A Compressed Sensing Receiver for Bursty Communication with UWB Impulse Radio

Anand Oka and Lutz Lampe
Dept. of Electrical and Computer Eng., University of British Columbia, Canada. {anando, lampe}@ece.ubc.ca

Invited Paper

***Abstract*—We propose a novel receiver for Ultra-Wideband Impulse-Radio communication in bursty applications like Wireless Sensor Networks. It is based on the principle of Compressed Sensing, and exploits the sparsity of the transmitted signal to achieve reliable demodulation. Instead of a full-fledged high-rate A/D, a modest number of projections of the received signal are acquired using analog correlators, and a joint decoding of the time of arrival and the data bits is performed from these under-sampled measurements via an efficient quadratic program. The receiver does not use wideband analog delay lines, and is robust to large timing uncertainty, hence the transmitter need not waste power on explicit training headers for timing synchronization. Moreover, the receiver can operate in a regime of heavy inter-symbol interference (ISI), and allows a very high baud rate (close to the Nyquist rate). Its performance is shown to remain close to the maximum likelihood receiver under every scenario of under-sampling, timing uncertainty, ISI, and delay spread.**

I. Introduction

Ultra-Wideband (UWB) radio [1] [2] [3], on account of its ability to trade bandwidth for a reduced transmit power, is widely regarded to be a promising candidate for the physical layer of power constrained multi-hop applications like wireless sensor networks. However, in such applications the traffic is often bursty with a low duty cycle, which implies that there is a large timing uncertainty at the start of each burst. The usual approach of using long training headers for accurate timing acquisition is not satisfactory due to its excessive overhead.

Another vexing problem is that typical indoor UWB channels have numerous multi-path components and a large temporal dispersion of $10-100$ nanoseconds [4]. The use of a high baud rate (close to the Nyquist rate) on such channels leads to heavy Inter-Symbol Interference (ISI). Maximum likelihood sequential estimation (MLSE) of such signals is prohibitively complex, and, on account of the large bandwidth, requires an expensive and power-hungry high-speed analog-to-digital (A/D) converter [3]. One pragmatic alternative is to avoid ISI all-together by using a sufficiently low baud-rate so that the MLSE simplifies to a matched filter (MF) that can be implemented as a maximum ratio combining (MRC) rake [3]. Unfortunately, the price paid is a low instantaneous data rate and long channel occupancy. Even non-coherent approaches like energy detecting (ED) receivers [5], transmit-reference (TR) receivers [6], and differential transmit-reference (DTR) receivers [7], while being robust to timing uncertainty, become significantly complicated in the regime of strong ISI [8], and hence usually remain restricted to low baud rates. Moreover, ED suffers a very large SNR penalty, as do TR and DTR to a lesser extent. TR/DTR also involve the use of very long analog delay lines, which are difficult to implement with the requisite accuracy.

In this paper we propose a novel receiver that combines the advantages of MLSE coherent receivers and non-coherent receivers, while avoiding their respective drawbacks. We propose a flexible and robust receiver architecture that performs a 'joint' decoding of the timing and amplitude information. This joint decoding is inspired by the principle of compressed sensing (CS) proposed by [9] [10]. The uncertainty in the arrival time of each burst is treated as 'sparsity' in the classical sense of [9] [10] and therefore tackled automatically in the reconstruction process. Furthermore, the fact that the amplitudes are antipodal $\{+1, -1\}$, and hence the overall transmit signal belongs to a relatively small discrete set rather than being a generic real-valued ultrawide-band signal, is also exploited by the reconstruction process. Hence, the architecture completely bypasses the requirement of high-rate A/D conversion. Instead we use an analog front-end consisting of a bank of correlators with tractable test functions, a low-rate A/D converter, and a DSP back-end based on a computationally efficient quadratic program (QP). The number of correlators can be significantly smaller than the requirement suggested by the Shannon-Nyquist sampling theorem, and the performance nevertheless degrades gracefully.

The proposed receiver works robustly even in significant ISI, and hence we are not restricted to a low baud-rate. At the same time, its complexity is far smaller than full-fledged MLSE. Moreover, we do not rely on long analog delay lines or any specific modulation format as in TR/DTR. The same architecture can operate with various levels of timing accuracy, ranging from a fraction of the pulse width T_{pulse} to many multiples of T_{pulse}, and in each case a performance close to the MLSE receiver is attained. Furthermore, as the burst size becomes moderately large, the receiver implicitly acquires perfect timing 'on the fly' and hence the penalty associated with timing uncertainty becomes negligibly small. Therefore we can send bursts *without* training headers, and yet attain a power efficiency comparable to genie-aided timing.

CS has been used previously by [11] for mitigation of narrow band interference and by [12] for channel estimation. In an approach analogous to ours, [13] [14] and [15] have used CS for direct detection of IR pulses without using a rake or a digital correlator. However [13] [14] consider the restricted regime of no ISI and perfect timing. Similarly, [15] also assumes an ISI-free regime and requires the solution of a set of linear equations that are often ill-conditioned.

Outline of the paper: In Section II we describe the system model and receiver architecture. In Section III we first formulate and analyze the maximum likelihood (ML) receiver, and then propose signal demodulation via a significantly simpler suboptimal QP optimization. Section IV presents simulation results and Section V makes concluding remarks. *Convention:* With an abuse of notation, $P(x)$ will denote the density or mass function of a random variable X. $U([a,b])$ will denote a uniform distribution over the interval $[a,b]$ of the real line or of integers, depending on the context. When x is a vector, x^T is the transpose, $\|x\|_2$ the L_2-norm (Euclidean length), $\|x\|_1$ the L_1-norm (largest absolute value), and $\|x\|_0$ the number of non-zero elements.

II. System Model and Receiver Architecture

A. Transmitter and Channel

The UWB-IR transmitter consists of three main blocks, namely, a timing block that generates a clock signal at a nominal frequency f_{baud}, a payload block that supplies the information bits, and an IR pulse generator. The baud clock provides the timing for the IR pulses within each burst, as well as the timing for the start of each burst after requisite down-sampling. A total of K pulses are transmitted in each burst after which the transmitter hibernates till the start of the next burst.[1] At the k-th strobe of the clock within a burst, the IR pulse generator sends on the air a pulse $\phi(t)$, amplitude modulated by the bit B^k provided by the payload, drawn equiprobably from $\{+1,-1\}$. The pulse $\phi(t)$ is nominally centered at the frequency f_c with a bandwidth Ω. There is no other RF processing at the transmitter, like heterodyning or filtering, which makes this transmitter very simple, small and inexpensive to build.

The UWB channel is known to be linear dispersive with tens or hundred of resolved multi-path components, depending on the radio environment. In indoor environments the temporal dispersion can often be as large as $\tau_{chan} = 100$ nanoseconds. In our simulations we will use realizations from the channel models CM1 through CM4 postulated in [4] for various indoor scenarios.

For example, consider the Hanning modulated RF pulse of [2] which we used in our simulations, with $f_c = 4.0$ GHz and a 6-dB bandwidth $\Omega = 2.0$ GHz. The pulse duration is small, $T_{pulse} = 1.0$ nanosecond. With $\tau_{chan} = 100$ nanoseconds, a conventional UWB-IR system needs to choose a small baud-rate, $f_{baud} \leq 1/\tau_{chan} = 10$ Mbaud, in order to avoid ISI. Our receiver, on the other hand, can tolerate significant ISI and therefore we may choose a baud-rate close to the Nyquist frequency, say $f_{baud} = f_{nyquist}/8 = 500$ Mbaud. Hence the interval between consecutive pulses is $T_{baud} \doteq 1/f_{baud} = 2.0$ nanoseconds, and a burst of $K = 64$ bits (say) will therefore last for 127 nanoseconds. In contrast, the interval between consecutive bursts may be as large as $T_{burst} = 100$ microseconds.

[1] It is often the case that a simple $(1, N_f)$ repetition code is used whose effect is simply to improve the BER vs SNR characteristic by a factor $10\log_{10}(N_f)$ dB, at the cost of a $\frac{1}{N_f}$ rate reduction. Since our receiver architecture remains applicable to the repetition coded scenario without any modification [16], for clarity we will assume $N_f = 1$ in the rest of this paper.

Since a practical inexpensive clock has a timing drift of $\rho \sim 40$ parts per million (p.p.m.) caused by random frequency modulation [17], the total drift from the beginning to the end of a burst is limited to $K\rho f_{baud} = 5.1$ picoseconds, which is negligible. On the other hand, the drift from one burst to the the next is very large, ~ 4.0 nanoseconds. Even with a coarse timing algorithm for predicting the start of the bursts, like a second-order tracking loop, a residual tracking error of the order of 1.0 nanosecond is unavoidable.

Without loss of generality we can concentrate on the reception of a single burst, and treat the estimated time of arrival (TOA) of that burst as the temporal origin, $t = 0$. The residual error of the coarse timing block is then perceived as a late arrival of the actual burst by an amount υ seconds. (By prefixing a sufficient guard interval in the coarse timing estimate, we can ensure that the true arrival can only be late but never early.) For simplicity suppose that the true TOA υ is distributed over the interval $[0,\gamma]$ according to a uniform density. From the point of view of the receiver, the output of the transmitter during the burst is then written as $S(t) = \sum_{k=0}^{K-1} B^k\,\phi(t - kT_{baud} - \upsilon)$.

Notice that in writing this equation we ignore the timing drift within a burst, since, as remarked earlier, it is negligible when f_{baud} is high. It is noteworthy, however, that our setup also subsumes the case of a low baud rate $f_{baud} \leq \frac{1}{\tau_{chan}}$, if we choose $K = 1$ (one pulse in each burst), treat the pulse-to-pulse drift as the burst-to-burst drift υ, and demodulate each pulse independently.

B. Receiver

The receiver consists of an analog front-end and a DSP back-end. The defining characteristic of our receiver is that we relieve the analog front-end of difficult tasks like fast A/D conversion and accurate delay lines, and instead compensate by using an elaborate but tractable DSP back-end.

1) Analog Front-end: Let the received signal at the antenna be denoted by $U(t)$. The first block in the analog front-end is a noise-limiting bandpass-pass filter $g(\cdot)$ centered at f_c, having noise equivalent bandwidth $\approx \Omega$. Its output is

$$R(t) = \sum_{k=0}^{K-1} B^k h(t - kT_{baud} - \upsilon) + W(t), \qquad (1)$$

where $h(t)$ denotes the total impulse response, which is the convolution of the transmit pulse $\phi(t)$, the channel $h_c(t)$, and the filter response $g(t)$, and $W(t) = \int V(t-\tau)g(\tau)d\tau$ is band-limited zero-mean additive Gaussian noise, modeled as the response of the filter to a white Gaussian thermal noise process $V(t)$ of power spectral density N_0. The SNR *per bit* at the output of a *hypothetical* MF is given by $\text{SNR} = \frac{\|h(t)\|_2^2}{N_0/2}$. Let λ_h denote the length of the total impulse response $h(t)$.

The signal $R(t)$ is fed to a bank of M parallel analog correlators, followed by M integrators. This module replaces other conventional structures like a rake receiver or a fast A/D converter. The test function used in correlator number m is denoted as $\psi_m(t)$, and the whole ensemble of test functions is denoted by $\{\psi_m(t)\}$. In Section III-C, we will discuss the criteria for selecting the ensemble. At this point, it suffices

to note that we do not need to tune the timing of these test functions (i.e. no analog delay lines), and hence they are relatively easy to implement. All we require is that the ensemble be known to the DSP back-end.

The integrators $m = 0, 1, \ldots, M-1$ are reset to zero at the epoch $t = 0$ and their output is sampled synchronously at the epoch $\lambda_h + \gamma + (K-1)T_{baud}$ when all of the energy of the burst is known to have arrived with high probability. Thus, for $m = 0, 1, \ldots, M-1$, we have the M *measurements*

$$Y_m = \int_0^{\lambda_h+\gamma+(K-1)T_{baud}} R(t)\psi_m(t)\,dt. \tag{2}$$

The vector of measurements $Y = [Y_1, Y_2, \ldots, Y_M]^T$ is then fed to the DSP back-end, which recovers the payload bits $B^k,\ k = 0, 1 \ldots, K-1$ via a tractable QP algorithm.

2) DSP Back-end: The demodulation of the payload by the DSP back-end relies on a consistent discrete time representation of the signal. Let f_s be a sufficiently large *virtual sampling frequency* [13] for the received UWB-IR signal. We would like to emphasize that this is only a 'thought-experiment' construction, and no A/D conversion is done at rate f_s in actuality. Let $T_s \doteq 1/f_s$, and define $h[n] \doteq h(nT_s),\ n = 0, 1, \ldots, \Lambda_h - 1$ and $h \doteq [h[0], h[1], \ldots, h[\Lambda_h - 1]^T$, where $\Lambda_h = \lceil \lambda_h f_s \rceil$ is the length of the discrete-time finite impulse response $h[n]$. A similar convention will apply to other signals like $g(t), \psi_m(t), W(t)$ etc. Let γ and T_{baud} be multiples of T_s, which can be achieved by construction. Now, expressed in rate f_s samples, the TOA uncertainty is $\Gamma \doteq \gamma f_s$ and the baud period (the interval between consecutive pulses) is $N_{baud} = T_{baud} f_s$. Define $\Lambda_X \doteq \Gamma + (K-1)N_{baud}$. Then the length of the total burst response including the timing uncertainty is $N \doteq \Lambda_h + \Lambda_X - 1$.

Let $\Upsilon = \text{round}(\upsilon f_s)$ be the burst TOA υ quantized to a step size of T_s. (The quantization error is negligible provided f_s is chosen large enough.) Now, the sampled version of $R(t)$ can be written as a vector $R \in \mathbb{R}^N$ given by

$$R = \mathcal{H}X + W. \tag{3}$$

Here the vector $X \in \mathbb{R}^{\Lambda_X}$ is a *virtual* discrete time information signal which has all samples equal to zero except for K non-zero samples. The k-th non-zero sample, for $k = 0, 1, \ldots, K-1$, has a random amplitude B^k drawn independently and equiprobably from $\{-1, +1\}$, and has a random location $\Lambda^k = \Upsilon + kN_{baud}$. On account of the modeling assumption made in Section II-A, it follows that $\Upsilon \sim U([0, \Gamma])$. The vector $W \in \mathbb{R}^N$ is the sampled version of the additive Gaussian noise $W(t)$, and the matrix $\mathcal{H} \in \mathbb{R}^{N \times \Lambda_X}$ is the convolutional matrix (Toeplitz form) of $h[n]$ [18]. In a similar vein we can further relate the actually sampled measurements Y at the output of the integrators to the virtual information signal X. Define the $M \times N$ *measurement matrix*

$$\Psi \doteq \frac{1}{f_s}\Big[\psi_0, \psi_1, \ldots \psi_{M-1}\Big]^T, \tag{4}$$

where $\psi_i \doteq \Big[\psi_i[0], \psi_i[1], \ldots, \psi_i[N-1]\Big]^T,\ \forall i = 0, 1, \ldots, M-1$. Then we can write the *measurement equation*

$$Y = \Psi R = \Psi\mathcal{H}X + \Psi W. \tag{5}$$

Let $B \doteq [B^0, B^1, \ldots, B^{K-1}]^T$. Then the aim of the DSP back-end is to optimally estimate B, Υ from the measurement Y, based on the relation in equation (5) and the a-priori statistical knowledge about B, Υ. Note that B contains the payload which is of primary interest, while the quantity Υ is a 'nuisance' parameter.

III. Bit Demodulation Based on Incomplete Measurements

The maximum likelihood (ML) demodulation of B, based on the measurement Y given by equation (5), will be described in Section III-A. In Section III-B we propose an alternative computationally efficient reconstruction via a QP. Both these algorithms assume that the total system response h is known, up to a random TOA υ. This can be achieved by a suitable identification algorithm as discussed in Section III-D.

A. ML Demodulation and BER Analysis

Let us define the set $\mathcal{X}$ as the set of all signals $x \in \mathbb{R}^{\Lambda_X}$ that satisfy the following properties:

1) $\|x\|_0 = K$ (burst length).
2) The first nonzero sample is located at $\ell^0 \in [0, \Gamma]$. The subsequent non-zero samples are located at positions $\ell^k = \ell^0 + kN_{baud},\ \forall k = 1, 2, \ldots, K-1$ (timing).
3) The magnitudes of all the nonzero samples are from $\{-1, +1\}$ (signaling alphabet).

Clearly $|\mathcal{X}| = 2^K(\Gamma+1)$, $\mathcal{X}$ is the finite equiprobable alphabet of the random information signal X (cf. Section II-B2), and there is a one-to-one mapping

$$\{-1, +1\}^K \times \{0, 1, \ldots, \Gamma\} \rightarrow \mathcal{X} \tag{6}$$

$$(B, \Upsilon) \mapsto X(B, \Upsilon). \tag{7}$$

Hence we can write $P(Y|B, \Upsilon) = P(Y|X)$, which implies that, without losing optimality, we may first make the ML estimate $\hat{X}$ of the information signal X, and then map it to the optimal estimates of the payload $\hat{B}(\hat{X})$ and TOA $\hat{\Upsilon}(\hat{X})$.

Let $\sigma^2 = \frac{N_0}{2f_s}$. It is easily shown that the likelihood of a candidate signal $x \in \mathcal{X}$ conditioned on the observation Y is given, up to a normalization factor, by

$$P(Y|x) \propto \exp\left\{\frac{-1}{2\sigma^2}(Y - \Psi\mathcal{H}x)^T(\Psi\mathcal{G}\mathcal{G}^T\Psi^T)^{-1}(Y - \Psi\mathcal{H}x)\right\}. \tag{8}$$

Hence, the ML demodulator declares the estimated signal as

$$\hat{X} = \underset{x \in \mathcal{X}}{\text{argmin}}\ (Y - \Psi\mathcal{H}x)^T(\Psi\mathcal{G}\mathcal{G}^T\Psi^T)^{-1}(Y - \Psi\mathcal{H}x). \tag{9}$$

Consequently we have the union bound on the BER [16]

$$P_e \leq \frac{1}{2^K(\Gamma+1)} \sum_{x^1, x^0 \in \mathcal{X}} \frac{\|\hat{B}(x^0) - \hat{B}(x^1)\|_0}{K} \mathbf{Q}\left(\frac{\sqrt{(x^0 - x^1)^T\mathcal{H}^T\Psi^T(\Psi\mathcal{G}\mathcal{G}^T\Psi^T)^{-1}\Psi\mathcal{H}(x^0 - x^1)}}{2\sigma}\right), \tag{10}$$

where $\mathbf{Q}(a) = \int_a^\infty \frac{1}{\sqrt{2\pi}} \exp\left\{\frac{-x^2}{2}\right\} dx$. Notice that if (i) there is no under-sampling (i.e. $M = N$), (ii) Ψ and G are invertible, (iii) the timing is ideal (i.e. $\Gamma = 0$), and (iv) there is only one bit per burst and hence no ISI (i.e. $K = 1$), expression (10) yields $P_e = \mathbf{Q}\left(\sqrt{\text{SNR}}\right)$, the familiar expression for a perfectly timed MF.

B. Suboptimal Computationally Efficient Demodulation

The complexity of the ML demodulation problem (9) scales as $2^{\min(K,\lambda_h f_{baud})}$), and is clearly intractable when K and $\lambda_h f_{baud}$ are even moderately large. Hence we will now propose an alternative suboptimal demodulation technique whose complexity is $O(K^3)$. Let the vector $\xi(a,\ell_1,\ell_2)$ be a positive penalty vector for the candidate information signals $x \in \mathcal{X}$. It incorporates the available timing information by giving more penalty to those locations of x where the occurrence of the non-zero samples is unlikely. That is, for all $n = 0, 1, \ldots, \Lambda_X - 1$,

$$\xi(a,\ell_1,\ell_2)[n] \doteq \begin{cases} & n = \ell + kN_{baud} \\ 1.0, & \ell \in [a+\ell_1, a+\ell_2] \\ & k \in \{0,1,\ldots,K-1\} \\ \mho, & \text{otherwise,} \end{cases} \tag{11}$$

where $\mho$ is some suitable large number like 10^3. Also define a corresponding diagonal penalty matrix as $\Xi(a,\ell_1,\ell_2) = diag(\xi(a,\ell_1,\ell_2))$. Now consider the following relaxation of the ML demodulation problem of equation (9):

$$\tilde{X} = \underset{x\in\mathbb{R}^N : \|\Xi(a,\ell_1,\ell_2)\, x\|_1 = K}{\operatorname{argmin}} (Y - \Psi\mathcal{H}x)^T (\Psi\mathcal{G}\mathcal{G}^T\Psi^T)^{-1} (Y - \Psi\mathcal{H}x). \tag{12}$$

Notice that the new constraint set is not a discrete set, but rather a continuous set of signals of adequately small L_1 norm. Therefore notice that $\mathcal{X} \subset \{x \in \mathbb{R}^N : \|\Xi(0,0,\Gamma)\, x\|_1 = K\}$. With some manipulation we can re-write (12) as [9]

$$\begin{aligned} \tilde{X}_n &= \tilde{Z}_n - \tilde{Z}_{n+N},\ n = 0,1,2,\ldots,N, \\ \tilde{Z} &= \begin{array}{l} \min f^T z + \frac{1}{2} z^T Q z \\ z \geq 0,\ [\xi(a,\ell_1,\ell_2)^T, \xi(a,\ell_1,\ell_2)^T] z = K, \end{array} \end{aligned} \tag{13}$$

$$Q = \begin{pmatrix} \mathcal{H}^T\Psi^T(\Psi\mathcal{G}\mathcal{G}^T\Psi^T)^{-1}\Psi\mathcal{H} & -\mathcal{H}^T\Psi^T(\Psi\mathcal{G}\mathcal{G}^T\Psi^T)^{-1}\Psi\mathcal{H} \\ -\mathcal{H}^T\Psi^T(\Psi\mathcal{G}\mathcal{G}^T\Psi^T)^{-1}\Psi\mathcal{H} & \mathcal{H}^T\Psi^T(\Psi\mathcal{G}\mathcal{G}^T\Psi^T)^{-1}\Psi\mathcal{H} \end{pmatrix}, \tag{14}$$

$$f = [-Y^T(\Psi\mathcal{G}\mathcal{G}^T\Psi^T)^{-1}\Psi\mathcal{H}, Y^T(\Psi\mathcal{G}\mathcal{G}^T\Psi^T)^{-1}\Psi\mathcal{H}]. \tag{15}$$

(13) is now a standard QP, which has several efficient large-scale techniques of solution, of which $O(K^3)$ interior point methods are generally the fastest [19]. We perform the demodulation in two stages. In the first stage we solve the QP in (13) using $\xi(a=0,\ell_1=0,\ell_2=\Gamma)$. The result of this stage, $\tilde{X}^{(1)}$, is used to extract an estimate $\hat{\Upsilon}$ of the arrival time via correlation with the template $\xi(0,0,0)[n]$ as follows:

$$\hat{\Upsilon} = \underset{n'\in\{0,1,\ldots,\Gamma\}}{\operatorname{argmax}} \sum_n |\tilde{X}^{(1)}[n-n']|\ \xi(0,0,0)[n]. \tag{16}$$

We then solve the QP in (13) again, using $\xi(a=\hat{\Upsilon},\ell_1=0,\ell_2=0)$. From the result $\tilde{X}^{(2)}$, we demodulate the payload

$$\hat{B}^k = sign(\tilde{X}^{(2)}[\hat{\Upsilon} + k\, N_{baud}]),\ k = 0,1,\ldots,K-1. \tag{17}$$

In summary, in lieu of the ML demodulation problem, which involves minimization over a large discrete set $\mathcal{X}$, we have formulated a relaxed continuous QP which jointly solves for the best payload and timing *without explicitly checking each timing epoch and bit pattern individually*. We would like to point out that our reconstruction differs from he constrained L_1-minimization used in classical noiseless CS [9], [10], as well as the regularized optimizations proposed for the noisy setting like the LASSO [20]. These methods exploit only generic sparsity $\|X\|_0 \leq K$, while in our application we know that $\|X\|_0 = K$ *exactly*. In fact once the location of the first non-zero sample in X is known, the locations of the others are also *fixed*, and in this sense the sparsity of X is not K but only 1. Moreover, the non-zero samples are from a known fixed (antipodal) alphabet and the measurement noise is not white and its covariance matrix is known. This allows us to recast the LASSO as a QP, whose performance is *invariant* w.r.t. the burst length K provided the number of front-end correlators M, scales *linearly* with K.

C. Choice of Measurement Ensemble

First it is worth noting that if we set set $M = N$, idealize the front-end filter be a low-pass Nyquist filter of bandwidth $f_s/2$ so that $g(t) = \frac{\sin \pi t f_s}{\pi t f_s}$, and let the test functions be the canonical functions $\psi_m(t) = \delta(t - mT_s)$, $m = 0,1,\ldots,M$, the correlator bank simply provides uniform samples of the incoming signal, i.e. an A/D conversion, which constitute sufficient statistics. Of course, if $M < N$, the choice of ensemble does become critical. This is because if the energy of the signal happens to fall in the null space of the ensemble with high probability, there is no hope of reconstructing the signal. Is it possible to devise universal ensembles that allow reliable reconstruction of any under-sampled signal provided the under-sampling is not too severe? Moreover, can they allow a graceful SNR penalty in the presence of receiver noise? The surprising answer to the first question is known to be in the affirmative [9], [10]. Our simulations in Section IV indicate that the answer to the second also seems to be affirmative.

The 'universal' measurement ensembles were shown to be sets of randomly generated noise-like signals [9], [10]. One example is that of binary pseudo-noise (PN) signals that transit independently and equiprobably between levels $\{\frac{+1}{\sqrt{N}}, \frac{-1}{\sqrt{N}}\}$ at intervals of T_s seconds. However, we actually do not need such strict universality since we known that our signal sparsity is always in the temporal domain [21]. The ensemble of *square waves* of amplitude $1/\sqrt{N}$ and frequencies selected *deterministically and uniformly* from the signal band $[f_c - \frac{\Omega}{2}, f_c + \frac{\Omega}{2}]$ is also seen to perform as well as the PN ensemble.

If we idealize $g(t)$ to be an ideal band-pass Nyquist filter of bandwidth Ω, so that $\mathcal{G}\mathcal{G}^T = I$, and pretend that the limiting property $\Psi\Psi^T = I$ holds, we have $\Psi\mathcal{G}\mathcal{G}^T\Psi^T = I$. Hence no matter which M measurement signals we choose from the underlying ensemble, we are very likely to capture roughly a fraction $\frac{Mf_s}{2\Omega N}$ of the received signal's energy. This implies that reliable demodulation of the UWB-IR signal is possible after paying an under-sampling penalty of roughly $10\log_{10}\frac{2\Omega N}{Mf_s}$ dB, and this penalty will (on an average) decrease monotonically and vanish as $M \uparrow N\frac{2\Omega}{f_s}$. Finally, note that CS also has an important robustness property that the generated test functions do not need to have an ideal waveform. For example the PN or square wave ensemble need not have rectangular level transitions. Imperfections like ringing and non-ideal rise time are well-tolerated, provided we know these effects in advance so that we can compensate for them by choosing an appropriately modified Ψ in the reconstruction algorithm.

D. Channel Identification

The demodulation algorithms presented in Sections III-A and III-B assume that the total system response h is known, up to a random TOA v. This can be easily achieved by an incremental channel identification algorithm based on maximum likelihood principles, as described in [16]. In particular note that the algorithm uses as inputs the same CS measurements Y that are available for bit-demodulation, and hence no additional analog hardware is needed. Similarly, its computational overhead is also negligible.

IV. Simulations

Let 'CS-ML' denote a receiver having the CS analog front-end and ML demodulation in the DSP back-end (cf. Section III-A) and let 'CS-QP' denote a receiver having the CS analog front-end and a QP demodulation in the back-end (cf. Section III-B). The performance of CS-ML is always an achievable lower bound with which we will compare the practical CS-QP receiver. Let 'Genie-MF' denote a receiver implementing a perfectly timed matched filter in an ISI-free environment. It gives an ultimate lower bound, but it is not necessarily achievable in ISI, timing uncertainty or under-sampling. Our discussion is divided into two parts. First, in Section IV-A, we will investigate the effect of incomplete measurements, timing error and burst length on the BER of CS-ML and CS-QP receivers. Then in Section IV-B we will demonstrate robustness to channels models and their random realizations. All simulations were performed with $f_s = 10$ GHz, $f_{baud} = 500$ Mbaud, and the IR pulse described in Section II-A. The measurement ensemble used was the square wave ensemble described in Section III-C. In the following, the quantity $\frac{Mf_s}{2\alpha\Omega N}$ will be called the *under-sampling factor*. To ensure that $\frac{Mf_s}{2\alpha\Omega N} = 1.0$ (*adequate sampling*) achieves a performance indistinguishable form MLSE under Nyquist rate samples, we empirically set $\alpha = 1.5$. (If $\phi(t)$ was strictly band-limited we would use $\alpha = 1.0$).

A. Effect of Under-Sampling, Timing Uncertainty and Multiple Interfering Symbols

Consider Figure 1 which shows, for a fixed CM1 channel, the effect of under-sampling (via M), timing uncertainty (via Γ) and the burst length K. Figures 1(a),(b) correspond to $\frac{Mf_s}{2\alpha\Omega N} = 1.0, 0.25$ under ideal timing $\Gamma = 0$, and Figures 1(c),(d) correspond to $\frac{Mf_s}{2\alpha\Omega N} = 1.0, 0.25$ under uncertain timing $\Gamma = 10$. In each sub-figure we simulate CS-QP with $K = 1, 2, 4, 8, 16$ bits per burst and plot it with dashed lines with circle markers. We plot with solid blue lines the analytical performance of CS-ML given by equation (10), for $K = 1, 2, 4, 8$. We do not give CS-ML curve for $K = 16$ because it is intractable. The figure allows the following interesting observations:

(i) We see that with ideal timing $\Gamma = 0$ and various amounts of under-sampling in sub-plots (a),(b), the CS-QP receiver performance is very close to CS-ML, for all K. This demonstrates that we can indeed recover the performance of an ideal coherent receiver. Furthermore with adequate sampling, all the CS-ML and CS-QP curves for various K coincide with the Genie-MF curve, implying that there is negligible loss due to the ISI. With under-sampling $\frac{Mf_s}{2\alpha\Omega N} = 0.25$, the curves of CS-QP and CS-ML for all K stay bunched together and have a consistent penalty of about 6.0 dB, as predicted in Section III-C.

(ii) Even with non-ideal timing $\Gamma = 10$, the CS-QP receiver performance is reasonably close to CS-ML, for each K. Note that now even in the adequate sampling case, the $K = 1$ curve of CS-ML suffers a penalty of ~ 7.0 dB w.r.t. the the corresponding curve of ideal timing from sub-plot (a). While this penalty is big, it is not catastrophic like the rake receiver which suffers a loss of 20 dB. More interestingly, as the number of bits in a burst K increases, the CS-ML curves start paring the loss and approach the ideal timing curve.

(iii) Finally notice that in sub-plot (d), where we have both under-sampling as well as non-ideal timing, the penalty suffered by the CS-ML receiver is approximately the *additive composition* of the two individual penalties, and this is seen to consistently hold for all K. The CS-QP performance is also seen to mimic this behavior.

In summary, with CS-ML as well as CS-QP, the effects of under-sampling and timing uncertainty are approximately de-coupled. The loss due to under-sampling is consistently $10\log_{10}\frac{Mf_s}{2\alpha\Omega N}$ dB. For lossless sampling we need $M = \frac{2\alpha\Omega}{f_s}N$ projections. Timing uncertainty can be combated by using sufficiently many bits per burst, and the associated penalty can be practically eliminated.

B. Robustness to Stochasticity of Channel Realizations

In the preceding discussion, we used one fixed realization from the CM1 channel model. Now we will study the effect of the stochasticity of the channel realizations and variations in the channel models. In Figure 2(a)-(d) we draw six random realizations from each model CM1 through CM4, and respectively plot the BER-vs-SNR characteristic of the various receivers CS-QP, CS-ML and Genie-MF. In all cases we use a *constant number of projections* $M = 128$ which corresponds to a significant amount of under sampling ranging from $\frac{Mf_s}{2\alpha\Omega N} = 0.15$ to 0.25, depending of the channel model and realization. We set $\Gamma = 10$ and $K = 8$.

First notice that, for every channel model, the CS-ML and CS-QP receiver performance does not vary by more than $2-3$ dB no matter what the realization of the channel is. This demonstrates the *universality* of the ensemble used in the analog front-end even in the under-sampled case. Obviously if we used adequate sampling (large enough M), all the curves would bunch together with no appreciable variation. The loss of CS-QP relative to CS-ML is the result of the sub-optimality of CS-QP and is in line with the loss observed in Figure 1(d). We can thus conclude that the proposed receiver is indeed very robust to various channel models and their realizations.

V. Concluding Remarks

We have proposed a novel receiver for UWB Impulse Radio transmission based on the principle of Compressed Sensing. It is very robust to timing uncertainty, ISI and under-sampling, and performs consistently close to an ML receiver. It allows the use of baud-rates comparable to the Nyquist rate,

and hence large network loading factors. The demodulation procedure is insensitive to the nature of the multi-path channel (CM1, CM2 etc). It is thus ideally suited to low-power indoor applications with bursty traffic, like wireless sensor networks.

REFERENCES

[1] M.Z. Win and R.A. Scholtz. Impulse Radio: How It Works. *IEEE Commun. Letters*, 2(2):36–38, February 1998.

[2] S. Roy, J.R. Foerster, V.S. Somayazulu, and D.G. Leeper. Ultrawideband Radio Design: The Promise of High-Speed, Short-Range Wireless Connectivity. *Proc. of the IEEE*, 92(2):295–311, February 2004.

[3] H. Arslan, Z. N. Chen, and M Di Benedetto. *Ultra Wideband Wireless Communication*. John Wiley & Sons, Inc., 2006.

[4] A. F. Molisch, D. Cassioli, C. C. Chong, S. Emami, A. Fort, B. Kannan, J. Karedal, J. Kunisch, H. G. Schantz, K. Siwiak, and M. Z. Win. A Comprehensive Standardized Model for Ultrawideband Propagation Channels. *IEEE Trans. on Antennas and Propagation*, 54(11):3151–3166, November 2006.

[5] A. A. D'Amico, U. Mengali, and E. Arias de Reyna. Energy-Detection UWB Receivers with Multiple Energy Measurements. *IEEE Trans. Wireless Commun.*, 6(7):2652–2659, July 2007.

[6] Yi-Ling Chao and R. A. Scholtz. Ultra-wideband transmitted reference systems. *IEEE Trans. Veh. Technol.*, 54(5):1556–1569, September 2005.

[7] Yi-Ling Chao and R. A. Scholtz. Optimal and Suboptimal Receivers for Ultra-Wideband Transmitted Reference Systems. *Proc. IEEE Global Telecom. Conf. (GLOBECOM)*, 2:759–763 Vol.2, December 2003.

[8] M. Pausini, G.J.M. Janssen, and K. Witrisal. Performance Enhancement of Differential UWB Autocorrelation Receivers Under ISI. *IEEE Journal on Selected Areas in Communications*, 24(4):815–821, April 2006.

[9] D. L. Donoho. Compressed Sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, April 2006.

[10] E. J. Candes and T. Tao. Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies? *IEEE Trans. Inform. Theory*, 52(12):5406–5425, December 2006.

[11] Z. Wang, G. R. Arce, B. M. Sadler, J. L. Paredes, S. Hoyos, and Z. Yu. Compressed UWB Signal Detection with Narrowband Interference Mitigation. *IEEE Int. Conf. on UWB*, 2:157–160, September 2008.

[12] J.L. Paredes, G. R. Arce, and Zhongmin Wang. Ultra-Wideband Compressed Sensing: Channel Estimation. *IEEE J. of Select. Topics in Signal Processing*, 1(3):383–395, October 2007.

[13] Z. Wang, G. R. Arce, B. M. Sadler, J. L. Paredes, and X. Ma. Compressed Detection for Pilot Assisted Ultra-Wideband Impulse Radio. *IEEE Int. Conf. on UWB*, pages 393–398, September 2007.

[14] Z. Wang, G. R. Arce, J. L. Paredes, and B. M. Sadler. Compressed Detection for Ultra-Wideband Impulse Radio. *IEEE Workshop on Sig. Proc. Advances in Wireless Communications*, pages 1–5, June 2007.

[15] J. Kusuma, I. Maravic, and M. Vetterli. Sampling with Finite Rate of Innovation: Channel and Timing Estimation for UWB and GPS. *Proc. IEEE Int. Conf. Commun. (ICC)*, 5:3540–3544 vol.5, May 2003.

[16] Anand Oka and Lutz Lampe. A Compressed Sensing Receiver for UWB Impulse Radio Communication in Wireless Sensor Networks. Submitted to the Elsevier - Physical Communications (Special Issue on Advances in Ultra-Wideband Wireless Communication). Preprint available at www.ece.ubc.ca/∼anando/.

[17] IEEE 802.15.4a-2007. Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs); Amendment 1: Add Alternate PHYs, March 2007.

[18] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 2005.

[19] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[20] R. Tibshirani. Regression Shrinkage and Selection via the Lasso. *J. of the Roy. Stat. Soc.*, 58 - Series B(1):267–288, 1996.

[21] E. J. Candes, J. Romberg, and T. Tao. Robust Uncertainty Principles: Exact Signal Reconstruction From Highly Incomplete Frequency Information. *IEEE Trans. Inform. Theory*, 52(2):489–509, February 2006.

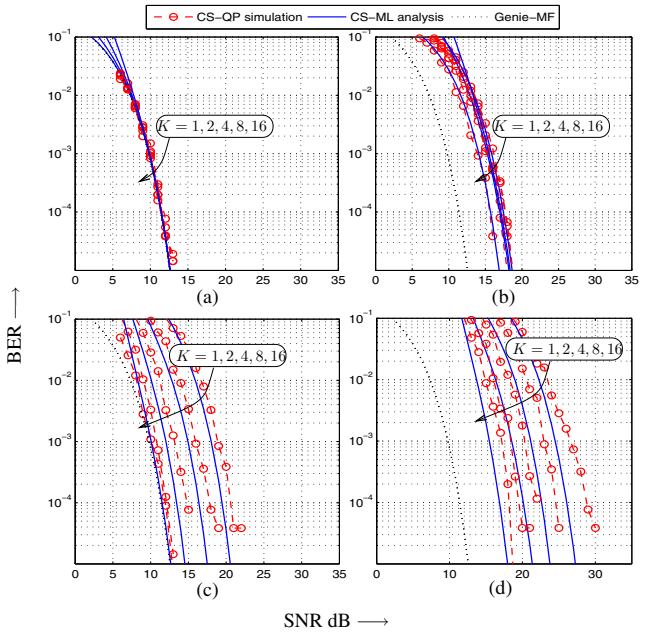

Fig. 1. Effect of under-sampling, timing uncertainty and burst length on the receiver performance. (a),(b): $\frac{Mf_s}{2\alpha\Omega N} = 1.0, 0.25$ and $\Gamma = 0$. (c),(d): $\frac{Mf_s}{2\alpha\Omega N} = 1.0, 0.25$ and $\Gamma = 10$. In each sub-figure $K = 1, 2, 4, 8, 16$ (CS-QP) and $K = 1, 2, 4, 8$ (CS-ML).

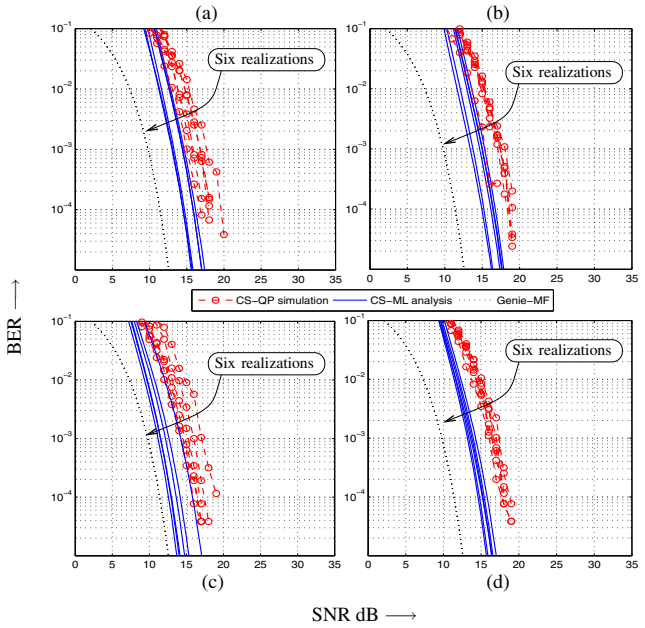

Fig. 2. Robustness to stochastic channel realizations. Sub-plots (a)- (d) correspond to channel models CM1-CM4 respectively. The BER vs SNR characteristics of six stochastic realizations are provided. $M = 128, \Gamma = 10$ and $K = 8$.