A/D PRECISION REQUIREMENTS FOR AN ULTRA-WIDEBAND RADIO RECEIVER

Puneet P. Newaskar, Raul Blazquez, Anantha P. Chandrakasan

Microsystems Technology Laboratory
Massachusetts Institute of Technology
Cambridge, MA

ABSTRACT

Ultra wideband radio (UWB) is a new wireless technology that uses narrow pulses to transmit information. Implementing an “all-digital” UWB receiver has numerous potential benefits ranging from low-cost and ease-of-design to flexibility. Digitizing an RF signal near the antenna, however, introduces its own set of challenges and has traditionally been considered infeasible. A high-speed, high-resolution analog-digital converter (ADC) is difficult to design, and is extremely power-hungry. The viability of an “all-digital” architecture, therefore, hinges upon the specifications of this block. In this paper, we demonstrate that 4 bits of resolution are sufficient for reliable detection of a typical UWB signal that is swamped in noise and interference.

1. INTRODUCTION

Ultra wideband radio is a promising new wireless technology, recently approved for commercial use by the Federal Communications Commission. A UWB signal is typically composed of a train of sub-nanosecond pulses[1], resulting in a bandwidth over 1 GHz. Since the total power is spread over such a wide swath of frequencies, its power spectral density is extremely low. This minimizes the interference caused to existing services that already use the same spectrum. On account of the large bandwidth used, UWB links are capable of transmitting data over tens of megabits per second. Other benefits include a low probability of interference and detection, precise locationing capability and the possibility of transceiver implementation using simple, “all-digital” architectures. (Some analog components are still needed at the front-end, but far fewer than what is required in conventional RF transceivers.)

A digital architecture means low cost, ease-of-design and most of the associated benefits of CMOS technology scaling. Furthermore, it allows for considerable flexibility. A single receiver could support different modulation schemes, bit-rates, qualities of service and operating ranges, and change these parameters dynamically. The vision of fully-configurable software radios is an exciting one, and it appears UWB may be more amenable to such a realization than conventional, narrowband systems.

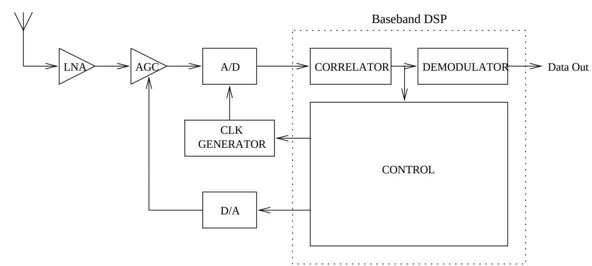

Fig. 1. Block Diagram of “All-Digital” UWB Receiver

Fig. 1 shows a generic block diagram for a digital UWB receiver. A key component of such a system is the analog-to-digital converter. In accordance with the Nyquist theorem, the ADC sampling rate for digitizing a UWB signal must be on the order of a few gigasamples/sec (GSPS). Even with the most modern process technologies, this constitutes a serious challenge. Most reported data converters operating at this speed employ interleaving[2], with each channel typically based on a FLASH converter. The latter is the architecture of choice for high-speed designs, but is not suitable for high-resolution applications[3]. An N-bit FLASH converter uses 2^N comparators so its power and area scale exponentially with resolution. Among recently reported high-speed ADCs (>1 GSPS) representing the state of the art[2],[4], none has a resolution exceeding 8 bits.

The minimum number of bits needed for reliable detection of a UWB signal is, therefore, a critical parameter. If excessively large, it can render an “all-digital” receiver infeasible. In this paper, we show that 4 bits are sufficient for reliable UWB reception. We first provide a theoretical framework for understanding this sufficiency, and then present simulation results to validate our hypothesis.

2. UWB SIGNALS

Information in such a system is typically transmitted using a collection of narrow pulses (0.2 ns to 1.5 ns) with a very low duty cycle ($\sim$ 1%)[1]. Each user is assigned a different pseudo-noise (PN) sequence that is used to encode the pulses in either position (PPM)[5] or polarity (BPSK[1])[6]. Channelization is thus based on the assigned code, as in the case of CDMA systems.

Suppose the bitstream is denoted by a sequence of binary symbols b_j (with values $+1$ or -1) for $j = -\infty, ..., \infty$. A single bit is represented using N_c pulses, where N_c refers to the length of the PN code c_i. For BPSK, the code modulates the polarity of a pulse within each frame. For PPM, it modulates the pulse positions (incrementing or decrementing them by multiples of T_c). Data modulation is achieved by setting the sign of the block of N_c pulses for BPSK. For PPM, we append an additional time-shift τ_{b_j} whose value depends on whether b_j is $+1$ or -1. Each frame has duration T_f; the duration of each bit is thus given by $N_c T_f$. Letting A denote the amplitude of each pulse p(t), the transmitted signal s(t) can be written as follows for the two different modulation schemes :

$$s_{BPSK}(t) = A \sum_{j=-\infty}^{\infty} \sum_{i=0}^{N_c-1} b_j c_i p(t - jN_cT_f - iT_f) \quad (1)$$

$$s_{PPM}(t) = A \sum_{j=-\infty}^{\infty} \sum_{i=0}^{N_c-1} p(t - jN_cT_f - iT_f - c_iT_c - \tau_{b_j}) \quad (2)$$

As mentioned above, each bit is represented by N_c pulses. This redundancy is one component of the signal's processing gain (PG). The other component is the duty cycle, a ratio of the short duration of each pulse to the large interval between successive ones. Processing gain refers to the boost in effective SNR as a UWB signal is processed by a correlating receiver.

$$PG/dB = 10log(N_c) + 10log(\frac{1}{duty\_cycle}) \quad (3)$$

3. RECEIVER DESIGN

Optimal detection of a noisy signal is based on matched filtering[7]. This would entail correlating the received signal r(t) against a template l(t) that is an exact replica of the original transmitted signal, and then feeding the correlator output to a slicer. However, generating exact replicas of sub-nanosecond pulses is a hard problem, and is highly susceptible to timing jitter. A more tractable approach is to use a template signal comprising a train of rectangular pulses that are, in general, wider than the actual received pulses but are coded with the same PN sequence. Fig. 2 illustrates the structure of the received and template signals, and depicts the operations necessary for demodulation. l_0(t) and l_1(t)

[1]The term BPSK (binary phase-shift keying) is somewhat of a misnomer in the context of a UWB signal; we are basically referring to antipodal signalling, but will continue to use this term for convenience.

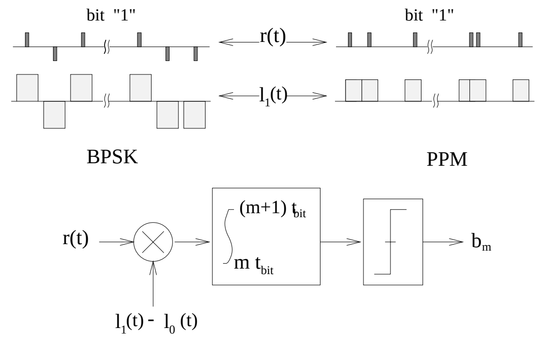

Fig. 2. Demodulating a UWB Signal

represent the template signals corresponding to a transmitted 0 and 1 respectively. They are related to one another by either a sign-inversion or a time-shift, depending on whether BPSK or PPM was employed. Equations describing l_1(t) and l_0(t) can be obtained from (1) and (2) by simply replacing the original pulse p(t) with a rectangular pulse rect(t), and setting b_j to $+1$ and -1 respectively. By using such template signals, we are essentially performing a form of windowing. In other words, we are correlating the received signal against its underlying PN code over narrow windows.

The digital implementation of such a receiver entails sampling the received signal (once it has been sufficiently amplified), converting it to digital form and then performing correlation. A figure analogous to Fig. 2 describing this process can be obtained simply be replacing r(t) and l(t) with their discrete-time counterparts r[n] and l[n].

4. ANALYSIS

Any form of distortion that is uncorrelated with the UWB signal's underlying PN code contributes to "noise" at the output of the correlator shown in Fig. 2. There are numerous sources of such noise : channel and circuit noise, quantization, other UWB signals, timing jitter and in-band narrowband interferers. In our model, we reduce all these disparate sources to just three categories : Additive white gaussian noise (AWGN), narrowband interference and quantization noise. For our analysis, we first discuss classical approaches to setting the ADC resolution in a narrowband radio receiver. We then present a method appropriate for a UWB system.

4.1. Traditional Approach

The minimum ADC resolution required is typically obtained [8] by either one of two methods :
(1) There is a certain minimum SNR (prior to detection) associated with the desired bit error rate (BER). The number of bits is chosen so that the overall SNR is above this

minimum value, assuming quantization to be the dominant source of noise.
(2) The desired signal is narrowband compared to the bandwidth digitized, and there are often powerful blocking interferers nearby. Although the blocker can be filtered out in the digital domain, there could be spurs that fall in the same band as the desired signal, generated by the non-linearity of the analog-to-digital conversion. The spurious-free dynamic range (SFDR) quantifies this non-linearity; it is defined as the ratio of the sinusoidal signal power to the peak-power of the largest spurious signal in the ADC output spectrum. For reliable detection of the desired signal, the power of the largest spur must be below signal power by SNR_{min}. This, in turn, requires the SFDR and the ADC resolution to be above a certain minimum.

Neither of the two methods outlined above is appropriate for a UWB system. The first one relies on the assumption that quantization noise dominates over all other sources. But a UWB signal is likely to be immersed in strong AWGN and interference due to stringent FCC restrictions on its transmit power. As for the second method, it is tied to the problem of a spur falling in the same band as the signal. This idea has little meaning when applied to a signal that is several gigahertz wide, relative to which practically all interferers and their spurs appear narrowband.

The bottomline is that the modeling and treatment of AWGN, interference and quantization noise must be done differently for a UWB system than for a narrowband system. The effect of each on the post-correlation SNR must be analyzed. Within this framework, we will consider the effect of quantization noise.

4.2. A New Framework

Upon sampling, the received signal is given by the following equation :

$$v[n] = s[n] + w[n] \tag{4}$$

where $s[n]$ is the desired signal. For BPSK[2], it has the following expression :

$$s[n] = A \sum_{j=-\infty}^{\infty} \sum_{i=0}^{N_c-1} b_j c_i p\left[n - jN_cN_f - iN_f\right] \tag{5}$$

N_f is equal to $\lfloor \frac{T_f}{T_s} \rfloor$ where T_f and T_s denote the frame duration and sampling interval respectively. $p[n]$ is taken to be a rectangular pulse of width W :

$$p[n] = \begin{cases} 1 & \text{if } 0 \le n < W \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

[2]For PPM, a similar analysis can be carried out. The only difference from BPSK is the appearance of a factor that divides the signal-to-noise ratio or signal-to-interference ratio in the expression for overall probability of error. This factor is 2 if the signals representing bits 1 and 0 are orthogonal.

In expression (4), $w[n]$ represents any form of distortion prior to the ADC. Following quantization by the ADC, we get an additional noise component $q[n]$:

$$r[n] = s[n] + w[n] + q[n] \tag{7}$$

We then perform correlation against a template signal $l[n]$ (same form as $s[n]$ but with wider rectangular pulses of width W_o). The signal at the output of the correlator $y[m]$ is given by :

$$y[m] = \sum_{n=mN_cN_f}^{(m+1)N_cN_f-1} r[n] \cdot l[n] = s'[m] + w'[m] + q'[m] \tag{8}$$

$s'[m]$, $w'[m]$ and $q'[m]$ represent the contributions of $s[n]$, $w[n]$ and $q[n]$ respectively to the final correlation value. $s'[m]$ can be viewed as the sum of samples from N_c pulses (each with amplitude A or -A, depending on the value of the information bit, and width W). Correlating against the right code ensures that these samples are added with the right signs. Thus, we obtain :

$$s'[m] = \begin{cases} -AN_cW & \text{if } b_m = 0 \\ +AN_cW & \text{if } b_m = 1 \end{cases} \tag{9}$$

$w'[m]$ and $q'[m]$ are both random variables, whose statistical properties and effect on post-correlation SNR we now discuss. In doing so, we treat separately the AWGN-limited and interference-limited cases.

4.3. AWGN-Limited Case

We begin by ignoring the quantization noise $q'[m]$ and consider only the effect of $w'[m]$. In theory, this is equivalent to using an infinite-resolution ADC.

4.3.1. No Quantization

During correlation, $N_c \cdot W_o$ uncorrelated samples of noise are added. Since each individual sample is a gaussian random variable with zero mean and variance σ^2, their sum $w'[m]$ is also gaussian, and has zero mean with variance equal to $\sigma^2 N_c W_o$. So $y[m] = s'[m] + w'[m]$ is gaussian with mean AN_cW and variance $\sigma^2 N_c W_o$. Thus the probability of error can be expressed as :

$$P_e = Q\left(\frac{AN_cW}{\sqrt{\sigma^2 N_c W_o}}\right) = Q\left(\sqrt{SNR \cdot N_c N_f \frac{W}{W_o}}\right) \tag{10}$$

Measured at the input of the receiver, SNR is equal to $\frac{A^2 d}{\sigma^2}$, where d is the duty cycle (given by the fraction W/N_f). We note that having a window larger than the pulse width thus results in a loss factor $L = \frac{W_o}{W}$. Capturing the full processing gain, therefore, requires setting $W_o = W$.

4.3.2. *Quantization Effects Included*

The ADC in our model has a fixed input range, from -1 to 1. So if the number of bits is b, the quantization step is given by $\Delta = \frac{1}{2^{b-1}}$. Due to the presence of gaussian noise at the input, we can reliably model the quantization noise as a uniform random variable of variance $\Delta^2/12$[3]. The implicit assumption here is that the ratio of the AWGN standard deviation to the quantization step falls within a certain range. If this ratio is too large, the probability of ADC saturation/clipping is high and the quantization noise added has variance well above $\Delta^2/12$. If the ratio tends to zero, on the other hand, the quantization noise power tends to $\Delta^2/4$. In order to avoid either of these sub-optimal regimes, we place an automatic-gain-control (AGC) circuit before the A/D converter. The AGC scales its noisy input signal by a factor α such that the A/D is fed an "optimal" input power of σ_o^2, given below in Table 1 for different resolutions. Due to this block, we can safely assume henceforth that the quantization noise power added by the ADC for all input SNR's is $\Delta^2/12$.

Table 1. σ_o Values Set by AGC

A/D Resolution	σ_o
2 bits	0.2850
3 bits	0.2025
4 bits	0.1425

Assuming the quantization noise is uncorrelated with the gaussian noise at the input, we can formulate the SNR after A/D conversion (but prior to correlation) as follows :

$$SNR_{\text{after ADC}} = \frac{\frac{A^2 d}{\alpha^2}}{\frac{\sigma^2}{\alpha^2} + \frac{\Delta^2}{12}} \tag{11}$$

After the ADC, we perform correlation during which we must account for the $N_c W_o$ samples of quantization noise that are added along with the signal and its input noise. We can assume these three components are all uncorrelated. Since $N_c W_o$ is typically a large number, the central limit theorem implies that the distribution of the summed quantization noise samples approximates a gaussian with zero mean, and variance equal to $\sigma_{quant}^2 = \frac{\Delta^2 N_c W_o}{12}$.

The σ^2 in (10) is now replaced by $\sigma^2 + \sigma_{quant}^2$. Combining the resulting expression with (11), we have the revised approximation :

$$P_e = Q\left(\sqrt{SNR_{\text{after ADC}} N_c N_f \frac{W}{W_o}}\right) \tag{12}$$

For an amplitude A of 1 and a duty cycle of 2%, Fig 3 shows the decrease in SNR after the ADC for various input SNR's and bit resolutions. This loss is marginal (less than

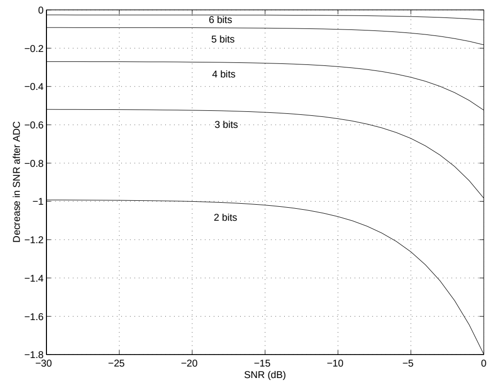

Fig. 3. Loss in SNR Due To Quantization

2dB[3]) over the range of SNR's within which we expect a UWB signal to operate. Furthermore, increasing the number of bits provides diminishing improvement.

4.4. Interference-Limited Case

Again, we start with no ADC and thus, zero quantization noise. First we must model the impact that the correlation process has on a narrowband signal at the input of the system. Let us assume it is given by $w(t) = B\cos(\Omega t + \phi_o)$. If this signal is sampled with sampling period T_s we obtain:

$$w[n] = w(nT_s) = B\cos(\Omega T_s n + \phi_o) = B\cos(\omega_o n + \phi_o) \tag{13}$$

These interference samples are then processed by the correlator This step can be viewed as running $w[n]$ through a filter $F(z)$ that yields an output of the form :

$$I[n] = BF_o \cos(\omega_o n + \phi_o + \theta_o) \tag{14}$$

After downsampling by a factor $N_c N_f$, this sinusoid may be aliased to a different frequency ω_o', but its amplitude does not change. We are left with the random variable :

$$y[m] = s'[m] + w'[m] = WN_c A + BF_o \cos(\Phi) \tag{15}$$

If the initial phase ϕ_o is a uniform random variable between 0 and 2π, then Φ is also a random variable between 0 and 2π, regardless of the distribution of the interferer's frequency. The probability of error can be shown to be :

$$P_e = \begin{cases} 0 & \text{if } WN_c A > BF_o \\ \frac{1}{2} - \frac{1}{\pi}\sin^{-1}\frac{N_c}{F_o}\sqrt{\frac{SIR \cdot N_f W}{2}} & \text{if } WN_c A < BF_o \end{cases} \tag{16}$$

where SIR (signal-to-interference ratio) is equal to $\frac{2dA^2}{B^2}$. As we can see, the behaviour depends on the length of the

[3]Without the AGC, the loss would be somewhat larger since the quantization noise power added by the ADC would be more than $\Delta^2/12$ as explained earlier.

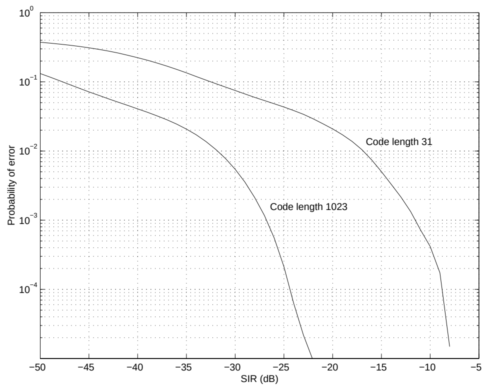

Fig. 4. Probability of Error vs SIR (Theoretical)

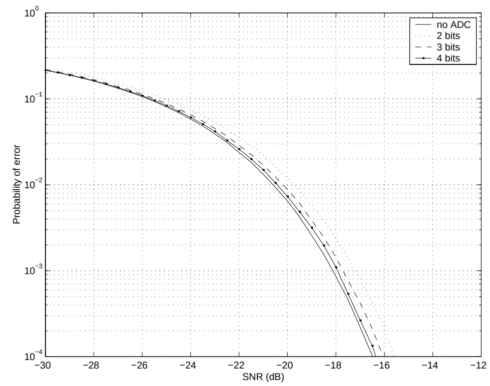

Fig. 5. Probability of Error vs SNR (Simulated)

code. For comparison purposes, simulations were run using the above equation for Gold codes of length 31 and length 1023. Their results are observed in Fig. 4. The longer code yields a probability of error of 10^{-3} for an SIR that is 15dB lower. This difference is equal to the difference in processing gains.

Developing closed-form expressions for the cumulative effects of interference and quantization noise on the error probability is cumbersome. We will thus rely on simulations to see the effect of increasing ADC resolution for the interference-limited case.

4.5. Summary of Analysis

Bit error rate (BER) in a UWB system is tied to the post-correlation SNR, which, in turn, is some function of signal power, AWGN noise, interference and quantization. High processing gain inherent in a UWB signal allows us to operate reliably (BER of 10^{-4} say) with a pre-correlation SNR that is considerably lower than for a typical narrowband system. This implies that the power levels of AWGN and interference at the input, corresponding to this target BER, are likely to be high. Our analysis confirms this. The models we have developed demonstrate the performance achievable in the absence of any quantization noise. We expect that introducing the latter should not cause a significant departure from the performance curves obtained without it. Simulations have been used to validate our hypothesis and to determine the minimum number of bits needed to get close enough to these "ideal" performance curves.

5. SIMULATIONS

Again, we treat separately the AWGN-limited and interference limited cases. The representation of the UWB signal is the same in each, however. Rectangular pulses of width 2 samples are used, with a pulse-to-pulse interval of 100 samples (ie. duty cycle of $2\ \%$). The PN code used is a Gold code of length 31 bits. Correlation is performed using a window of 10 samples per pulse.

5.1. AWGN-Limited Case

Noise samples are uncorrelated. The ADC is preceded by an AGC which sets the power fed to the converter based on the "policy" described in the previous section. For each value of SNR simulated, 10^6 independent trials were carried out, based on which a probability-of-error P_e was assigned. We ran Montecarlo simulations that provided a standard deviation under under $10\ \%$ for a P_e of 10^{-4}, and less than $1\ \%$ for a P_e of 10^{-3} (or higher). Our results are shown in Fig. 5 for ADC resolutions 2,3,4 and ∞ (no ADC).

The simulations closely match the results of our analysis. The shape of the curves agrees with the predictions of (10). For a P_e of 10^{-3}, the gap in SNR terms from the infinite resolution case decreases from $0.9\ dB$ to $0.4\ dB$ in going from 2 to 3 bits. Moving up to 4 bits lowers this gap to a mere $0.2\ dB$. There is little to be gained by increasing the resolution further.

5.2. Interference-Limited Case

The interference we wish to model is narrowband in the real sense of the word, and is thus described by a pure sinusoid, rather than as a modulated carrier with a finite data bandwidth. Thus, there are no abrupt changes of phase over the duration of one bit (ie. $N_c N_f$ samples). Its frequency is a uniform random variable in the range from 0 to half the sampling rate. Its initial phase is an independent uniform random variable from 0 to 2π. Our Montecarlo simulation provided a standard deviation of $10\ \%$ for a P_e of 10^{-4}, and less than $1\ \%$ for a P_e of 10^{-3} (or higher). The simulation results are shown in Fig. 6 for bit resolutions of 2,3,4 and ∞.

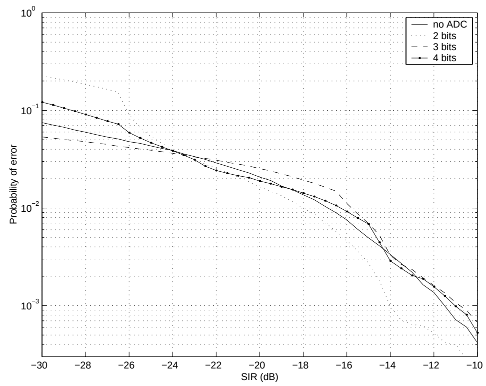

Fig. 6. Probability of Error vs SIR (Simulated)

The curve shape for the no-ADC case agrees with the results of our analysis (summarized by (16)). For 2 and 3 bits, we see kinks in the curve stemming from the severe nonlinearity of the ADC characteristic at low bit resolutions. We observe that for SIR's between -23dB and 0dB, the probability of error is actually lower for 2 bits than infinite bits (no ADC). There is an explanation for this. Firstly, note that when the signal-to-interference ratio is higher than -14dB, the signal amplitude is larger than the interference amplitude (the low duty cycle makes the ratio of average powers considerably lower than the ratio of instantaneous voltages). When the signal is larger than or comparable to the interferer in amplitude terms, a low bit resolution actually *suppresses* the effect of the latter. Consider the case of 2 bits. During the portion of the window where there is signal, the ADC output saturates (its sign reflecting the sign of the pulse) and is completely deterministic. During the rest of the window, there is only interference and no signal. The ADC output is now either a plus or a minus LSB (least-significant bit) and exhibits less variation than it would for higher bit resolutions. The above observation, however, applies only for a limited range of high SIR's. In terms of immunity to noise and large interferers, 4 bits is better than 2 or 3 bits and thus, a more appropriate design choice. Going for resolutions higher than this is unnecessary, however, due to the diminishing returns. Simulated curves for 5 and 6 bits have been omitted from Fig. 6 because they are indistinguishable from the ∞ resolution curve.

A clarification should be made regarding the pulse shape used in our analysis and simulations. As shown in Fig. 2, we employ rectangular pulses. In reality, such a pulse cannot be generated and represents an abstraction. More practical shapes include a gaussian, monocycle[1] or more generally, a wavelet. We contend, however, that there is no loss of generality in assuming a rectangular pulse, primarily because the various sources of noise we have identified can be assumed to be uncorrelated with the signal. Our claim that 4 bits are sufficient can thus be extrapolated to other pulse shapes.

6. CONCLUSIONS

In this paper, we address the problem of determining the minimum ADC resolution needed for reliable detection of a UWB signal. The unique nature of these signals requires a different approach for arriving at such a specification than traditional methods. Based on analysis and simulations, we have demonstrated that 4 bits of resolution are sufficient. This result eases one of the main obstacles to implementing an all-digital, software-defined UWB radio, namely high-speed A/D conversion within a reasonable power budget.

7. ACKNOWLEDGEMENTS

This work is funded in part by Hewlett-Packard under the HP/MIT Alliance. Puneet Newaskar is supported by an MIT Presidential Fellowship, and Raul Blazquez is supported by the Caixa Fellowship.

8. REFERENCES

[1] Time Domain Corporation, *www.time-domain.com.*

[2] Ellersick, W., Chih-Kong Ken Yang, Horowitz, W. and Dally, W., "GAD: A 12Gs/s CMOS 4-bit A/D converter for an equalized multi-level link," in *Symposium on VLSI Circuits, 1999, Digest of Technical Papers*, 1999, pp. 49–52.

[3] Razavi, B., *Principles of Data Conversion System Design*, IEEE Press, 1995.

[4] Poulton, K. et al., "An 8-GSample/s 8-bit ADC system," in *VLSI Symposium on Circuits*, June 1997.

[5] Le Martret, C.J. and Giannakis, G.B., "All digital PPM impulse radio for multiple-access through frequency-selective multipath," in *Sensor Array and Multichannel Signal Processing Workshop, 2000, Proceedigns of the 2000 IEEE*, 2000, pp. 22–26.

[6] Le Martret, C.J. and Giannakis, G.B., "All digital PAM impulse radio for multiple-access through frequency-selective multipath," in *Global Telecommunications Conference, 2000, IEEE*, 2000, vol. 1, pp. 77–81.

[7] Proakis, J.G., *Digital Communications*, McGraw Hill, Inc., 1983.

[8] Salkintzis, A.K., Hong Nie, Mathiopoulos, P.T., "ADC and DSP challenges in the development of software radio base stations," *IEEE Personal Communications*, vol. 6, no. 4, pp. 47–55, Aug. 1999.