processing, but one with a concise treatment is the text by Levanon [2]. Specific topics related to digital implementations can be found in Chapter 5 of [1] and Chapter 13 in [8].

PROJECT 1: PROPERTIES OF THE LFM CHIRP SIGNAL

In this project, the characteristics of the LFM chirp signal are investigated: its time-domain appearance and its Fourier transform. In Project 2 we investigate its autocorrelation function, which is the output of a "pulse compression" matched filter.

The *chirp* radar signal is defined by the formula

$$s(t) = e^{j\pi W t^2/T} \qquad -\tfrac{1}{2}T \le t \le \tfrac{1}{2}T \tag{1-1}$$

Since the phase of $s(t)$ varies quadratically versus t, and the derivative of phase determines the instantaneous frequency of the signal, the frequency changes linearly versus t. The signal is complex valued because it is the baseband form of the linear frequency modulation. The LFM signal is a pulse whose time duration equals T seconds. Over the life of the pulse, the changing frequency sweeps from $-\frac{1}{2}W$ to $+\frac{1}{2}W$ hertz.

An intuitive guess about the frequency spectrum $S(f)$ leads one to suspect that most of the energy in the frequency domain will be concentrated in the range $|f| < \frac{1}{2}W$. This is, in fact, true if the frequency sweeps slowly enough; equivalently, if T is large enough. In the examples that follow, the dependence of the Fourier transform on the time–bandwidth product (TW) is studied.

Hints

Since the chirp signal is complex-valued and is processed by a complex-valued matched filter, all plots must be made of either the real part of (1-1) or the magnitude of the Fourier transform to show the correct behavior.

Two types of chirps are considered in the exercises that follow: a discrete-time chirp and a continuous-time chirp. MATLAB can deal only with a sampled version of the LFM signal, so the analog chirp is simulated by oversampling $s(t)$. For the discrete-time case, the sampling frequency is taken nearly equal to the swept bandwidth W, but for the continuous-time case, oversampling by a factor of 5 or more is needed for an accurate simulation.

EXERCISE 1.1

Sampled Chirp Signal

In MATLAB the chirp signal must be represented as a discrete-time signal. Therefore, the formula for $s(t)$ in (1-1) must be sampled at the rate $f_s = 1/T_s$. The sampling rate can be tied to W, the swept bandwidth of the chirp. In many cases the chirp is more or less bandlimited to a frequency extent of W. Therefore, it is convenient to let $f_s = pW$, where $p \ge 1$ represents the oversampling factor.

a. Convert the continuous-time chirp formula (1-1) into a discrete-time signal by sampling at a rate $f_s = pW$. Give the equation for the discrete-time signal in the form

$$s[n] = \exp\left[j2\pi\alpha\left(n - \tfrac{1}{2}N\right)^2\right] \qquad 0 \le n < N \tag{1-2}$$

Determine the correct formulas for α and N, and show that these parameters depend only on p and TW, the time–bandwidth product. (*Note*: It may not be possible to make the discrete-time chirp symmetric, depending on how the sampling times are defined. Starting at $t_n = -\frac{1}{2}T$ may not be the best strategy if $t = 0$ is not included in the sampling grid.)

b. Write an M-file to synthesize a discrete-time chirp. The function should have only two inputs, p and TW, and should return the complex-valued signal $s[n]$, as specified by the following comments.

```
function    s = dchirp( TW, p )
%DCHIRP       generate a sampled chirp signal
%  usage   s = dchirp( TW, p )
%      s :  samples of a digital "chirp" signal
%             exp(j(W/T)pi*t^2)   -T/2 <= t < +T/2
%     TW :  time-bandwidth product
%      p :  sample at p times the Nyquist rate (W)
```

c. Generate a sampled chirp signal whose time–bandwidth product is $TW = 50$. Plot the real and imaginary parts of the chirp signal, and observe how the apparent frequency changes versus time. The "chirped" pulse should have the characteristics given in Table 10.1.

TABLE 10.1
Desired Parameters for a Chirp Signal

Parameter	*Value*	*Units*
Pulse length	25	μs
Swept bandwidth	2	MHz
Sampling frequency	20	MHz
TW product	50	dimensionless

d. Redo the previous plot, but perform the sampling at $1.2W$. In this case the behavior of the chirp near its ends should exhibit the characteristics of a sampled sinusoid.

e. To show that the assumption of a "large" TW product is necessary, generate a chirp with $TW = 9$ and plot its Fourier transform. Use significant oversampling to simulate an analog chirp. Measure the fraction of the energy that lies outside the region $|f| < \frac{1}{2}W$.

f. *Optional*: Construct a plot of out-of-band energy versus TW over the range from 3 to 90. Although the out-of-band energy is quite small, it is possible to identify the value of TW where there is a "knee" in the curve.

EXERCISE 1.2

Fourier Transform of a Chirp

The Fourier transform magnitude of a chirp is *approximately* a rectangle in frequency *if the time–bandwidth product is large* (i.e., $TW > 50$). If we assume that $s(t)$ is the chirp and $S(f)$ is its Fourier transform, we can approximate $|S(f)|$ with a rectangle that extends from $f = -\frac{1}{2}W$ to $f = +\frac{1}{2}W$. Figure 10.1 shows that this rectangular approximation is not bad.

a. Find a formula for the height of the approximate spectrum $\tilde{S}(f)$ by using Parseval's theorem to equate the energies in the time and frequency domains.

b. Compute and plot the Fourier transform of the chirp for an oversampled case. Use the parameters in Table 10.1 from Exercise 1.1. This case should approximate the behavior of a continuous-time chirp, so use a long FFT (with zero padding) to get a smooth plot in the frequency domain. Scale the spectrum so that its magnitude near dc is correct.

c. Now consider the discrete-time case where the sampled chirp $s[n]$ is obtained by sampling just above the minimum rate—use $f_s = 1.2W$ and then $f_s = 2W$. Compute the DTFT of $s[n]$ and plot versus ω. Determine the appropriate cutoff frequency relative to $\omega = \pi$ in both cases. Make sure to use sufficient zero padding with the FFT to get a smooth plot of the DTFT.

d. Give a general formula for the height of the DTFT magnitude in terms of TW, N, and p. Again, Parseval's theorem should provide the link between the time and frequency domains.

e. The DFT of a chirp sometimes has remarkable properties, when $p = 1$. Since the number of nonzero samples in the discrete-time chirp signal is N, we can compute its N-point DFT. Show by example that whenever N is a multiple of 4 and $p = f_s/W = 1$, the DFT $S[k]$ is also an exact chirp. Determine a formula for the quadratic phase of $S[k]$, and verify that the magnitude of $S[k]$ is constant for all k.

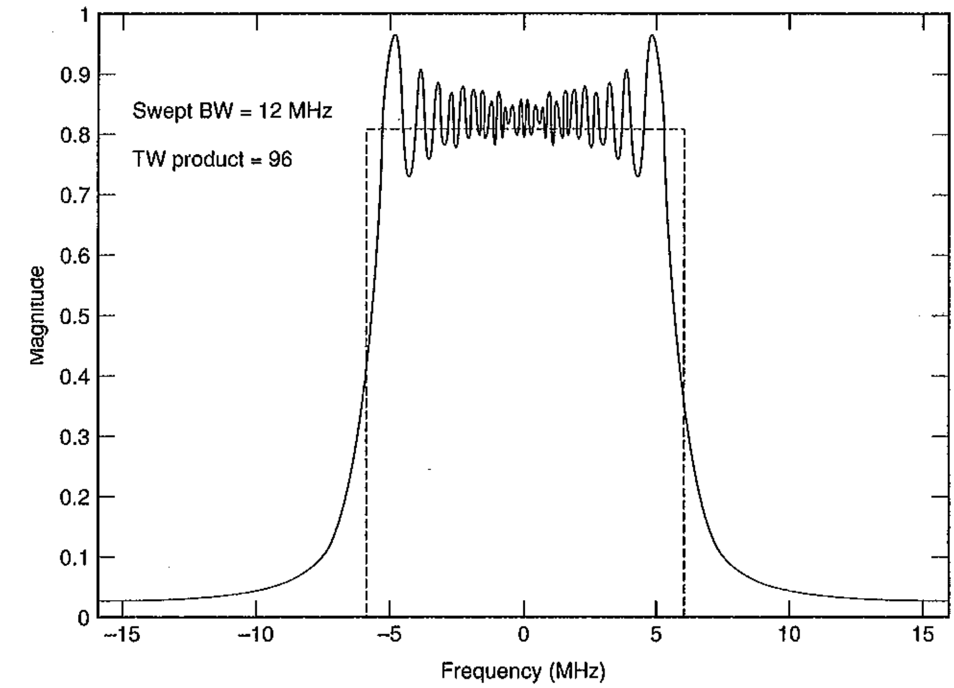

Figure 10.1
Fourier transform of a continuous-time chirp. Notice that most of the energy is concentrated in the frequency region between $-\frac{1}{2}W$ and $+\frac{1}{2}W$. The dashed line is the magnitude of an approximate transform $|\tilde{S}(f)|$, which is perfectly bandlimited.

PROJECT 2: RANGE PROCESSING

The transmitted signal in a radar must be designed differently for range estimation and for velocity estimation. In the case of range processing, the primary concern is to maximize the output SNR and range resolution from a matched filter, so large TW-product chirps are used. In this project, the LFM pulse and its matched filter are examined. The resulting output is a very narrow "compressed" peak that has large amplitude, which makes the detection of echoes easier. For a theoretical development of some equations describing the received radar signal, see the section *Theory of Radar Returns* in Project 4.

Hints

The matched filter involves a convolution that can be computed in either the time domain, via direct convolution (see `conv`), or in the frequency domain, via fast convolution with FFTs (see `fftfilt`).

EXERCISE 2.1

Pulse-Compression Matched Filtering

The matched filter is defined by either its frequency response, $H(f) = S^*(f)$, or by its impulse response, which, in turn, is determined by the transmitted waveform:

$$h(t) = s^*(-t) = e^{-j\pi W t^2/T} \qquad -\tfrac{1}{2}T \le t \le +\tfrac{1}{2}T \tag{2-1}$$

The output of the matched filter is

$$y(t) = h(t) * Gs(t - T_d) \tag{2-2}$$

where T_d is a time delay due to a target at range $R = \frac{1}{2}cT_d$. Since the matched filter is time-invariant, it is sufficient to make plots for the case where $T_d = 0$.

A discrete-time matched filter would involve a sampled version of $h(t)$. Since the discrete-time matched filter is FIR, its output can be computed via direct convolution. The purpose of this exercise is to study the form for the output of the matched filter.

a. In the continuous-time case, derive an expression for the matched filter output (2-2) by plugging into the convolution integral. Show that $y(t)$ can be written with an envelope function that is a slightly modified "sinc" function. This result has to be true because if the Fourier transform were approximated with a rectangle $\tilde{S}(f)$ as in Fig. 10.1, the output of a pulse-compression matched filter would be a "sinc" function whose width would be inversely proportional to W.

b. To verify the form for the matched filter's output, generate an oversampled chirp signal with $TW = 50$ and $p = 8$ or 10. Use this signal in a matched filter, creating the output by direct convolution (e.g., via `conv`). Plot the output, especially near the peak, and verify that it has the correct mainlobe width (distance between first zero crossings). Make sure to label the time axis in correct units.

c. The entire output will be created by `conv` and will require a significant amount of computation, but only the region near the peak is of interest. Prove that the matched filter output (2-2) can also be written as the autocorrelation of $s(t)$. Then use the M-file `acf` from Appendix A to generate only the region near the mainlobe.

d. Generate the matched filter output for the $f_s = 1.2W$ case. Verify that the mainlobe width is correct, but observe that there are just a few samples on the mainlobe. For the general case ($f_s = pW$), determine how many samples will be on the mainlobe.

EXERCISE 2.2

Processing Gain

The matched filter enhances SNR and, as a result, it is able to detect chirp signals even when they are "buried" beneath the noise level.

a. Generate a signal vector that is 700 points long but which contains in the middle a chirp with a TW product of 64 and $f_s = 3W$. Add (complex) white Gaussian noise to the signal so that the SNR is -10 dB (i.e., standard deviation is $\sqrt{10}$ times the signal amplitude). Plot the real part of the raw signal.

b. Process the signal through a digital matched filter and plot the output versus n. Explain how the location of the peak in n is related to the beginning of the echo and the starting index of the digital matched filter. In other words, how would the peak location index be converted to a range measurement?

c. Calculate the processing gain (in dB) by subtracting the input SNR from the output SNR. Measure the peak output versus the noise floor to compute an output SNR. Relate this output signal-to-noise ratio to the TW product and the oversampling factor p. It might be necessary to experiment with different values of TW and p to uncover a simple formula for processing gain.

EXERCISE 2.3

Range Estimation

The compressed-pulse output of the matched filter is significant because its increased height makes it easier for the radar processor to detect the echo and estimate its location when calculating range, according to $R = \frac{1}{2}cT_d$. There is, however, some uncertainty in a range estimate due to the inherent uncertainty in locating the peak of the echo. In the discrete-time case, this uncertainty could be fairly large.