/s for 1 GHz of bandwidth. Reception consists of gain, matched to the antenna impedance, followed immediately by sampling and digitization. The digital samples are then fed to the digital backend for processing; e.g., acquisition, synchronization, and detection.

III. SYSTEM MODELING

To explore the performance tradeoffs for this system, a linear model is created that incorporates signal, noise, and interference,

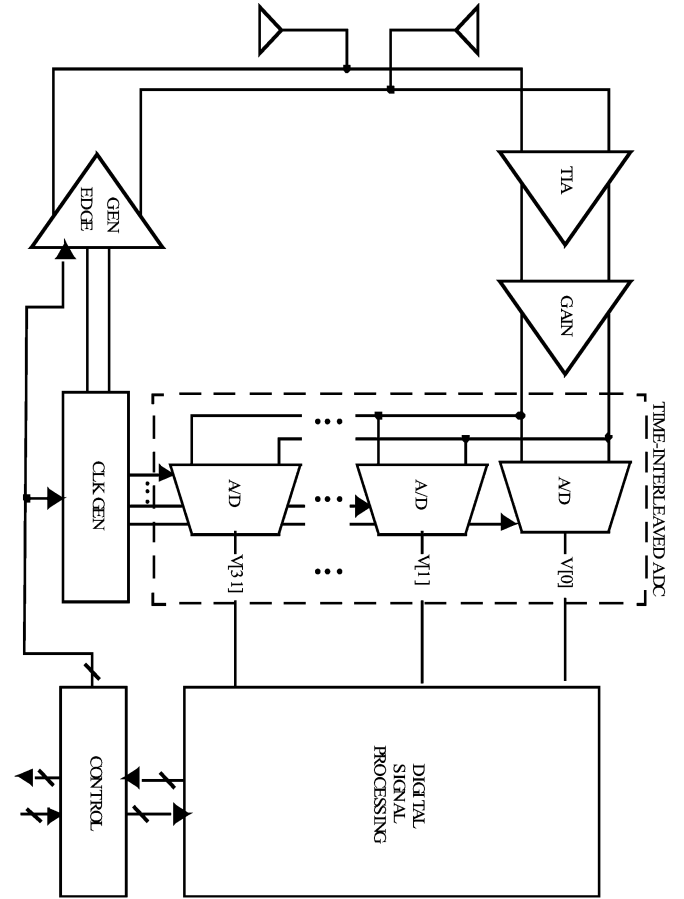


Fig. 1. Transceiver architecture.

including circuit nonidealities such as limited gain, filtering effects, quantization noise, and noise figure. For this model, the metric chosen is the "signal to noise + interference ratio" (SNIR) at the output of a pulse template filter. In the absence of interference, this can represent the optimal matched filter response [14]. For the purpose of a simple, low power digital implementation, no channel estimation beyond the knowledge of the pulse template is assumed and no interference cancellation is utilized. This provides an estimate for the degradation due to interference, and allows us to examine the impact on system performance from ADC resolution and the subsequent digital correlation precision.

We define the sampled, received signal after the ADC as

$$V = S + N + I + X$$

where S is K samples in time of the desired pulse; equal to the received pulse after gain and filtering

$$S = [s[0] \quad s[1] \quad \cdots \quad s[k-1]]$$

and N is K samples in time of Gaussian noise; variance set by the background noise floor times the system power gain and noise factor of the front-end

$$n[k] = \mathbf{N}(0, A_v^2 \cdot NF \cdot kTBR)$$

and I is K samples in time of the total narrowband interference seen at the ADC input

$$I = [i[0] \quad i[1] \quad \dots \quad i[k-1]]$$

where a narrowband interferer is modeled as a sinusoid with the equivalent power and uniform random phase

$$i[k] = \sum_{n=0}^{N-1} A_n \cos(\omega_n T_{\text{sample}} k + \theta_n)$$

and, X represents the quantization error; assumed to be zero mean and uniform over ± 0.5 lsb

$$x[k] = \mathbf{U} \left(0, \frac{\Delta_{A/D}^2}{12} \right).$$

Defining the matched filter coefficients as

$$W = S + Y$$

where S is again K samples of the desired pulse, and Y represents the quantization error for the matched filter coefficients; also assumed to be zero mean and uniform over ± 0.5 lsb

$$y[k] = \mathbf{U} \left(0, \frac{\Delta_{\text{MF}}^2}{12} \right).$$

Then the output of the matched filter, Z is equal to

$$Z = VW^t$$

and we may define the SNIR as

$$\text{SNIR} = \frac{E[Z]^2}{\text{Var}[Z]}.$$

Recall that the noise is zero mean, hence

$$E[Z] = (SS^t) = \sum_k s[k]^2 = P_s.$$

Then, **SNIR** is

$$\frac{P_s^2}{P_s (\sigma_{NX}^2 + \sigma_Y^2) + (SR_{II}S^t) + K\sigma_y^2 \left(\sigma_{NX}^2 + \sum_{n=0}^{N-1} \frac{A_n^2}{2} \right)}$$

where

$$\sigma_{NX}^2 = \sigma_N^2 + \sigma_X^2$$

and R_{II} is a $K \times K$ matrix whose elements are given by

$$r_{II}[i, j] = \left(\sum_{n=0}^{N-1} \frac{A_n^2}{2} \cos(\omega_n T_{\text{sample}}(i - j)) \right)$$

for $i, j = [0, 1, \dots, K-1]$.

A. A Comment on Modulation Choice

The previous equation predicts the SNIR at the output of a pulse template filter responding to an input pulse in the presence of noise and interference. This may be used to quantify the performance of PAM, but to investigate OOK, PPM or bi-orthogonal modulation, the output of the template filter must be determined in the absence of signal as well as the presence. Due

to the expectation that the channel will be interference dominated, and hence low SNIR, only the simple variants of these modulation schemes will be discussed.

In the absence of a received signal, we define U as

$$U = N + I + X$$

where the output of the matched filter is now Z_o

$$Z_o = UW^t$$

which is zero-mean, and has variance

$$\text{Var}[Z_o] = P_s (\sigma_{NX}^2) + (SR_{II}S^t) + K\sigma_y^2 \left(\sigma_{NX}^2 + \sum_{n=0}^{N-1} \frac{A_n^2}{2} \right)$$

which is

$$\text{Var}[Z_o] = \text{Var}[Z] - P_s (\sigma_Y^2) \approx \text{Var}[Z].$$

Assume σ_Y^2 is $\ll \sigma_X^2$ (i.e., that the matched filter's resolution is at least 3 bits larger than the ADCs).

As the variance of Z_o is essentially equal to the denominator of the SNIR per pulse, we expect OOK and 2-PPM performance to be approximately 3 dB worse than binary antipodal (2-PAM). Furthermore, receiver power consumption for 2-PPM will increase as the analog front-end is on for two reception windows during a pulse repetition period, assuming the PPM separation is larger than the delay spread of the channel to ensure orthogonality. This implies that OOK is also undesirable compared to binary antipodal as it achieves worse performance for the same receiver power consumption. 2-PPM is even worse compared to binary antipodal and OOK for modulation as it has both worse BER performance and doubled power consumption. Note that transmit power is severely limited by FCC specification, hence the receiver power dominates the total power consumption even for low transmit efficiency.

A variant on PAM and PPM modulation is bi-orthogonal signaling which combines binary antipodal and 2-PPM. While this achieves a higher data rate at 2 bits/symbol, this is cancelled out by doubling the power consumption, and the BER performance is predicted to be similar to 2-PPM, due to the shorter minimum distance between the orthogonal signal components. This implies that bi-orthogonal signaling is not as power efficient as binary antipodal. For this reason binary antipodal was chosen as the preferred modulation scheme.

IV. DESIGN CONSTRAINT MAPPING

A. ADC Resolution

Using the linear analysis derived in Section III, we may explore the system specification and quantify the design tradeoffs inherent in the resolution of the ADC. The ADC resolution is one of the most important of the design requirements, as the high sampling rate, on the order of 2 GHz, may preclude low power operation entirely. Low resolution, high sample rate ADC designs published in recent years may be surveyed to get a predictive estimate of power consumption for a given specification. Using the results from [15], ADCs are compared using

TABLE I
LOW RESOLUTION ADC FIGURES OF MERIT

F_{sample} (GSa/s)	P_{diss} (mW)	# bits	FOM	Ref	Year
1.0	70	4	229e9	[16]	2004
2.0	310	6	413e9	[17]	2003
1.6	328	6	312e9	[18]	2002
0.5	200	6	160e9	[19]	2002
1.3	545	6	153e9	[19]	2001
1.1	363	6	194e9	[19]	2001
0.7	187	6	240e9	[19]	2000
0.8	400	6	128e9	[19]	2000
0.5	400	6	80e9	[19]	1999
0.5	375	6	85e9	[19]	1999
0.4	188	6	136e9	[19]	1998




Fig. 2. SNIR versus interference over ADC resolution.

the following figure of merit (FOM):

$$\text{FOM} = \frac{2^{N_{\text{bits}}} \cdot F_{\text{sample}}}{P_{\text{diss}}}$$

Table I shows ADC FOM performance for recent publications. Using the best figure of merit of approximately $4e11$, we estimate the power consumption of a 4 bit 2 GSa/s ADC as 80 mW! As we move to lower resolutions, the ADC simplifies and the power decreases roughly as $1/2^N$ as most high speed ADCs are a flash architecture, and power scales roughly as the number of comparators. While this implies that even a 1 bit ADC (degenerate case) is on the order of 5 mW, in reality, the value will be less than 1 mW. For a given process there is a sense of the “natural” dynamic range based on the matching performance of the devices, usually on the order of 2 or 3 bits. Below this resolution, only simple comparators are needed (i.e., without offset cancellation or averaging). However, given the power concern, there is a strong drive to lower the ADC resolution, if possible, to save on power consumption. Hence, we need to quantify the system ADC resolution requirements.

To examine this, the SNIR per pulse is calculated at a given level of interference, noise, gain, etc. The results are shown in Fig. 2, plotted against the total received interference power.




Fig. 3. Simulated versus predicted error from SNIR for 1 bit ADC resolution.

Calculations used a Gaussian monocycle pulse [2] sent at a 5 MHz rate. Interference was generated based on measurements taken in our lab with a spectrum analyzer to represent “typical” levels and then scaled over the range shown. The UWB channel model was for a 3 m path; derived from an in-house ray tracing tool which estimates the impulse response using a 3-D indoor building model [20]. A input-referred noise figure of 10 dB was assumed for the gain stages, and the gain was set w/AGC to allow only an infrequent amount of limiting. To model finite bandwidth of the input gain stages, the pulse is filtered with a 5-pole rolloff at 1 GHz.

In Fig. 2, we see that only at low levels of interference, where thermal noise dominates, does extra resolution in the ADC improve SNR. As interference increases, the impact of higher resolution in the ADC decreases. This realization, that ADC resolution in an interference dominated environment is not critical, allows us to simplify the ADC design to 1 bit to save power without incurring a tremendous penalty in performance. Typical values for the aggregate interference over 0 to 1 GHz measured in our labs are around -40 dBm, predicting about ~ 7 dB of loss relative to a higher resolution ADC. This result agrees with previous work on time-based mono-bit digital receiver modeling using a matched filter in the presence of AWGN [21]. Note that this performance is not optimal, as it makes no effort to cancel the interference or gather more pulse energy from reflections. For the goal of low power operation and low cost/complexity, though, the predicted performance is adequate for the application requirements of sensor network radios.

The linear model calculation over ADC resolution show a trend that indicates 1 bit is sufficient; however, linear modeling of quantization noise breaks down at very low resolutions (1 to 2 bits). To verify these results, a time domain BER simulation was also run with and without the ADC resolution limit for the same conditions. The results are shown in Fig. 3. The time-domain simulation matches the linear model well, deviating by only a dB or so over the range. This confirms that the linear model accurately evaluated the effect of ADC resolution and

that a 1 bit ADC may be used to save system power without an excessive degradation in performance.

To compensate for the loss due to ADC quantization in the system link budget, it would be more efficient to increase the transmit power by 4 given the low transmit power regulation: approximately -10 dBm total of average power over DC to 960 MHz. Unfortunately, the FCC limits the power spectral density for UWB emissions, so transmit power is fixed for a fixed pulse rate. This results in a loss of data throughput (by approximately $1/4$) to compensate for the choice of a 1 bit A/D converter. However, high throughput is not critical for sensor network applications which require data rates on the order of 10–100 kbps [22].

B. Sampling Clock Generation

The use of impulse UWB signaling may imply tight timing tolerances, but we will show that the requirements are reasonable. The main issues associated with sampling clock generation are the jitter performance of the system clock, matching between the TX and RX clocks, and identifying a sampling clock generation architecture suitable for a low power implementation.

1) *Oscillator Jitter*: The allowable jitter variance may be approximately mapped to a phase noise requirement for the oscillator [23]. The lowest frequency we need to consider for jitter is the symbol rate since the digital backend will track any frequency variations slower than that. Assuming the mean square phase deviation over a symbol is much less than 1 rad and taking the phase noise spectral density to be of the form

$$L(\Delta f) = \frac{K}{(\Delta f)^2}.$$

Then the corresponding phase noise, given the total accumulated jitter σ_T over T_{symbol} , is

$$L(\Delta f) \approx 2\pi^2 \sigma_T^2 \left(\frac{f_c}{\Delta f} \right) \frac{1}{T_{\text{symbol}}}.$$

For an accumulated jitter of 75 ps over a $100 \mu\text{s}$ symbol, which allows for a -0.14 dB degradation in the SNR for an ideal matched filter with a Gaussian monocycle, we would require -103 dBc/Hz at a 100 kHz offset for $f_c = 100$ MHz. This level of performance seems achievable, as [24] reports a low power oscillator, digitally trimmable to 0.3 PPM with phase noise -100 dBc at a 100 Hz offset. Fig. 4 shows these phase noise requirements versus symbol rate along with boundaries for performance from common implementations based on reported results in the literature. One can see that the jitter specification is relaxed enough to allow for a ring oscillator implementation, suggesting that complete integration is possible. However, as we will see, the matching requirement between the TX and RX oscillators will be more stringent and will likely preclude a ring oscillator without an external precision component or crystal.

2) *Clock Matching*: The matching between the transmit and receive clocks must be accurate enough to allow the digital backend to track the drift. In our design the correlation results are compared at the symbol rate, thus requiring the drift over a symbol's reception to be a fraction of a sampling bin to keep




Fig. 4. Allowable phase noise for 75 ps standard deviation jitter over one symbol.

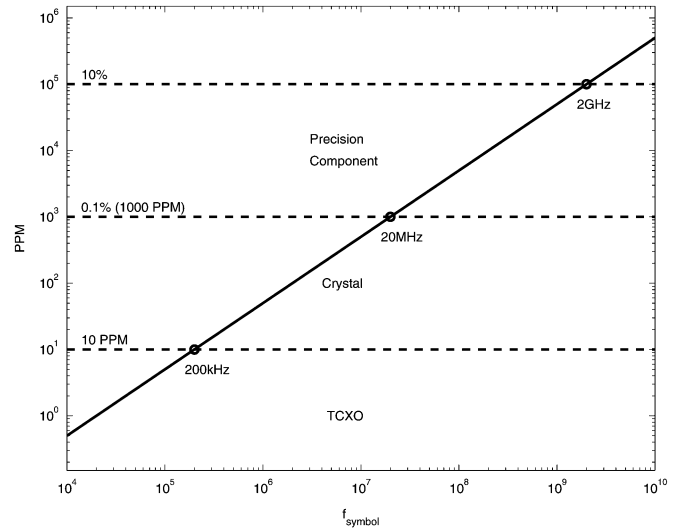


Fig. 5. Allowable mismatch for 100 ps drift over one symbol.

the energy within that correlator. Defining $f_c = 0.5 * (f_{RX} + f_{TX})$ and $\Delta f = |f_{RX} - f_{TX}|$, we may express this constraint as

$$\frac{\Delta f}{f_c} \approx \frac{1}{2} f_{\text{symbol}} \Delta T_{\text{bin}}.$$

Given a minimum symbol rate of 10 kHz, then for a -0.15 dB degradation from an ideal matched filter, 100 ps of drift is the worst-case allowable. From Fig. 5 we see that this requires very stringent matching with $\Delta f/f_c$ equal to 0.5 PPM. This value is only necessary to support the slowest symbol rate; i.e., a 1 MHz pulse rate with a length 1000 spreading-gain sequence. This implies that a crystal oscillator will be necessary if longer transmit ranges (and hence longer spreading codes) or slower pulse rates (for heavily duty-cycled power savings) will be used. For our system, a lower cost and lower precision crystal was selected in conjunction with tuning through oscillator pulling, to meet the matching specification.

TABLE II
ERROR PROBABILITY GIVEN ADC OFFSET

σ_N/σ_{VOS} :	10	20	33	50	100
p(Error)	3.2%	1.6%	1.0%	0.6%	0.3%

3) *Clock Generation Architecture*: To chose an architecture for clock generation, first we observe that we want the oscillator frequency to be as low as possible to save power in the oscillator, as g_m , and hence I_{bias} in submicron CMOS scales roughly as f_c^2 [24]. To accommodate this, either a phase locked loop (PLL) or delay locked loop (DLL) may be employed. A DLL was chosen as it is possible to drive the entire system from a slower clock if the pulse rate for the system is chosen to be at or below the delay spread of the channel. By selecting the delay line length and f_c appropriately, we can get well-controlled 0.5 ns steps between consecutive delay cells and generate a virtual 2 GHz effective sampling rate by combining the delayed clock phases. An additional benefit of a DLL is that it does not accumulate jitter [25] and hence should have better jitter performance than a PLL. The per-stage divider jitter is proportional to the output slope [26], and with careful design, the total jitter (oscillator plus DLL plus control logic generation for the ADC sampling) can be kept below the 75 ps target. From the oscillator, we also may derive the pulse repetition clock for transmission using a programmable divider for flexibility.

C. Gain, Offset, and Noise Figure

Ideally, a comparator switches exactly when one input is infinitesimally larger than the other. With this level of accuracy, no gain stages would be necessary as one could simply sample the antenna voltage directly. In practice, the offset voltage seen at the input of the comparator will determine the minimum amount of gain necessary to ensure accurate sampling. Modeling the offset, the probability of a comparator making a mistake is calculated as

$$P(\text{Error}) = \int P(\text{Error} | V_{OS})P(V_{OS}) dV_{OS}.$$

Assuming V_{OS} is Gaussian with a mean systematic offset μ_{VOS} and variance σ_{VOS}^2 , and taking the input as Gaussian, zero mean with variance σ_N^2 , we can calculate the probability of a comparator making an error. The exact impact of a comparator error depends on a particular set of \mathbf{Y} , the matched filter coefficients, hence the probability of a sampling error is analyzed for different σ_N/σ_{VOS} ratios (assuming the comparators are designed without systematic error).

For an error rate of 1%, the minimum gain necessary may be determined relative to the expected offset voltage variance. Mismatch simulations indicate that simple differential sampling with near minimum sized devices yields offsets on the order of tens of mVs for a 1 GHz tracking bandwidth. Incorporating offset cancellation into the comparator can bring this number down to several mV. The data in Table II implies the input signal to a 1 bit comparator must be on the order of 33 mV. The maximum gain condition would be a minimum input signal; i.e., thermal noise at room temperature over 1 GHz of bandwidth at

the 50 Ω input to the gain stages times the input noise figure, corresponding to 75 dB of gain. However, the expected levels of interference are much higher than this minimum; e.g., for a -40 dBm input, we need only about 25 dB. The minimum gain required depends upon the maximum interference level we wish to accommodate without clipping. At very high interference levels, no gain would be required at all, but performance would be very poor due to the large amount of interference. A reasonable range was chosen from about 50 to 10 dB, with a gain stage architecture that allows the ability to directly trade current consumption for gain. Especially for high gain, but even for low gain, offset will need to be controlled. Note that offset arises not only from the comparators, but also from the preceding gain stages. Without the use of offset cancellation techniques and/or capacitive coupling between stages, the systematic offset would saturate the comparators.

In addition to offset, the noise figure of the front end is often considered a critical design parameter. In an interference-dominated channel, though, this is not the case. The presence of large interferers dwarf the extra noise contribution from the front end circuits themselves. Thus, to save power, we may relax the LNA requirement without degrading the overall system significantly. The primary LNA design constraint becomes one of impedance matching to the antenna. This helps reduce power, as often the only way to achieve a low noise figure is to consume a large amount of current. In order to reject digital switching noise, a differential topology was chosen even though this doubles the power consumption. The gain stages, and wideband input-matching stage in particular, are predicted to consume the most power in the system.

D. Digital Signal Processing Requirements

As the processing has been moved into the digital domain, one concern may be that the computational load will either balloon in area or dominate power consumption itself. To evaluate this, the pulse template filter resolution and the issue of acquisition and area are examined. Low rate, impulse UWB has an innate problem with fast acquisition as one must search over the entire cycle to find the pulse. This is made more difficult if a spreading sequence is overlaid, as the spreading phase must also be determined. To accommodate this, a hybrid parallel/serial architecture is chosen consisting of a bank of pulse template filters, with each filter followed by an independent bank of despreading correlators.

Using the same conditions as Section IV, the SNIR per pulse is recalculated against the template filter coefficient bit-width, as shown in Fig. 6. In this case we see that a 1 bit coefficient is not necessarily adequate, as performance is predicted to be worse over all levels of interference. Note that template coefficient resolution is a separate issue from ADC resolution. For a 1 bit ADC input, the filter is correlating the zero crossings of the input against the expected pulse shape. Intuitively, the more accurate our knowledge of the expected pulse shape, the better we can weight those zero crossings to estimate the presence or absence of a pulse. This is distinct from the case where the ADC becomes overwhelmed in noise and increased ADC resolution




Fig. 6. SNIR versus interference over template filter resolution.




Fig. 7. Simulated versus predicted error for template filter output given 1 bit ADC.

only captures the noise more accurately, which doesn't aid in estimating the presence of a pulse unless we are able to subtract the noise.

Fig. 7 depicts the result of a time domain simulation to double check the linear system model. Results for this case agree that more than 4 bits of coefficient resolution in the template filter produce no improvement in the system performance.

In order to explore the implementation cost of the digital backend, we assume a binary antipodal receiver with a spreading code overlaid on top of the pulses to improve reception range as discussed in [27] and shown in Fig. 8. Data from the front end enters the digital backend which aggregates several consecutive windows of data, 16 ns long each sampled at 2 GSa/s into a block of up to 256 samples (128 ns). To speed acquisition, 128 samples are searched in parallel by 128 matched filters. To guarantee that a pulse doesn't straddle the boundary between steps, 256 samples are needed. The pulse-matched filters are of length 128 samples (64 ns), sized to the expected delay spread for an UWB indoor channel [28] to allow for future experimentation with channel estimation and/or interference cancel-

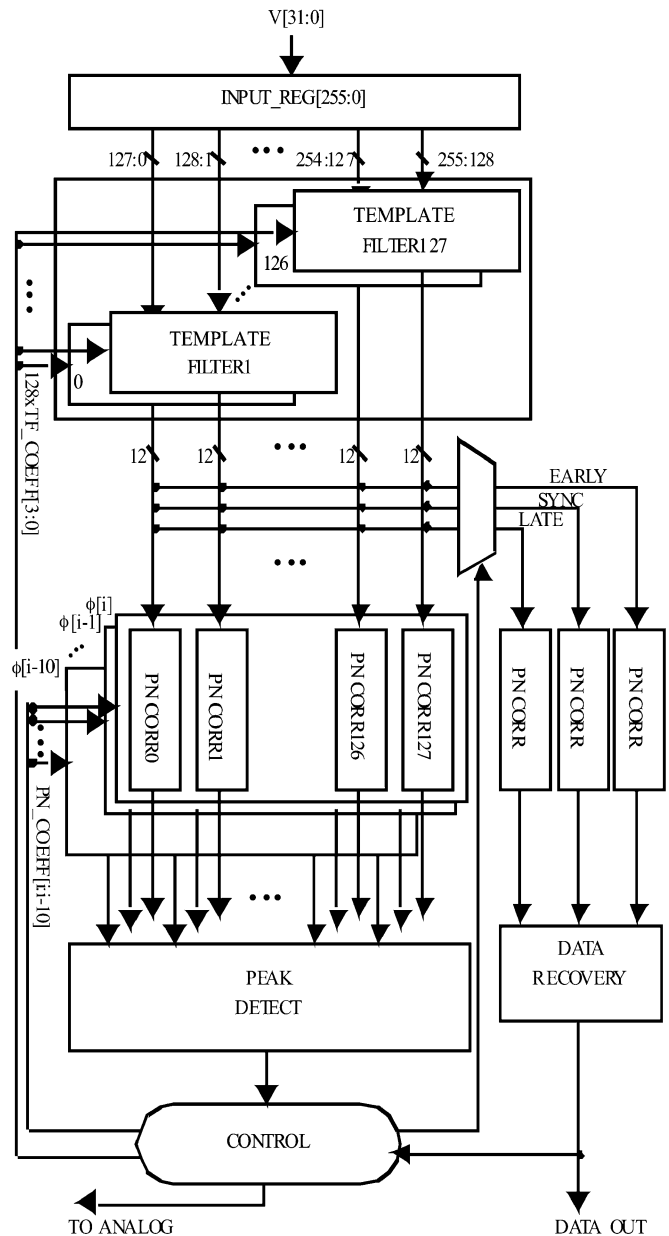


Fig. 8. Example digital backend architecture.

lation, even though a smaller pulse template is being used. The matched filter outputs are then sent to either an acquisition or synchronization block. For synchronization, as only three values, "early," "on-time" (or "sync"), and "late" are needed, all of the other matched filter inputs are disabled to save power. For acquisition, we search over all 128 samples and 11 spreading code phases at a time as a compromise between area and search time. Once a correlation peak above the programmable threshold is found by the peak detector logic, the backend switches from acquisition to tracking mode. Because binary antipodal is used, the data recovery block is a simple slicer based on a programmable threshold. In the interest of flexibility, two different spreading sequences may be used: one for acquisition and one while synchronized. Both sequences may be of length 1 to 1024.




Fig. 9. Template filter area versus resolution over filter length.




Fig. 10. Digital area versus worst-case acquisition time.

Fig. 9 depicts the impact of template filter coefficient resolution in area per filter block for different filter lengths. Increasing filter length causes a geometric increase in area. Increasing the matched filter resolution (bit-width) linearly increases the area with a tap-size dependent slope. While the delay spread of the channel may be larger than 64 ns (128 samples), often a majority of the energy is concentrated in 32 ns or less of time, which saves area at the expense of worst case performance. Note that the filter coefficients are fully programmable, in the interest of maintaining flexibility for experimentation.

For longer distance operation, we trade data rate for range by overlaying a spreading code; i.e., in a direct-sequence spread spectrum (DSSS) manner. This additional coding increases the acquisition burden on the digital backend and there is a geometric increase in complexity for correlating in parallel over the spreading sequence. Searching all phases simultaneously is prohibitively large. Fig. 10 shows the tradeoff between the total

digital area, template filter bank plus correlator banks, and acquisition time versus the number of phases of the spreading code correlated in parallel. For our design example, an area of around 10.2 mm² is predicted, using a window size of 256 samples for a 128 sample pulse size, with 4 bit coefficients in the matched filter, pulses sent at a 5 MHz rate and a maximal spreading length of 1024 chips. Note that to search a bigger (or smaller) window in parallel, these curves will move up (or down) by the same factor; i.e., 2x window is 2x area. Likewise, if either the pulse rate is sped up, or the spreading length decreased, then the acquisition time will improve by the square of the factor. For example, a shorter spreading sequence means fewer phases to search over in addition to a shorter wait for each sequence. Depending upon the desired conditions, these curves may be scaled or shifted to predict the area consumption.

V. CONCLUSION

In this paper, a system architecture for an integrated, CMOS, impulse ultra-wideband transceiver suitable for sensor network applications is presented and modeled. In particular, the system design constraints are quantified and traded off against implementation options with a strong focus on reducing power consumption. A mostly digital architecture with a 1 bit ADC front end is advocated due to the large power savings, simplicity, flexibility, and scalability it provides with only moderate performance penalty. Binary antipodal modulation is recommended as the best choice for power consumption efficiency for a fixed BER. The issues of jitter and matching are explored for the sampling time base, and matching was found to be the critical concern, with jitter predicted to less important. For sampling clock generation, a delay locked loop architecture driven from an off-chip tunable crystal reference was chosen. Bounds on the necessary amount of gain in the front end were calculated based on the offset at the ADC input. Due to the large amount of expected in-band interference, the noise figure of the front end is deemed less important and may be sacrificed to save power consumption without deteriorating system performance. Finally the backend digital system area and acquisition time are explored for a simple digital architecture, demonstrating that the burden created by moving the computation into the digital domain is reasonable if not sensible.

ACKNOWLEDGMENT

The author would like to thank M. S. W. Chen for valuable discussions and assistance with the digital backend design [27], and S. B. T. Wang (U.C. Berkeley) for useful comments.

REFERENCES

- [1] First Report and Order, FCC, FCC 02-48, Feb., 2002.
- [2] R. A. Scholtz, "Multiple access with time-hopping impulse modulation," in *Proc. MILCOM 1993*, vol. 2, New York, NY, Oct. 1993, pp. 447-450.
- [3] P. I. Withington and L. W. Fullerton, "An impulse radio communications system," in *Proc. Int. Conf. Ultra-Wideband, Short-Pulse Electromagnetics*, New York, NY, Oct. 1993, pp. 113-120.
- [4] R. A. Fleming and C. E. Kushner, "Spread spectrum localizers," U.S. Patent 5 748 891, Nicasio, CA, May 5, 1998.


- [5] C. J. Le Martret and G. B. Giannakis, "All-digital PAM impulse radio for multiple-access through frequency-selective multipath," in *Proc. GLOBE-COM'00*, vol. 1, Piscataway, NJ, pp. 77–81.
- [6] G. M. Maggio, N. Rulkov, and L. Reggiani, "Pseudo-chaotic time hopping for UWB impulse radio," *IEEE Trans. Circuits Syst. I*, pp. 1–12, Dec. 2001.
- [7] R. Hoctor and H. Tomlinson, "Delay-hopped transmitted reference RF communications," in *Proc. IEEE Conf. Ultra-Wideband Systems and Technologies 2002*, Piscataway, NJ, May 2002.
- [8] W. Namgoong, "A channelized DSSS ultra-wideband receiver," in *Proc. RAWCON 2001*, Piscataway, NJ, Aug. 2001, pp. 105–108.
- [9] H. J. Lee, D. S. Ha, and H. S. Lee, "A frequency-domain approach for all-digital CMOS ultra wideband receivers," in *Proc. IEEE Conf. Ultra Wideband Systems and Technologies 2003*, Piscataway, NJ, Nov. 2003, pp. 86–90.
- [10] S. Hoyos, B. M. Sadler, and G. R. Arce, "Analog to digital conversion of ultra-wideband signals in orthogonal spaces," in *Proc. IEEE Conf. Ultra Wideband Systems and Technologies 2003*, Piscataway, NJ, Nov. 2003, pp. 47–51.
- [11] I. D. O'Donnell and R. W. Brodersen, *A Highly-Integrated, Low-Power, Ultra-Wideband Transceiver for Low-Rate, Indoor Wireless Systems*. Berkeley: Qualifying Exam, Dept. Elect. Eng., Univ. of CA, Nov. 2000.
- [12] R. Blazquez, P. P. Newaskar, F. S. Lee, and A. P. Chandrakasan, "A baseband processor for pulsed ultra-wideband signals," in *Proc. IEEE Custom Integrated Circuits Conf.*, Piscataway, NJ, Oct. 2004, pp. 587–590.
- [13] I. D. O'Donnell, S. W. Chen, S. B. T. Wang, and R. W. Brodersen, "An integrated, low power, ultra-wideband transceiver architecture for low-rate, indoor wireless systems," in *Proc. IEEE CAS Workshop Wireless Communications and Networking*, Sep. 2002.
- [14] J. G. Proakis, *Digital Communications*, 3rd ed., New York: McGraw Hill, 1995.
- [15] R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 4, pp. 539–550, Apr. 1999.
- [16] L. Y. Nathawad, R. Urata, B. A. Wooley, and D. A. B. Miller, "A 20 GHz bandwidth, 4 b photoconductive-sampling time-interleaved CMOS ADC," in *Proc. Int. Solid-State Circuits Conf. Digest of Technical Papers*, vol. 46, Piscataway, NJ, Feb. 2003, pp. 320–321.
- [17] X. Jiang, Z. Wang, and M. F. Chang, "A 2 GS/s 6 b ADC in 0.18 μm CMOS," in *Proc. Int. Solid-State Circuits Conf. Digest of Technical Papers*, vol. 46, Piscataway, NJ, Feb. 2003, pp. 322–323.
- [18] P. Scholtens and M. Vertregt, "A 6b 1.6 GSample/s flash ADC in 0.18 μm CMOS using averaging termination," in *Proc. Int. Solid-State Circuits Conf. Digest of Technical Papers*, vol. 45, Piscataway, NJ, Feb. 2002, pp. 168–169.
- [19] C. Donovan and M. P. Flynn, "A 'digital' 6-bit ADC in 0.25- μm CMOS," *IEEE J. Solid-State Circuits*, vol. 37, no. 3, pp. 432–437, Mar. 2002.
- [20] H. Tang, "A unified approach to wireless system design," Ph.D. dissertation, Dept. Elect. Eng., Univ. of CA, Berkeley, 2003.
- [21] S. Hoyos, B. Sadler, and G. Arce, "Mono-bit digital receivers for ultra-wideband communications," *IEEE Trans. Wireless Commun.*, vol. 4, no. 4, Jun. 2005.
- [22] J. M. Rabaey, J. Ammer, T. Karalar, S. Li, B. Otis, M. Sheets, and T. Tuan, "Picoradios for wireless sensor networks—the next challenge in ultra-low power design," in *Proc. Int. Solid-State Circuits Conf. Digest of Technical Papers*, vol. 45, Piscataway, NJ, Feb. 2002, pp. 200–201.
- [23] J. A. Crawford in *Frequency Synthesizer Design Handbook*, Norwood, MA: Artech House, 1994.
- [24] Q. Huang and P. Basedau, "A 200 μA , 78 MHz CMOS crystal-oscillator digitally trimmable to 0.3 ppm," in *Int. Symp. Low Power Elect. and Design*, New York, NY, Aug. 1996, pp. 305–308, 390.
- [25] B. Kim, T. C. Weigandt, and P. R. Gray, "PLL/DLL system noise analysis for low jitter clock synthesizer design," in *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 4, New York, NY, Jun. 1994, pp. 31–34.
- [26] T. C. Weigandt, B. Kim, and P. R. Gray, "Analysis of timing jitter in CMOS ring oscillators," in *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 4, Jun. 1994, pp. 27–30.
- [27] M. S. W. Chen "Ultra-Wide-Band Baseband Design and Implementation", M.S. thesis, Dept. of Elect. Eng., Univ. of CA, Berkeley, 2002.
- [28] D. Cassioli, M. Z. Win, and A. F. Molisch, "The ultra-wide bandwidth indoor channel: from statistical model to simulations," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 6, pp. 1247–1257, Aug. 2002.



Ian D. O'Donnell (M'98) received the B.S. and M.S. degrees in electrical engineering and computer science from the University of California, Berkeley, in 1993 and 1996, respectively. His master's topic was in the area of digital, low power, CMOS circuit design for a wireless LAN receiver as part of the InfoPad project.

From 1996 to 1999 he worked at Silicon Graphics, Inc. as a Digital ASIC Designer, and in 1999 he joined NVIDIA, Inc. where he worked on high speed serial design. In 1998, he returned to Berkeley, joining the

Berkeley Wireless Research Center, to work in the area of low power, integrated, picocellular radios. His Ph.D. research focus is the design and implementation of an impulse based, low power, base-band Ultra-Wideband transceiver in 0.13 micron CMOS suitable for sensor network applications.



Robert W. Brodersen (M'76–SM'81–F'82) received the B.S. degree in electrical engineering and mathematics from the California State Polytechnic University, Pomona, CA, in 1966; the Engineering and the M.S. degree from the Massachusetts Institute of Technology (MIT), Cambridge, MA in 1968; and the Ph.D. degree in engineering from MIT in 1972. He received the award Honor Doctor of Technology (*Technologie Doctor Honoris Causa*) from the University of Lund, Sweden, 1999.

He is the John R. Whinnery Distinguished Professor in the Department of Electrical Engineering and Computer Science (EECS) at the University of California, Berkeley. He is also the Co-Scientific Director of the Berkeley Wireless Research Center (BWRC), where his research focus is new applications of integrated circuits as applied to personal communications systems with emphasis on wireless communications, low power design, and the CAD tools necessary to support these activities. From 1972–1976, he was a member of the Technical Staff, Central Research Laboratory, Texas Instruments, Dallas. He joined the EECS faculty at the University of California in 1976. Professor Brodersen is a member of the National Academy of Engineering.

With P. R. Gray and D. A. Hodges, he received the 1983 IEEE Morris N. Liebmann Award for pioneering work on switched-capacitor circuits. He was awarded the Technical Achievement Award from the IEEE Circuits and Systems Society in 1986, the Technical Achievement Award from the Signal and Processing Society in 1991, the IEEE Solid-State Circuits Award in 1997, the Mobicom Award: ACM Sigmobility's 1998 Outstanding Contribution to Mobile Computing, and the IEEE Golden Jubilee (Millennium) Medal in 2000 for exceptional contributions toward advancing the Society's goals.

He is the recipient of many awards for outstanding papers, has authored or contributed to numerous journals, conference papers, and books, and has served on the editorial board or as reviewer for scholarly journals and publications.