Frequency domain processing of ultra-wideband signals

Robert D. Weaver
Communication Sciences Institute
University of Southern California
Mail Code 2565 / EEB500
Los Angeles, California 90089-2565

***Abstract*— The central challenge of ultra-wideband (UWB) radio is to overcome its intrinsic complexity. This paper argues for the use of high-speed sampling and a front-end FFT as the paradigm of choice for robust performance and low energy consumption in certain ultra-wideband designs. Assuming the need to approach optimal performance, it highlights the complexity advantage of frequency domain (FD) over time domain processing. Anechoic chamber and realistically propagated indoor UWB measurements, combined with simulated noise and interference, are used to illustrate FD techniques for initial synchronization and estimation of channel response.**

I. Introduction

Absent interference, when the propagation path is short and line-of-sight, ultra-wideband (UWB) radio is simple in concept. For both engineer and layman, that is part of its appeal. In practice however, dense indoor multipath, massive in-band interference by conventional services and the low allowable UWB transmit power all conspire to complicate the implementation of robust UWB systems.

We shall take for granted that the processing required to achieve near-optimal performance in a UWB receiver is not reasonably implementable in analog form. Supported by the knowledge that the all-digital receiver preserves the information required for optimal processing [1], we consider the hypothesis that it is most efficient for a UWB receiver to sample its passband at high speed, perform an FFT on an appropriate time window and then operate entirely in the frequency domain (FD). Although we cannot fully explore here this hypothesis, which will not be true for all applications, it is the goal of this paper to present the basic argument as motivation for further research into UWB receiver architecture and to apply it in the contexts of UWB signal timing acquisition, channel estimation and interference rejection.

We will see that this approach is most appealing for the class of UWB applications involving a relatively narrow post-correlator bandwidth, i.e. long signal integration times. For instance, a UWB positioning system can, in principle, easily track a dynamically moving emitter at received C/N_0 levels of 60 dB-Hz, a number that a Global Positioning System (GPS) receiver designer will immediately recognize as very generous. Indeed, the tracking precision that could be achieved with UWB signals is orders of magnitude better than GPS, simply on account of UWB's much greater bandwidth. However, this begs the question of how the UWB receiver is to acquire such a weak signal to begin with. Consider that the autocorrelation peak of the UWB signal is a fraction of a nanosecond wide, while the search space, due to FCC spectral uniformity requirements, will be at least a few microseconds and the integration time required to achieve theoretical sensitivity is of the order of ten microseconds as well. Measured in hypothetical A/D samples taken at Nyquist rate (several GHz), so that the sample time is comparable to the timing resolution, the search space then consists of tens or even hundreds of thousands of independent initial timing hypotheses. To compound the problem even further, there may be severe narrowband interference to deal with.

By constructing this (hopefully) plausible application scenario, I wish to make the point that the sensitivity of a low-data rate UWB application will be squarely limited by the amount of signal processing horsepower that we are willing to bring to bear for signal acquisition and, particularly relevant to this discussion, the efficiency of the algorithms employed.

II. FD convolution

As is well known, the time domain (TD) convolution operation,

$$r \quad = \quad x \otimes h$$

is equivalent to the FD operation

$$\mathsf{R} = \mathsf{X}\mathsf{H} \tag{1}$$

where the functions of frequency R, X and H are the Fourier representations of the time functions r, x and h, respectively, whose domains may be taken as either discrete or continuous, and either finite and periodic or infinite. A digital receiver will necessarily implement the discrete and finite variant, the DFT

$$\begin{aligned} \mathsf{X}(k) &= \mathcal{F}_k\{x(n)\} \\ &\triangleq \sum_{n=0}^{N-1} x(n)\exp(-j\frac{2\pi}{N}kn). \end{aligned}$$

However, DFT convolution using (1) does not serve unconditionally to duplicate the function of a FIR filter, because its output behaves as if time flowed periodically, modulo the size of the DFT, N. In general, this causes a wrap-around effect when convolution is naïvely attempted. The solution of this problem, for signals and system responses of finite length, is zero-filling, as was reported by Helms [2] soon after the invention of the FFT. Rather than using his development, a matrix-based representation of this convolution is more appropriate to what follows.

If the vectors x and h represent two discrete-time signals of finite lengths M and L, respectively, their aperiodic convolution

$$r(n) \quad = \quad \sum_{m=0}^{L-1} x(n-m)h(m)$$

may be represented by the matrix transformation

$$\vec{r} \quad = \quad X\vec{h} \tag{2}$$

where X is a Toeplitz matrix [3].

If $\vec{x}$ and $\vec{h}$ are zero-filled to length $N \geq M+L-1$, X may be made circulant,

$$X = \begin{pmatrix} x_0 & 0 & \dots & x_{M-1} & \dots & x_1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_{M-2} & \dots & x_0 & 0 & \dots & x_{M-1} \\ x_{M-1} & \dots & x_1 & x_0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \dots & x_{M-1} & \dots & x_0 & 0 \\ 0 & \dots & 0 & x_{M-1} & \dots & x_0 \end{pmatrix}$$

with no impact other than zero-filling the result $\vec{r}$. Therefore, we have just implemented an aperiodic convolution of two finite-length sequences as a circular convolution with time treated modulo N. The relationship with the FFT becomes clear when we realize that X, being circulant, may be diagonalized as

$$X = FDF^H \tag{3}$$

where F is the inverse Fourier transform operator matrix whose elements are

$$F_{kn} = \exp(j\frac{2\pi}{N}kn),$$

F^H represents its Hermitian transpose (the Fourier transformation itself), and D is a diagonal matrix whose diagonal elements are $\{\mathsf{X}(k)\}$, the DFT of the sequence $\{x(n)\}$. This leads to $F^H\vec{r} = DF^H\vec{h}$, which is a statement of (1) for DFT's of appropriately zero-filled vectors.

FFT convolution may also be used in segmented fashion if $\vec{h}$ does not have length less than N overall, but this involves truncation of $\vec{r}$ after each segment is processed (again, see [2] and also [4]) and therefore does not lead to the most useful result for our purposes. We will be applying FFT convolution to the acquisition process, which we will treat as a batch process on an isolated signal.

The FD algorithm for a FIR filter of length L, implemented with a power-of-two FFT of length N, is less complex (measured in multiply-accumulate operations, MAC's) than

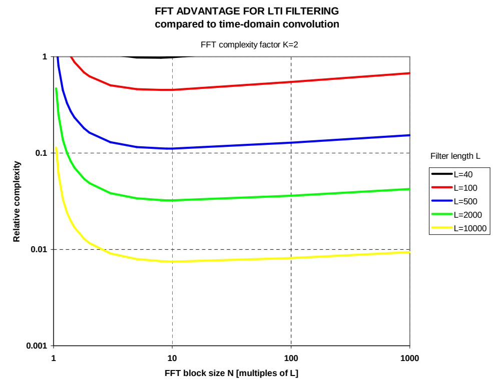

Fig. 1. Complexity advantage of FD over TD convolution

the equivalent TD filter by roughly $O(log_2(N)/L)$. More precisely, the complexity ratio for convolution in general is

$$\frac{2K\log_2(N)+1}{(1-\frac{L-1}{N})L}, \tag{4}$$

where N is the size of the FFT, L is the length of $h(n)$ and $x(n)$ is assumed to be of length less than or equal to $N-L+1$ to allow FFT convolution to work as above. K is a constant factor related to the efficiency of the particular FFT implementation and may be conservatively assumed to be about 2. Figure 1 plots relative complexity vs. L and N. It suggests that FFT convolution is probably not worthwhile for filter lengths less than about 40 samples. However, for lengths of a few hundred samples or more, the FD processing option becomes hard to ignore.

This algorithmic complexity advantage of FFT-based convolution translates directly into lower energy consumption. However, a direct hardware implementation of a large FFT is a complex undertaking, with many trade-offs available between serial, parallel and pipelined processing approaches [5]. Software implementation in a digital signal processor of even large FFT's is a well-developed art [6]. It is beyond the scope of this paper to study any further these important and difficult implementation issues.

III. UWB SIGNAL PROCESSING

Given the need to detect and process a signal from a UWB transmitter with a template pulse waveform known *a priori*, propagating through a poorly known and perhaps changing channel, in a background of thermal noise and multiple nearby narrowband interferers, let us consider the nature of the processing required. Figure 2 shows the UWB pulse waveform used for this study. It was acquired in an anechoic chamber. From this pulse, a simulated randomly time-hopped sequence of 128 pulses was then constructed, with a good impulsive autocorrelation function. This composite waveform, shown in Figure 3, constituted the template waveform for the signal to be acquired.

In similar fashion, simulated received waveforms, some corrupted by simulated noise and interference, were generated

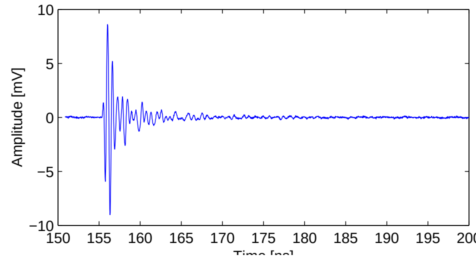

Fig. 2. Anechoic chamber characterization of UWB pulse. Measurement noise was reduced with 256-fold sweep averaging and low-pass filtering.

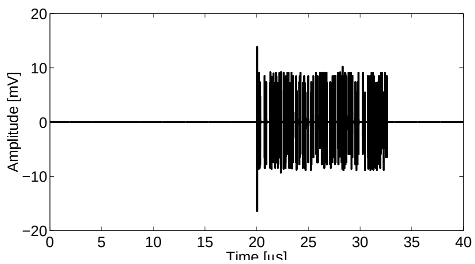

Fig. 3. Synthesized template waveform. It consists of 128 pulses, randomly time-hopped at an average PRF of 10 MHz.

from a measured pulse, realistically propagated in a laboratory environment using the same pulse generator and antennas. The objective was to demonstrate the channel estimation algorithm to be discussed below.

A. Maximum likelihood timing acquisition

The time-average cross-correlation of two real signals r and x being defined in discrete time as

$$\lambda_{rx}(n) \;\; = \;\; \sum_{m} r(m)x(m-n),$$

we find that $\lambda_{rx}(n) = r(n) \otimes x(-n)$, so that the corresponding FD expression is

$$\begin{aligned} \Lambda(k) &= \mathcal{F}_k\{\lambda_{rx}(n)\} \\ &= \mathsf{R}(k)\mathsf{X}^*(k). \end{aligned} \tag{5}$$

Note that, usually, the FFT of x, the template, can be pre-computed. Once we recover $\lambda_{rx}(n)$ with an inverse FFT of $\Lambda(k)$, we have effectively obtained from a matched-filter the likelihood function of the delay of the signal r relative to x, for $N - L$ multiples of the sample interval, where L is the length of the template. An illustration of this matched filtering operation, including the effect of interference rejection, is shown in the following section. For this test, L had a length of $64,000$ samples, while the simulated signal spanned $200,000$ samples. On a standard personal computer, the Matlab software package's FFT routines executed the full correlation in about 0.85 seconds, including template FFT computation. In a dramatic demonstration of the advantage of the FFT correlation approach, an equivalent time-domain correlation of the same signals took 19.8 minutes.

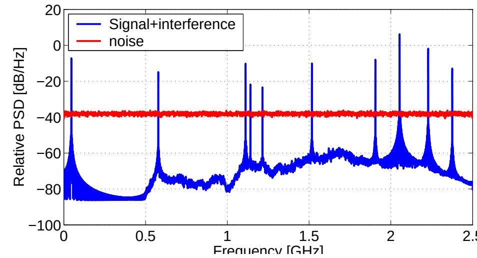

Fig. 4. Visualization of signal, noise and interference in the frequency domain. In this plot, the SNR is 18 dB and J/S is 40 dB. The noise level is shown separately to better depict its relationship to the signal.

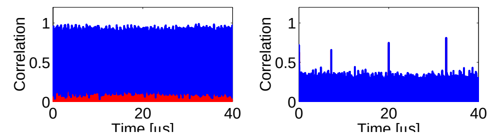

Fig. 5. Correlator output, in the presence of interference, with and without FD interference excision. J/S is 60 dB in both plots. Excision is turned on in the right-hand plot, allowing recovery of the three correlation peaks.

B. Interference rejection

Since the spectrum occupied by a UWB application may overlap other narrowband users, severe interference to the UWB receiver is to be expected. While the UWB transmitter power will typically be limited to a few tens of microwatts [7], the interferers may collectively generate *watts*, and perhaps at a closer range. Thus the jam-to-signal power ratio (J/S) may exceed 60 dB. When we view multiple narrowband interferers in the frequency domain, a collection of large impulses will result, whereas the UWB spectrum is distributed, as simulated in Figure 4. This suggests either notch-filtering or non-linearly clipping the frequencies exceeding a mask. Both alternatives are easily applied to a signal in the frequency domain, at an additional cost of $O(1)$ operation per sample. Closely similar ideas were studied by Milstein and Das, who reported useful results with simple FD-based excision algorithms [8]. For this demonstration, I made no attempt to derive an optimum algorithm. My FD excision method was simply to search for and set to zero all peaks in the FFT exceeding 10 dB above the noise floor, estimated as the average power per FFT cell. Even this simple-minded algorithm did remarkably well, as depicted in the comparison plots of Figure 5.

In the time domain, the cost per sample is proportional to the size of the filter, as previously discussed. Although I did not attempt to synthesize the TD filter, it should be obvious that it must be very selective, to avoid excessive loss of signal energy, and therefore quite complex, especially if there are multiple distinct interferers.

C. Channel estimation

The need to do channel estimation is not obvious. For communication, choosing among M hypotheses in Gaussian noise requires knowledge of the M possible received template waveforms s_i, which may be done by filtering a training

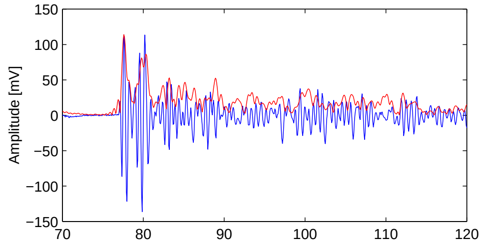

Fig. 6. The same template transmission as in Figure 2, but propagated in an indoor environment. 256-fold sweep averaging was used here as well. The estimated channel response envelope is shown in red.

sequence of such received templates.

$$r_i(n) = s_i(n) + w(n) \tag{6}$$

where

$$s_i(n) = x_i(n) \otimes h(n) \tag{7}$$

However, extracting channel response information from the templates may be appropriate, if we wish to model the channel, either for its own sake or in order to extract a channel model to aid or simplify tracking and later decisions.

Given the measured response vector $r(n)$ and a known transmitted signal $x(n)$, estimation of the channel impulse response in noise $w(n)$ results from solving

$$\vec{r} \quad = \quad X\vec{h} + \vec{w}$$

for $\vec{h}$. For non-singular X and gaussian $\vec{w}$, the least squares solution is

$$\hat{\vec{h}} \quad = \quad X^{-1}\vec{r}.$$

However, in practice, X is always singular. Following in the footsteps of [9] and [10], this is best handled by performing a singular value decomposition (SVD) of X. Once again, the frequency domain approach greatly simplifies this task, since (3) tells us that $\mathsf{X}(k)$ contains the eigenvalues of the matrix X at the frequencies indexed by k.

$X(k)$ is small or zero for frequencies outside the UWB passband (and perhaps also for other frequencies, depending on the details of the pulse sequence). Therefore, in the FD view, a well-behaved least-squares estimate of $\mathsf{H}(k)$ is accomplished very simply with point-by-point division and windowing:

$$\hat{\mathsf{H}}(\mathsf{k}) \quad = \quad \frac{\mathsf{R}(k)}{\mathsf{X}(k)}\mathsf{W}(k).$$

where $\mathsf{W}(k)$ is a window function determined in advance from knowledge of X. Figure 6 shows a measured indoor propagated signal, with the impulse response envelope computed from it in the above manner. For this example, the window function $\mathsf{W}(k)$ was not optimized. A rectangular window, including roughly the spectral content of the original pulse waveform, was used. As a result, I did not add noise or interference for this illustrative example.

IV. Conclusion

The feasibility of maximally sensitive low-data rate UWB systems is primarily limited by signal processing complexity. Two key functions of these systems, interference rejection and correlative signal acquisition, are made computationally feasible by implementing the necessary long convolutions and correlation searches in the frequency domain, where the $Nlog(N)$ complexity of the FFT results in a large reduction in the number of operations. Similarly, the large singular value decompositions involved in channel estimates by least-squares deconvolution are greatly simplified by carrying out these estimations using the FFT with appropriately time-limited signals. The potential savings in the time, delay and energy costs of these computations suggest that it may be particularly advantageous to taylor low-data rate UWB signals and system designs for frequency domain implementation of initial signal acquisition and synchronization.

Acknowledgment

The author gratefully acknowledges the assistance of the Ultra Lab of the University of Southern California, for providing the test facilities and instrumentation that produced the experimental data shown in this paper.

References

[1] H. Meyr, M. Oerder, and A. Polydoros, “On sampling rate, analog prefiltering, and sufficient statistics for digital receivers,” *IEEE Trans. Commun.*, vol. 42, no. 12, pp. 3208–3214, Dec. 1994.

[2] H. D. Helms, “Fast Fourier transform method of computing difference equations and simulating filters,” *IEEE Trans. Audio Electroacoust.*, vol. AU-15, no. 2, pp. 85–90, 1967.

[3] R. M. Gray, *Toeplitz and Circulant Matrices: A review*, August 2002 ed. http://ee.stanford.edu/ gray/toeplitz.pdf.

[4] T. J. Cavicchi, *Digital signal processing*. New York: Wiley, 2000.

[5] E. H. Wold and A. M. Despain, “Pipeline and parallel-pipeline FFT processors for VLSI implementations,” *IEEE Trans. Comput.*, vol. C-33, no. 5, pp. 414–425, May 1984.

[6] M. Frigo and S. G. Johnson, “FFTW: an adaptive software architecture for the FFT,” in *IEEE Int. Conf. Acoust., Speech, and Sig. Proc.*, May 1998, pp. 1381 –1384.

[7] Federal Communications Commission, “Revision of part 15 of the commission’s rules regarding ultra-wideband transmission systems,” FCC, ET Docket 98-153, Apr. 2002.

[8] L. B. Milstein and P. K. Das, “An analysis of a real-time transform domain filtering digital communication system–part i: Narrow-band interference rejection,” *IEEE Trans. Commun.*, vol. COM-28, no. 6, pp. 816–824, June 1980.

[9] A. Thorpe and L. Scharf, “Reduced-rank methods for solving least squares problems,” in *Proc. 23rd Asilomar Conf. Signals, Syst. and Computers*, Nov. 1989, pp. 609–613.

[10] M. Barton and D. W. Tufts, “Reduced-rank least squares channel estimation,” *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 38, no. 8, pp. 1403–1410, Aug. 1990.