CHAPTER 17

Doppler Processing

Mark A. Richards

Chapter Outline

17.1 INTRODUCTION

Doppler processing refers to the use of Doppler shift information to achieve one or both of two goals. The first is to enable detection of targets in environments where clutter is the dominant interference. The second is to measure Doppler shift, and thus radial velocity, of targets. In this chapter, two general classes of Doppler processing will be discussed: *moving target indication* (MTI) and *pulse-Doppler processing*. MTI processing addresses the first goal; pulse-Doppler processing addresses both.

Like target returns, clutter signals are echoes of external objects, so increasing radar power does not improve the signal-to-clutter ratio (SCR), as is seen in the signal-to-clutter form of the radar range equation in Chapter 2, equation (2.31). The SCR can be improved by decreasing the antenna azimuth beamwidth and, if the range cell extent is beam-limited, decreasing the elevation beamwidth. In a scenario where the ground clutter cell extent is pulse-limited or where volume clutter (weather, chaff) is a limiting factor, improving the range resolution with a shorter pulse or a higher time-bandwidth pulse will also help. Once the antenna and waveform parameters are set, however, Doppler processing is the principal means of improving SCR.

Chapter 8 described the Doppler effect in radar systems and typical components of the Doppler spectrum observed by a pulsed radar. Also discussed were the way a

pulsed radar collects the coherent range/pulse number (or fast-time/slow-time) data matrix and measures the Doppler spectrum via the "spatial Doppler" pulse-to-pulse phase history.

This chapter focuses on processing the slow-time data in a given range bin to analyze the Doppler content of this signal, to reduce interference from clutter, and to enable detection and parameter estimation of moving targets. The emphasis is on basic concepts for detecting moving targets in clutter and on the generic slow-time filtering and spectral analysis techniques and measurements used to make this possible. Introductions to some more specialized considerations such as ambiguity resolution and blind zones are included.

17.2 | REVIEW OF DOPPLER SHIFT AND PULSED RADAR DATA

17.2.1 Doppler Shift

As discussed in Chapter 8, the Doppler shift (*change* in radar frequency) observed by a monochromatic or narrowband radar illuminating a target with a radial velocity component of v m/s is, to a very good approximation,

$$f_d = \frac{2v}{c} f = \frac{2v}{\lambda} \tag{17.1}$$

where f and λ are the transmitted frequency and wavelength, respectively, and c is the speed of light. A positive value of v indicates an approaching target and thus a positive Doppler shift.

The numerical values of Doppler shift are small compared with the radar frequencies because typical target velocities are small compared with the speed of electromagnetic wave propagation. Table 17-1 (a reprint of Table 8-1) gives the magnitude of the Doppler shift corresponding to a radial velocity of 1 meter per second, knot, or mile per hour at representative radar frequencies. A Mach 2 aircraft observed with an L-band radar would cause a Doppler shift of only 4.4 kHz in a 1 GHz carrier frequency.

TABLE 17-1 ■ Doppler Shift as a Function of Velocity and Frequency

Radar Frequency f		Doppler Shift f_d (Hz)		
Band	Frequency (GHz)	1 m/s	1 knot	1 mph
L	1	6.67	3.43	2.98
S	3	20.0	10.3	8.94
C	5	33.3	17.1	14.9
X	10	66.7	34.3	29.8
K_u	16	107	54.9	47.7
K_a	35	233	120	104
W	95	633	326	283

17.3 | PULSED RADAR DOPPLER DATA ACQUISITION AND CHARACTERISTICS

17.3.1 Review of Pulsed Radar Data Matrix and Doppler Signal Model

Measurement and processing of Doppler data in a pulsed radar begins with the fast time/slow time (range/pulse number) matrix $y[l, m]$ of coherent, complex baseband data shown in Figure 17-1. The sampling rate in the fast time or range dimension (vertical in this figure) is at least equal to the fast time signal bandwidth, which in most cases equals the transmitted pulse bandwidth, $f_s \geq B$. For a radar using a simple pulse of length τ, $B \approx 1/\tau$. If pulse compression waveforms such as linear frequency modulation (LFM) chirps are used, B is the bandwidth of the modulated pulse.

The slow-time or pulse number dimension (horizontal in the figure) is sampled at the pulse repetition interval (PRI) T of the radar. Thus, the sampling rate in this dimension is the pulse repetition frequency (PRF). Each row of the matrix represents a series of measurements from the same range bin over M successive pulses. The total amount of time MT represented by the data matrix is called the *coherent processing interval* (CPI). In the absence of windowing, the Rayleigh (peak-to-null) Doppler resolution is $1/MT$ Hz. There may be several CPIs in a *dwell*, which is the amount of time a given target is within the antenna mainbeam on a single scan.

Suppose a target is present in range bin l_0, approaching the radar with a radial velocity v. It was shown in Chapter 8 that the model of the baseband slow-time target signal (not including noise or clutter) is the *spatial Doppler* signal

$$y[l_0, m] = A \exp\left(-j\frac{4\pi}{\lambda}R_0\right) \exp\left[+j2\pi\left(\frac{2v}{\lambda}\right)mT\right], \quad m = 0, \ldots, M-1 \quad (17.2)$$

That is, the slow-time data sequence forms a complex sinusoid at the Doppler frequency $2v/\lambda$ Hz. Equation (17.2) is valid under the "stop-and-hop" assumption discussed in

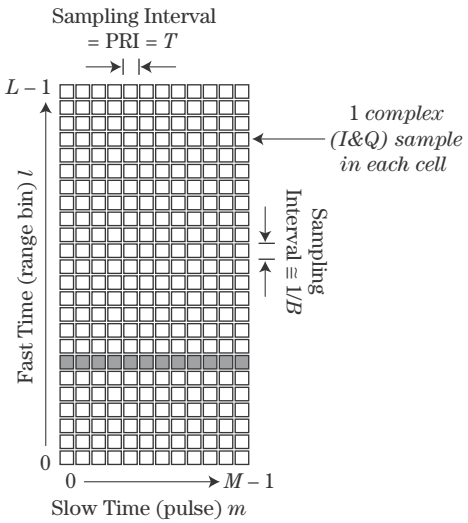

FIGURE 17-1 ■ Notional two-dimensional pulse-Doppler data matrix. The shaded samples are the slow-time signal for the seventh range bin, $y[6,m]$.

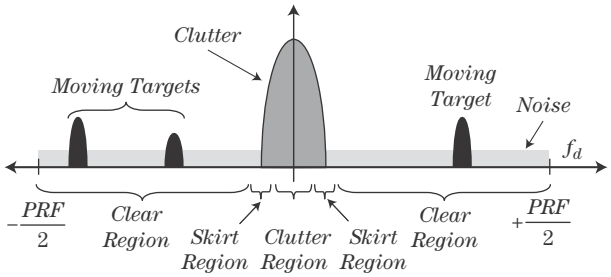

FIGURE 17-2 ■ The principal period of a notional generic Doppler spectrum containing noise, clutter, and target components. A stationary radar has been assumed.

Chapter 8 and [1]. For the velocities and CPI durations seen in conventional Doppler processing, it is usually safe to assume that range migration over the CPI is less than a range bin and can be ignored.

17.3.2 Generic Doppler Spectrum for a Single Range Bin

In general, the spectrum of the slow-time signal from a single range bin consists of noise, clutter, and one or more target signals. The distribution of these signals in range and velocity (or Doppler) was discussed in some detail in Chapter 8. Figure 17-2, which repeats Figure 8-16, shows a notional generic Doppler spectrum as observed from a stationary radar for a single range bin containing clutter, noise, and three moving targets.

In many situations, the relative amplitudes of the clutter, target, and noise signals are generally as shown: the target returns are above the noise floor (*signal-to-noise ratio* [SNR] > 1) but below the clutter (SCR < 1). In this case, targets cannot be detected reliably based on amplitude in the slow-time domain alone because the presence or absence of the target makes little difference to the overall power of the received signal, which is dominated by the clutter. Doppler processing, however, can separate moving target signals from the clutter signals in the frequency domain. The clutter can be explicitly filtered out, leaving the target returns as the strongest signal present, or the spectrum can be computed explicitly so that targets outside of the clutter region can be located by finding frequency components that significantly exceed the noise floor.

As shown in Chapter 8, the simple spectrum of Figure 17-2 is complicated by its distribution over range; moving platforms, which introduce sidelobe clutter and altitude lines; and ambiguities, which result in range and velocity foldovers. Nonetheless, it has all the features needed to introduce the basic concepts and algorithms of Doppler processing.

17.3.3 Review of Range and Velocity Aliasing and Coverage

What happens if the radar views a target having a Doppler shift magnitude greater than $PRF/2$? As was discussed in Chapter 14, the Doppler spectrum obtained by computing the discrete-time Fourier transform (DTFT), discrete Fourier transform (DFT), or fast Fourier transform (FFT) of a slow-time data sequence is periodic in frequency, with the principal period ranging from $-PRF/2$ to $+PRF/2$ Hz, corresponding to a velocity range of $\pm\lambda PRF/4$. Consequently, all frequency components are aliased, or replicated, every PRF Hz. Figure 17-3 illustrates this periodicity for a spectrum similar to Figure 17-2. Targets at a Doppler shift f_d outside of the $\pm PRF/2$ range will appear in the principal period at an apparent Doppler frequency

$$f_{d_a} = f_d - k_a PRF \tag{17.3}$$

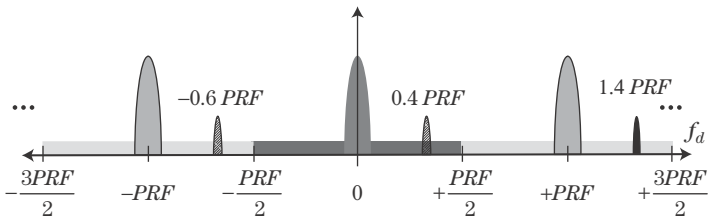

FIGURE 17-3 ■ Replication of Doppler spectrum of sampled slow-time data, showing aliasing of a high-speed target.

where k_a is an integer chosen such that f_{d_a} is between $-PRF/2$ and $+PRF/2$. For example, the target having a Doppler shift of $1.4PRF$ will have the aliases shown at $0.4PRF$ and $-0.6PRF$ as well as others at all frequencies $1.4PRF + k_aPRF$ for any integer k_a. The $k_a = -1$ alias at $0.4PRF$ will be the apparent Doppler shift for this target because it falls in the principal period of the spectrum.

The *unambiguous Doppler coverage* of the radar is the width of the range of Doppler shifts that can be measured without aliasing and is simply

$$f_{d_{ua}} = PRF \tag{17.4}$$

Caution must be used to avoid confusing the unambiguous Doppler coverage $f_{d_{ua}}$ with the limits $\pm f_{d_{ua}}/2$ of the principal period of the velocity spectrum. The range extent that can be represented unambiguously given T or PRF is

$$R_{ua} = \frac{c\,T}{2} = \frac{c}{2PRF} \tag{17.5}$$

This quantity is called the *unambiguous range* or the *range coverage*. Targets at range $R > R_{ua}$ will alias to the apparent range

$$R_a = R - k_a R_{ua} \tag{17.6}$$

where k_a is an integer chosen such that R_a is between zero and R_{ua}.

Combining equations (17.4) and (17.5) gives the combined range-Doppler coverage as

$$R_{ua} f_{d_{ua}} = \frac{c}{2}. \tag{17.7}$$

The equivalent range-velocity coverage is

$$R_{ua} v_{ua} = \frac{\lambda c}{4} \tag{17.8}$$

Equation (17.7) makes it clear that the total unambiguous range-Doppler coverage is independent of the PRF. Increasing the unambiguous range coverage by increasing the PRI reduces unambiguous Doppler coverage and vice versa. For a constant PRI system, the total unambiguous range-Doppler coverage can only be increased by increasing the wavelength—that is, using a lower radar frequency. Section 17.4.2 will show that the use of variable PRIs offers another way to improve total range-Doppler coverage.

17.4 MOVING TARGET INDICATION

Doppler processing is the term applied to filtering or spectral analysis of the slow-time signal $y[m]$ received from a fixed range bin over a period of time corresponding to several pulses. Doppler processing is applied independently to each range bin of interest. There are two major classes of Doppler processing: moving target indication and pulse-Doppler processing. In this chapter, MTI refers to the case where the slow-time signal is processed

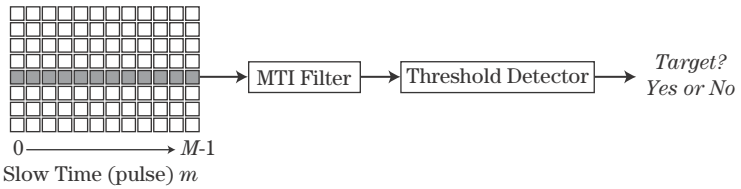

FIGURE 17-4 ■ MTI filtering and detection process.

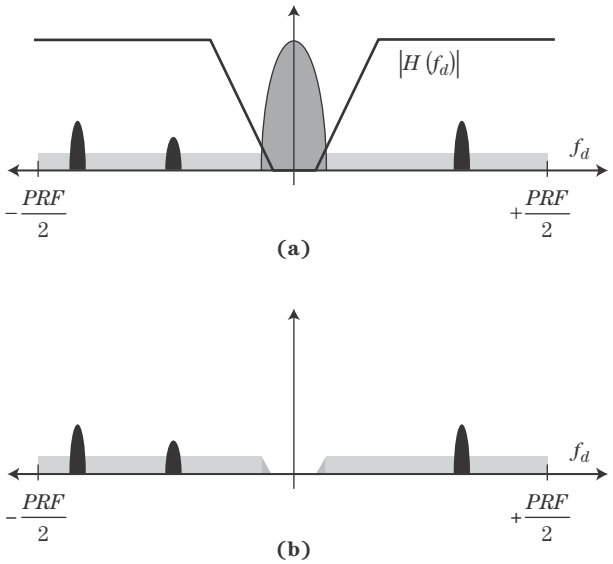

FIGURE 17-5 ■ Desired effect of the MTI filter. (a) Slow-time spectrum before MTI filtering and notional MTI filter frequency response $|H(f_d)|$. (b) After MTI filtering.

entirely in the time domain, usually using a single high-pass filter. Pulse-Doppler processing refers to the case where the signal is processed in the frequency domain, usually using an FFT.[1] As will be seen, MTI processing produces limited information at very low computational cost; pulse-Doppler processing requires more computation but produces more information and greater *signal-to-interference ratio* (SIR) improvement. Only coherent Doppler processing using digital implementations is considered, since this is the approach taken in most modern radars. Alternative systems using noncoherent Doppler processing and implementations based on analog technologies are described in [2–6], among many others.

MTI processing applies a linear filter to the slow-time data sequence to suppress the clutter component. Figure 17-4 illustrates the process. The type of filtering needed can be understood by considering the notional spectrum of the slow-time data shown in Figure 17-5. In this figure, it is assumed that knowledge of the platform motion and scenario geometry has been used to center the clutter spectrum at zero Doppler frequency. Clearly, some form of high-pass filter is needed to attenuate the clutter without filtering out moving targets in the clear portions of the Doppler spectrum.

The output of the high-pass MTI filter will be a modified slow-time signal containing components due to noise and, possibly, one or more targets. This signal is passed to a detector, typically based on the amplitude or squared amplitude of the data and possibly involving noncoherent detection as well (see Chapters 3 and 15). If the amplitude of the filtered signal exceeds the detector threshold (i.e., its energy is too great to likely be the

[1] Skolnik [2] distinguishes MTI and pulse-Doppler by defining pulse-Doppler as a system that uses a PRF high enough to avoid blind speeds. Here, the two cases are distinguished based on the type of processing operations used and the information obtained.

result of noise alone), a target will be declared; otherwise, the data are declared to represent interference only.

Note that in MTI processing, the presence or absence of a moving target is the only information obtained. The filtering process of Figures 17-4 and 17-5 does not provide any estimate of the Doppler frequency at which the target energy causing the detection occurred; thus, it "indicates" the presence of a moving target but does not determine whether the target is approaching or receding or at what radial velocity. Furthermore, it provides no indication of the number of moving targets present. If multiple moving targets are present in the slow-time signal from a particular range bin, the result will still be only a "target present" decision from the detector. On the other hand, MTI processing is very simple and computationally undemanding.

17.4.1 Pulse Cancellers

The major MTI design decision is the choice of the particular MTI filter to be used. MTI filters are typically low-order, simple *finite impulse response* (FIR; also called tapped delay line or nonrecursive) designs [7]. Indeed, some of the most common MTI filters are based on very simple heuristic design approaches. For example, suppose a stationary radar illuminates a stationary clutter scatterer. After demodulation, the measured sample of the received signal in the appropriate range bin will be of the form $Ae^{j\phi}$ for some amplitude A and phase ϕ. If the measurement is repeated, the same value will be measured again (ignoring noise). Subtracting the echoes from successive pairs of pulses would cancel the clutter return completely.

Now consider the same scenario, but with a moving target. While the amplitude of the successive echoes may be identical, the range to the target will change between pulses by an amount $\delta R = vT$ m, where v is the radial velocity of the target and T is the pulse repetition interval. Consequently, the phase of the echo will change by $(4\pi/\lambda)\delta R$ radians as can be seen in equation (17.2). Subtracting these two measurements will not result in a zero signal due to the different phases.

This reasoning motivates the *two-pulse MTI canceller*, also referred to as the *single canceller* or *first-order canceller*. Figure 17-6a illustrates the flowgraph of a two-pulse canceller, which is an especially simple FIR digital filter. The input data are a sequence of baseband complex (in-phase and quadrature, or I and Q) data samples from the same range bin over successive pulses, forming a discrete-time sequence $y[m]$ with a sampling

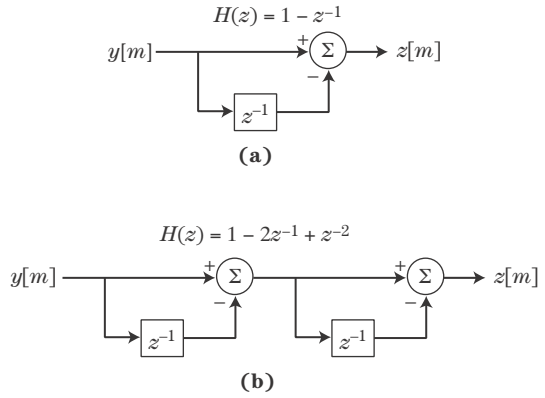

FIGURE 17-6 ■ Flowgraphs and transfer functions of basic MTI cancellers. (a) Two-pulse canceller. (b) Three-pulse canceller.

interval T equal to the pulse repetition interval. The discrete-time transfer function of this filter is simply $H(z) = 1 - z^{-1}$. The frequency response as a function of analog Doppler frequency f_d in hertz is obtained by setting $z = e^{j2\pi f_d T}$:

$$\begin{aligned} H(f_d) &= (1 - z^{-1})\big|_{z=e^{j2\pi f_D T}} = 1 - e^{-j2\pi f_d T} \\ &= e^{-j\pi f_d T}(e^{+j\pi f_d T} - e^{-j\pi f_d T}) \\ &= 2je^{-j\pi f_d T}\sin(\pi f_d T) \end{aligned} \tag{17.9}$$

The frequency response may also be expressed in terms of normalized frequency $\hat{f} = f_d T$ cycles/sample or the radian equivalent, $\hat{\omega} = \omega T = 2\pi f_d T$ radians/sample. For example, in terms of normalized radian frequency, the frequency response of the two-pulse canceller is

$$H(\hat{\omega}) = 2je^{-j\hat{\omega}/2}\sin(\hat{\omega}/2) \tag{17.10}$$

Note that as f_d ranges from $-PRF/2$ to $+PRF/2$ ($-1/2T$ to $+1/2T$), $\hat{f}$ ranges from -0.5 to $+0.5$ and $\hat{\omega}$ from $-\pi$ to $+\pi$.

Figure 17-7a plots the magnitude of this frequency response. Note that the filter is high-pass in nature, with a null at zero frequency to suppress the clutter energy. Spectral components representing moving targets may either be partially attenuated or amplified,

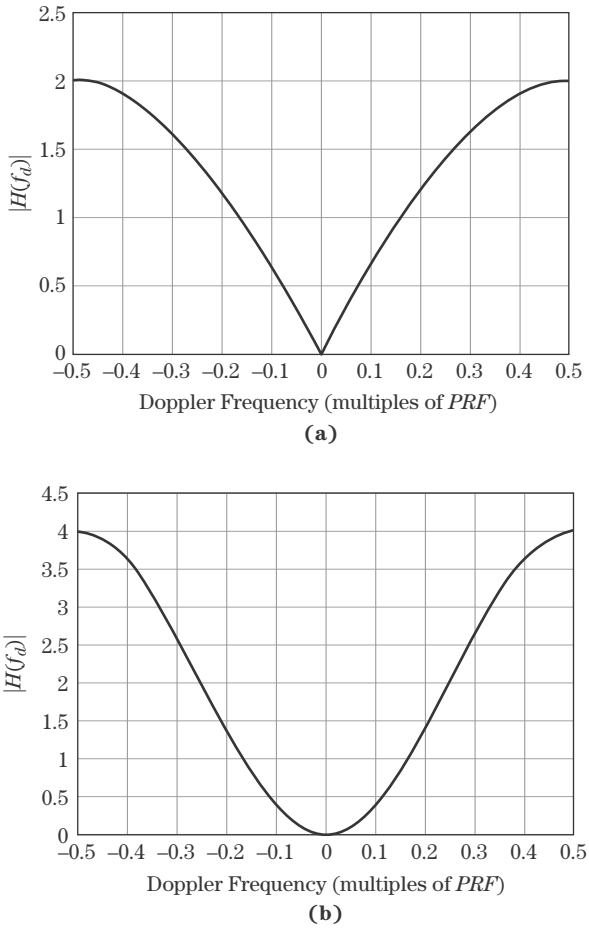

FIGURE 17-7 ■ Frequency response of basic MTI cancellers. (a) Two-pulse canceller. (b) Three-pulse canceller.

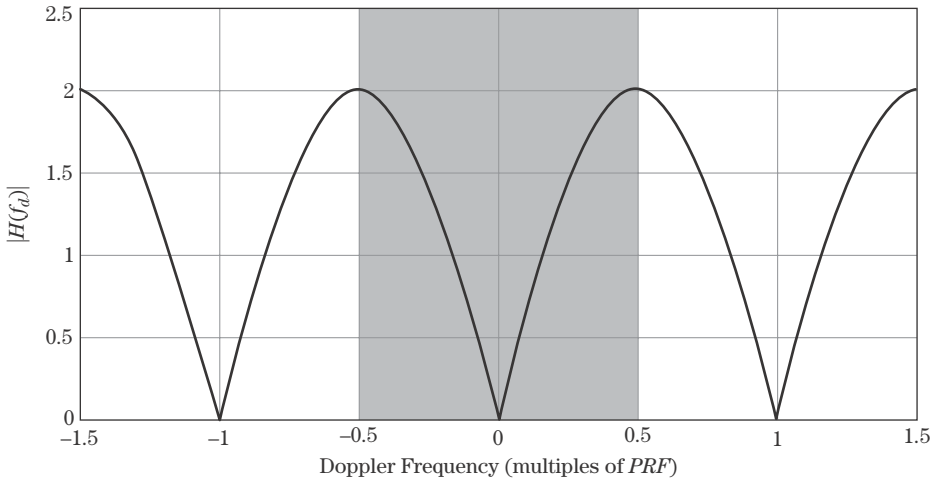

FIGURE 17-8 ■ "Blind" Doppler frequencies.

depending on their precise location on the Doppler frequency axis. Also recall that, like all discrete-time filters, the frequency response is periodic with a period of 1 cycle/sample, corresponding to a period of 2π in normalized radian frequency $\hat{\omega}$ or $1/T = PRF$ in analog frequency in hertz. This is illustrated in Figure 17-8, which extends Figure 17-7a to three periods. The shaded area highlights the principal period from $-PRF/2$ to $+PRF/2$; this is all that is normally plotted. The implications of the periodicity will be considered in Section 17.4.2.

The two-pulse canceller imposes a very low computational load; Figure 17-6a shows that its implementation requires no multiplications and only one subtraction per output sample. As Figure 17-7a shows, however, it is a poor approximation to an ideal high-pass filter for clutter suppression. The next traditional step up in MTI filtering is the three-pulse (second-order or double) canceller, obtained by cascading two two-pulse cancellers. The flowgraph and frequency response are shown in Figures 17-6b and 17-7b. The three-pulse canceller improves the null depth and width in the vicinity of zero Doppler and requires only two subtractions per output sample, but there is still a large variation in the filter gain or attenuation for moving targets at various Doppler shifts away from zero Doppler.

Despite their simplicity, the two-and three-pulse cancellers can be very effective against clutter with moderate to high pulse-to-pulse correlations. This is because highly correlated clutter corresponds to a narrow power spectrum, so that a high fraction of the clutter energy falls within the filter notch at zero Doppler shift. Section 17.4.3 addresses the clutter suppression obtainable with pulse cancellers.

The idea of cascading two-pulse canceller sections to obtain higher-order filters can be extended to the N-pulse canceller, obtained by cascading $N - 1$ two-pulse canceller sections. The transfer function of the N-pulse canceller is therefore

$$H_N(z) = (1 - z^{-1})^{N-1} \tag{17.11}$$

Other types of digital high-pass filters could also be designed for MTI filtering. For example, an FIR high-pass filter could be designed using standard digital filter design techniques such as the window method or the Parks-McClellan algorithm [7]. Alternatively, *infinite impulse response* (IIR) high-pass filters could be designed. However, most MTI filters are low order, and, with only a few filter coefficients to optimize, more elaborate MTI filter designs provide only modest performance improvements over pulse cancellers. Two- or

three-pulse cancellers are commonly used for initial MTI filtering due to their reasonable effectiveness and computational simplicity.

The N-pulse cancellers previously described are widely used. Nonetheless, they are motivated by heuristic ideas. Can a more effective pulse canceller be designed? Since the goal of MTI filtering is to maximize the signal-to-clutter ratio, it should be possible to apply the optimum filter concept, used previously in Chapter 14 to develop the matched filters, to this problem. The details are worked out in [1]. Assuming the clutter power is σ_c^2 and the noise power is σ_n^2, the resulting filter coefficients in the two-pulse case are

$$\boldsymbol{H} = \hat{k}\left[\sigma_c^2 + \sigma_n^2 - \rho^* \sigma_c^{2}\right]^T \tag{17.12}$$

where $\hat{k}$ absorbs all scale factors. In this equation, ρ is the first normalized autocorrelation lag of the clutter, and the superscript T denotes matrix transpose.

To interpret this result, consider the case where the clutter is the dominant interference and is highly correlated from one pulse to the next. Then σ_n^2 is negligible compared with σ_c^2, and ρ is close to 1. Absorbing σ_c^2 into $\hat{k}$, the matched filter coefficients are then 1 and approximately -1, that is, nearly the same as the two-pulse canceller. Despite its simplicity, the two-pulse canceller is therefore nearly a first-order matched filter for MTI processing when the clutter-to-noise ratio is high and the successive clutter pulses are highly correlated. In the limit of very high clutter-to-noise ratio and perfectly correlated clutter, the two-pulse canceller is exactly the first-order matched MTI filter.

17.4.2 Blind Speeds and Staggered PRFs

The frequency response of all discrete-time filters is periodic, repeating with a period of one in the normalized cyclical frequency, corresponding to a period of $PRF = 1/T$ Hz of Doppler shift. Since MTI filters are designed to have a null at zero frequency, they will also have nulls at Doppler frequencies that are multiples of the pulse repetition frequency. Consequently, a target moving with a radial velocity that results in a Doppler shift equal to a multiple of the PRF will be filtered out by the MTI filter. This is illustrated for a two-pulse canceller in Figure 17-8, which shows three periods of the frequency response. The first positive and negative blind Doppler frequencies are

$$f_b = \pm PRF \tag{17.13}$$

Velocities that result in these unfortunate Doppler shifts are called *blind speeds* because the target return will be suppressed by the MTI filter; the system is "blind" to such targets. From a digital signal processing point of view, blind speeds represent target velocities that will be aliased to zero frequency. The first blind speeds are obtained by scaling the blind Doppler frequencies:

$$v_{blind} = \frac{\lambda}{2} f_b = \pm \frac{\lambda}{2} PRF \tag{17.14}$$

As the PRF is increased for a given radio frequency (RF), the unambiguous range decreases, and the first blind Doppler frequency increases. Blind Doppler frequencies could be avoided by choosing the PRF high enough so that the first blind Doppler frequency exceeds any actual velocity likely to be observed for targets of interest. Unfortunately, sometimes no PRF will allow unambiguous coverage of both the range and Doppler intervals of interest. For example, suppose a designer requires an unambiguous range

of $R_{ua} = 100$ km and $v_{ua} = 100$ m/s (± 50 m/s) of unambiguous velocity coverage. The maximum RF at which this is possible is 2.25 GHz, as can be seen from equation (17.8). If the radar is required to be at X-band (10 GHz), then the combination of 100 km unambiguous range coverage and 100 m/s unambiguous velocity coverage is not obtainable with any PRF, and some ambiguity must be accepted in range, Doppler, or both.

The use of *staggered PRFs* is a data collection and processing technique that raises the first blind speed significantly without significantly degrading unambiguous range [5,8]. PRF staggering can be performed on either a pulse-to-pulse or CPI-to-CPI basis; the latter is also called a block-to-block or dwell-to-dwell basis. The CPI-to-CPI case is common in airborne pulse-Doppler radars and is discussed in Section 17.5.8.

Pulse-to-pulse stagger varies the pulse repetition interval, or equivalently the PRF, from one pulse to the next within a single CPI. The resulting slow-time data sequence is then passed through a conventional MTI filter such as a two- or three-pulse canceller. As will be seen, this increases the Doppler coverage within a single CPI. One disadvantage is that the slow-time data sequence in a given range bin is now nonuniformly sampled in slow time, making it more difficult to apply Doppler filtering to the data and greatly complicating analysis. Another is that range-ambiguous mainlobe clutter, if any, can cause large pulse-to-pulse amplitude changes as the PRF varies, since the range of the second-time-around clutter that folds into each range cell will change as the PRF changes. Consequently, pulse-to-pulse PRF stagger is generally used only in low PRF modes in which no range ambiguities are expected.

Consider a system using a set of P staggered PRFs $\{PRF_p\} = \{PRF_0, PRF_1, \ldots, PRF_{P-1}\}$. The corresponding set of pulse repetition intervals is $\{T_p\} = \{1/PRF_p\}$. Assume that each of the PRFs can be expressed as an integer multiple of the *greatest common divisor* (gcd) of the set, f_g:

$$\begin{aligned} PRF_p &= k_p \,\text{gcd}(PRF_0, \ldots, PRF_{P-1}) \\ &\equiv k_p f_g \end{aligned} \tag{17.15}$$

The set of integers $\{k_p\}$ is called the *staggers*,[2] and the ratio k_m:k_p of any two of them is called a *stagger ratio*. A CPI of data is collected by transmitting the first pulse (pulse number 0) and sampling the desired range bins. The transmitter then waits T_0 seconds and transmits pulse number 1 to get the second measurement in each range bin. The transmitter then waits T_1 seconds before transmitting pulse number 2 and so forth. If the CPI contains more than P pulses, then after pulse number $P - 1$ the transmitter waits T_{P-1} seconds, transmits the P-th pulse, and then cycles back to the first PRI, waiting T_0 seconds and transmitting pulse number $(P + 1)$ and so on until the full CPI of M pulses has been collected.

For a fixed PRF, any MTI filter will exhibit blind Doppler frequencies at all integer multiples of the PRF. Similarly, the first true blind Doppler frequency of a system using staggered PRFs will be the lowest frequency that is blind at all of the individual PRFs—that is, the *least common multiple* (lcm) of the set [1]:

$$\begin{aligned} f_b &= \text{lcm}\,(PRF_0, \ldots, PRF_{P-1}) \\ &= f_g \text{lcm}\,(k_0, \ldots, k_{P-1}) \end{aligned} \tag{17.16}$$

[2]Some authors work in terms of the PRIs instead of the PRFs and use the term *staggers* to refer to the ratio of the $\{PRI_p\}$.

A measure of effectiveness of this technique is how much the blind speed of the staggered system is increased relative to that of an unstaggered system with the same average PRI. It is straightforward to show that the ratio of the first blind Doppler frequency f_s of the staggered PRF system and the blind Doppler frequency of an unstaggered system with a PRF f_{us} corresponding to the average PRI is [1]

$$\frac{f_s}{f_{us}} = \frac{1}{P}\text{lcm}\,(k_0, \ldots, k_{P-1}) \left(\sum_{p=0}^{P-1} \frac{1}{k_p} \right) \tag{17.17}$$

For example, a two-PRF system with a stagger ratio of 3:4 would have a first blind Doppler frequency 3.5 times that of a system using a fixed PRI equal to the average of the two individual PRIs. If a third PRF is added to give the set of staggers $\{3, 4, 5\}$, the first blind Doppler frequency will be 15.67 times that of the comparable unstaggered system.

The unambiguous range R_{us} of the unstaggered system is the range corresponding to the unstaggered PRF, f_{us}, which is just $c/2f_{us}$. The unambiguous range R_{min} of the staggered PRF system is the shortest of the unambiguous ranges corresponding to the individual PRFs. It is easy to show that the ratio of these two ranges is

$$\frac{R_{min}}{R_{us}} = \frac{P}{\max\{k_p\} \left(\sum_{p=0}^{P-1} \frac{1}{k_p} \right)} \tag{17.18}$$

For the two-PRF system with staggers $\{3, 4\}$, this ratio is 6/7, representing a reduction of unambiguous range of 14% in exchange for the increase in velocity coverage by a factor of 3.5. For the $\{3, 4, 5\}$ case, the reduction in unambiguous range is a factor of 45/47, or 8%, while the Doppler coverage is increased by a factor of 15.67.

The slow-time data in a staggered PRF MTI system has a nonuniform sampling interval from one sample to the next. Unfortunately, the response of a digital filter to such an input, while still linear, is not time-invariant, so the frequency response of a pulse-to-pulse staggered system cannot be determined using conventional Fourier analysis techniques. Instead, an approach based on first principles can be used to explicitly compute the frequency response of a specified MTI filter with staggered PRF data by determining the average amplitude of the filter output when the input is a pure complex sinusoid of arbitrary frequency and random initial phase [5,8,9]. The result for the squared magnitude of the frequency response of the two-pulse canceller with staggered PRFs is [1]

$$\left|H_{2,P}(f)\right|^2 = \frac{4}{P}\sum_{p=0}^{P-1} \sin^2(\pi f T_p) = \frac{4}{P}\sum_{p=0}^{P-1} \sin^2(\pi f / PRF_p) \tag{17.19}$$

where the notation $H_{N,P}(f)$ indicates the frequency response of an N-pulse canceller using P staggers. The response of more general MTI filters can be obtained using a similar technique. The actual frequency in hertz rather than normalized frequency $\hat{\omega}$ or $\hat{f}$ must be used in equation (17.19) because the nonuniform sampling rate invalidates the usual definition of normalized frequency.

Figure 17-9 compares the frequency response of a two-pulse canceller using two ($P = 2$) PRFs versus conventional single-PRF operation. The staggered case uses PRFs of 750 and 1,000 pulses per second; thus, $f_g = 250$ Hz and the set of staggers k_p is $\{3, 4\}$. The first blind Doppler frequency occurs at the least common multiple of 750 and 1,000 Hz,

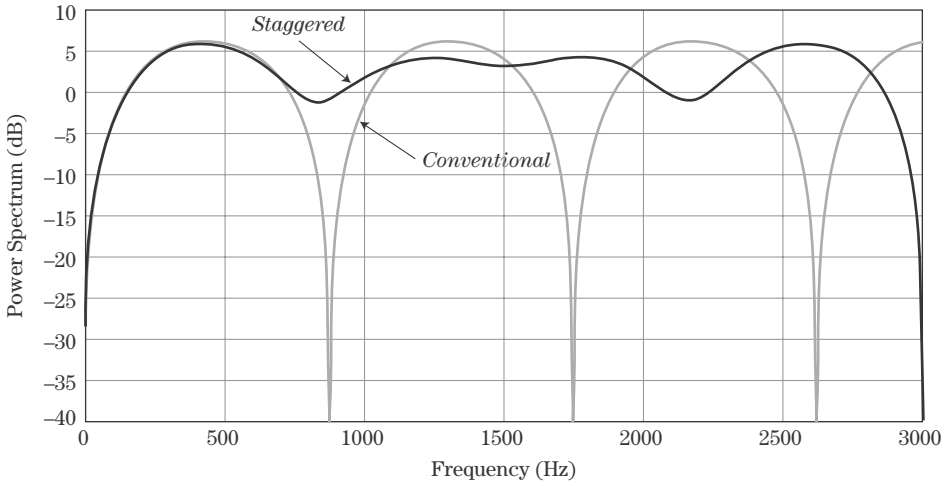

FIGURE 17-9 ■ Comparison of two-pulse canceller frequency response with conventional unstaggered waveform and 3:4 staggered waveform.

which is 3,000 Hz. f_{us}, which is the reciprocal of the average PRI $T_{avg} = 1.167$ ms, is 857.14 Hz. The conventional unstaggered response collected with the PRI $= T_{avg}$ shows blind Doppler shifts equal to integer multiples of 857.14 Hz. Thus, staggering the PRF has increased the blind Doppler frequency by a factor of 3.5 ($= 3{,}000/857.14$), consistent with equation (17.17). Note also that the unambiguous range corresponding to the highest PRF used (1,000 Hz) is $R_{min} = 150$ km, while the unambiguous range of the equivalent unstaggered system having the PRF 857.14 pulses/second is $R_{us} = 175$ km, a reduction by a factor of 6/7 (14%), as previously predicted.

The reader is cautioned that no effort has been made in this simple example to optimize the number or choice of PRFs. The "passband" region of the staggered response (about 300 to 2,700 Hz) shows frequency response variations of about 9 dB. Good staggered PRF systems spend considerable effort on PRF selection. Use of more than two PRIs in the stagger sequence and careful selection of their ratios can result in overall MTI frequency responses with much less variability than shown in this simple example [10,11]. Another design approach uses randomized PRIs rather than the fixed schedule described here to extend the blind speed [12].

17.4.3 MTI Figures of Merit

The goal of MTI filtering is to suppress clutter. In doing so, the MTI filter also attenuates or amplifies the target return, depending on the particular target Doppler shift. The change in signal and clutter power then affects the probabilities of detection and false alarm achievable in the system in a manner dependent on the particular design of the detection system.

There are three traditional MTI filtering figures of merit in wide use [13]. *Clutter attenuation* measures only the reduction in clutter power at the output of the MTI filter compared with the input but is simplest to compute. *Improvement factor* quantifies the increase in signal-to-clutter ratio due to MTI filtering; as such, it accounts for the effect of the filter on the target as well as on the clutter. *Subclutter visibility* is a more complex measure that also takes into account the detection and false alarm probabilities and the detector characteristic. Because of its complexity, it is less often used. In this chapter, attention is concentrated on clutter attenuation, *CA*, and improvement factor, *I*. Due to their common use in evaluating adaptive processing systems, two newer metrics,

minimum detectable velocity (MDV) and *usable Doppler space fraction* (UDSF), are introduced later in the chapter in the context of pulse-Doppler processing [14].

Calculation of clutter attenuation and improvement factor can be approached in several ways: frequency domain approaches using clutter power spectra and MTI filter transfer functions; an approach using the autocorrelation functions of the input and output of the MTI filter; and a method based on vector analysis. Only the frequency domain approach, which is perhaps the most intuitive, is described here. A description and examples of the other two methods are available in [1].

Clutter attenuation directly evaluates the MTI filter's effectiveness at its main function of suppressing the clutter energy. It is simply the ratio of the clutter power at the input of the MTI filter to the clutter power at the output, which can be calculated by integrating the clutter power density spectrum before and after the MTI filter is applied:

$$CA = \frac{\sigma_{ci}^2}{\sigma_{co}^2} = \frac{\int_{-PRF/2}^{PRF/2} S_{cc}(f_d)df_d}{\int_{-PRF/2}^{PRF/2} S_{cc}(f_d)|H(f_d)|^2 df_d} \tag{17.20}$$

Here σ_{ci}^2 and σ_{co}^2 are the clutter power at the filter input and output, respectively; $S_c(f_d)$ is the sampled clutter power spectrum; and $H(f_d)$ is the discrete-time MTI filter frequency response. In equation (17.20) each term is expressed in terms of analog frequency, f_d, but could also be expressed in terms of normalized frequencies $\hat{\omega}$ or $\hat{f}$. Since the MTI filter presumably reduces the clutter power, *CA* will be greater than 1. In fact, clutter attenuation can be 20 dB or more in favorable conditions. However, it also depends on the clutter itself through $S_c(f_d)$. A change in clutter power spectrum due to changing terrain or weather conditions will alter the achieved clutter cancellation. The shape of the clutter power spectrum and its spread in meters per second are determined by the physical phenomenology (type of clutter, clutter motion, weather conditions) and are not under the radar engineer's control. However, the percentage of the discrete-time spectrum width in hertz to which a given clutter power spectrum is mapped depends on the PRF and wavelength and therefore is determined by the system design. Thus, the width of the clutter spectrum relative to the PRF is a combination of factors, some influenced by the radar design and some not.

Improvement factor *I* is defined formally as the SCR at the filter output divided by the SCR at the filter input, averaged over all target radial velocities of interest [13]. Considering for the moment only a specific target Doppler shift, the improvement factor can be factored into the form [8]

$$I = \frac{(SCR)_{out}}{(SCR)_{in}} = \left(\frac{S_{out}}{S_{in}}\right)\left(\frac{C_{in}}{C_{out}}\right) = G \cdot CA \tag{17.21}$$

where *G* is the MTI filter *gain* at the Doppler shift of interest. Figure 17-7 makes clear that the effect of the MTI filter on the target signal is a strong function of the target Doppler shift. Thus, *G* is a function of target velocity, whereas clutter attenuation *CA* is not. The value of *CA* in equation (17.21) is obtained from equation (17.20). The value of target power gain, $G(f_{d_0})$, for a target at some specific Doppler frequency, f_{d_0}, is determined by

the frequency response of the MTI filter at that frequency:

$$G(f_{d_0}) = \left(\frac{S_{out}}{S_{in}}\right) = \left|H(f_{d_0})\right|^2 \tag{17.22}$$

With this definition of G, the improvement factor is a function of the target Doppler shift f_{d_0}.

To reduce I to a single number, the definition calls for averaging uniformly over all target Doppler shifts "of interest" [13]. It is more common to assume the target velocity is completely unknown a priori and use the average target power gain over a full period of the Doppler spectrum, which is just

$$G = \frac{1}{PRF}\int_{-PRF/2}^{PRF/2} |H(f_d)|^2 df_d \tag{17.23}$$

This equation gives an average value of G of 2 for a two-pulse canceller and 6 for a three-pulse canceller [4]. Combining (17.23) and (17.20) in (17.21) gives the improvement factor as

$$I = \frac{\left\{\int_{-PRF/2}^{PRF/2} |H(f_d)|^2 df_d\right\}\left\{\int_{-PRF/2}^{PRF/2} S_{cc}(f_d) df_d\right\}}{PRF\left\{\int_{-PRF/2}^{PRF/2} S_{cc}(f_d)|H(f_d)|^2 df_d\right\}} \tag{17.24}$$

Table 17-2 shows the improvement factor predicted by equation (17.24) for various clutter spectral widths, assuming the clutter spectrum is Gaussian in shape.[3] If the clutter spectrum is narrow compared with the PRF, then the improvement factor can be 20 dB or more even for the simple two-pulse canceller. If the clutter spectrum is wide, much of the clutter power will be in the passband of the MTI high-pass filter, and the improvement factor will be small.

TABLE 17-2 ■ Improvement Factor for Gaussian Clutter Power Spectrum

	Improvement Factor (dB)	
Standard Deviation of Clutter Power Spectrum (Hz)	Two-Pulse Canceller	Three-Pulse Canceller
PRF/6	3.7	5.7
PRF/10	7.5	12.5
PRF/20	13.2	21.7
PRF/100	24	51

[3] Clutter spectra having a standard deviation wider than *PRF*/6 do not decay to nearly zero at $f_d = \pm PRF/2$ and therefore will not be Gaussian in shape due to the replication of the spectrum with sampled data. The Gaussian spectrum assumed for this table is not valid for these very wide clutter spectra.

Additional MTI metrics can be defined. Improvement factor is the average of the improvement in signal-to-clutter ratio over one Doppler period. At some Doppler shifts, the target is above the clutter energy, while at others it is below the clutter and therefore not detectable. I does not indicate over what percentage of the Doppler spectrum a target can be detected. The concept of *MTI visibility factor* or *target visibility*, V, has been proposed to quantify this effect [15]. V is the percentage of the Doppler period over which the improvement factor for a target at a specific frequency is greater than or equal to the average improvement factor I.

17.4.4 Limitations to MTI Performance

The basic idea of MTI processing is that repeated measurements (pulses) of a stationary target yield the same echo amplitude and phase; thus, successive echo samples, when subtracted from one another, should cancel. Any effect internal or external to the radar that causes the received echo from a stationary target to vary will cause imperfect cancellation, limiting the improvement factor.

The simplest example is transmitter amplitude instability. If two transmitted pulses differ in amplitude by 10% (equivalent to $20\log_{10}(1.1/1) = 0.83$ dB), then the signal resulting from subtracting the two echoes from a perfectly stationary target will have an amplitude that is 10% that of the individual echoes. Consequently, clutter attenuation can be no better than $20\log_{10}(1/0.1) = 20$ dB. For a two-pulse canceller with an average signal gain G of 2 (3 dB), the maximum achievable improvement factor is 23 dB.

A more realistic analysis of the limitations due to pulse-to-pulse amplitude variations can be obtained by modeling the amplitude of the m-th transmitted pulse as $A[m] = k(1 + a[m])$, where $a[m]$ is a zero mean, white random process with variance σ_a^2 that represents the percentage variation in transmitted amplitude, and k is a constant. The received signal will have a complex amplitude of the form $k'(1+a[m])\exp(j\phi)$, where ϕ is the phase of the received slow-time sample and the constant k' absorbs all the radar range equation factors. The average power of this signal, which is the input to the pulse canceller, is

$$\mathrm{E}\{|y[m]|^2\} = k'^2\mathrm{E}\{1 + 2a[m] + a^2[m]\} = k'^2\left(1 + \sigma_a^2\right) \tag{17.25}$$

where $\mathrm{E}\{\cdot\}$ is the expected value operator. The expected value of the two-pulse canceller output power will be

$$\begin{aligned}\mathrm{E}\left\{|(y[m] - y[m-1])|^2\right\} &= \mathrm{E}\left\{\left|k'e^{j\phi}(a[m] - a[m-1])\right|^2\right\} \\ &= k'^2\mathrm{E}\{a^2[m]\} - 2\mathrm{E}\{a[m]a[m-1]\} + k'^2\mathrm{E}\left\{a^2[m-1]\right\} \\ &= 2k'^2\sigma_a^2\end{aligned} \tag{17.26}$$

The average clutter cancellation is thus

$$CA = \frac{\text{input power}}{\text{output power}} = \frac{k'^2\left(1 + \sigma_a^2\right)}{2k'^2\sigma_a^2} = \frac{1 + \sigma_a^2}{2\sigma_a^2} \quad \text{(amplitude jitter)} \tag{17.27}$$

For example, an amplitude variance of 1% ($\sigma_a^2 = 0.01$) limits two-pulse clutter cancellation to a factor of 50.5, or 17 dB. Because the average target gain G of the two-pulse canceller is $G = 2$ (3 dB), the limit to the improvement factor I is $17 + 3 = 20$ dB.[4]

[4]The effect of the jitter on the gain, G, should also be computed to be completely correct, but that effect is small and can usually be neglected.

Another limiting factor is phase drift in either the transmitter or receiver. This can occur, for example, due to instability in coherent local oscillators used either as part of the waveform generator on the transmit side or in the demodulation chains on the receiver side. A stationary radar viewing a stationary scatterer would expect the same phase as well as amplitude of two successive measurements. However, phase is measured by reference to an oscillator within the receiver. If the reference phase changes between measurements, the apparent measured phase will change, resulting again in imperfect cancellation of the two measurements when subtracted. This effect can be analyzed by modeling the error in the measured phase as a zero mean Gaussian random error. An analysis approach similar to that used for amplitude (see [16] for details) shows that the limitation on two-pulse cancellation due to the phase noise is

$$CA = \frac{\text{input power}}{\text{output power}} = \frac{1}{2\left(1 - e^{-\sigma_\phi^2}\right)} \text{ (phase jitter)} \tag{17.28}$$

Other sources of limitation due to radar system instabilities include instability in transmitter or oscillator frequencies; transmitter phase drift; coherent oscillator locking errors; PRI jitter; pulse width jitter; and quantization noise. Simple formulas to bound the achievable clutter attenuation due to each of these error sources are given in [2, 4]. These formulas can be used to construct an error budget and determine allowable tolerances on each error source. Another approach to improving canceller performance is to measure the actual errors when possible and compensate for them in the processing. For example, the actual power of each transmitted pulse could be measured by the radar and used to adjust the received signal voltages on a pulse-by-pulse basis to improve MTI cancellation.

External to the radar, the chief factor limiting MTI improvement factor is simply the width of the clutter spectrum itself. Wider spectra put more clutter energy outside of the MTI filter null so that less of the clutter energy is filtered out. This effect was illustrated numerically in Table 17-2. The clutter spectrum width is determined first and foremost by the inherent Doppler spread of the clutter scatterers. It can be increased by radar system effects and instabilities such as the amplitude and phase jitters previously discussed and by radar platform motion, as shown in the next subsection. For instance, a scanning antenna adds some amplitude modulation due to antenna pattern weighting to the clutter return that broadens the observed spectral width somewhat. In some cases, the clutter power spectrum may not be centered on zero Doppler shift. A good example is rain clutter: moving weather systems will have a nonzero average Doppler representing the rate at which the rain cell is approaching or receding from the radar system. Unless this average motion is detected and compensated, the MTI filter null will not be centered on the clutter spectrum, and cancellation will be poor.

17.4.5 MTI from a Moving Platform

The largest source of clutter offset and spreading is radar platform motion; formulas for these two effects were given in Chapter 8. The offset in center frequency of the clutter spectrum can be as much as a few kHz for fast aircraft, while the motion-induced spectral spread can be tens to a few hundreds of Hz. For space-based radars, both can be another order of magnitude larger. This clutter spreading adds to the intrinsic spread of the clutter spectrum due to internal motion and can often be the dominant effect determining the observed clutter spectral width and the resulting MTI performance limits.

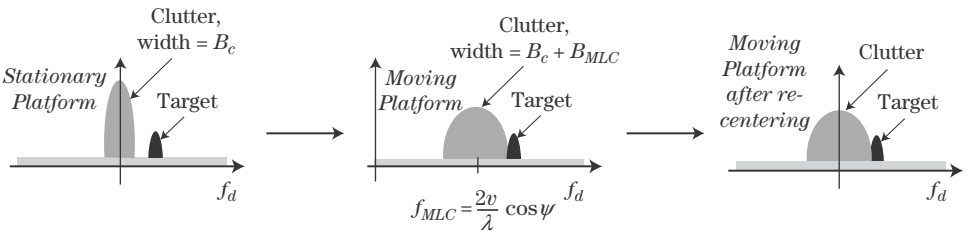

FIGURE 17-10 ■ Illustration of the effect of a moving radar platform on the Doppler spectrum and the detection of "slow movers."

This phenomenon is illustrated in Figure 17-10, which shows the shifting and spreading of the mainlobe clutter by the platform motion followed by the recentering of the spectrum at zero Doppler. The clutter spectrum width is B_c Hz when observed from a stationary radar, but when observed from the moving platform it increases to $B_c + B_{MLC}$ Hz, where $B_{MLC} = 2v\theta_3 \sin\psi/\lambda$ Hz. Because of this spreading, a relatively slow-moving target ("slow mover," typically a ground target) that was in the clear region (left portion of the figure) may now have to compete with clutter energy as well (center or right portion of the figure), making its detection more difficult.

Displaced phase center antenna (DPCA) processing is a technique for countering the effect of platform-induced clutter spectral spreading on MTI filtering. It is a special case of the more general *space-time adaptive processing* (STAP) [1,14].

In its simplest form, DPCA applies two-pulse MTI cancellation to data collected using two receive apertures on the side of the platform in a sidelooking configuration. The two apertures are denoted the "fore" and "aft" apertures. As the aircraft flies forward, the aft aperture passes through the same coordinates in space as the fore aperture, just delayed by $\Delta x/v$ seconds, where Δx is the spacing of the two receivers. If a data sample for a particular range bin collected on the aft aperture is subtracted from an earlier sample of the same range bin collected on the fore aperture when it was at the same location, the effect will be to subtract two samples of the same range bin taken from the same point in space. These two data samples will therefore be relatively well-correlated so that two-pulse cancellation is effective. Combining the data from two subapertures in this way implicitly avoids clutter spreading effects, improving clutter cancellation and slow target visibility. References for basic DPCA are [2,17,18].

Figure 17-11 illustrates the concept for a typical implementation using a sidelooking phased array antenna divided into two subapertures. Each half of the antenna has its own receiver, so there are in effect two receive apertures having respective phase centers R1 and R2, which are Δx m apart. Assume for the moment that the antenna also transmits using only one of the subapertures at a time.

Suppose the fore subaperture of the antenna is used to transmit and receive the first pulse and that its phase center R1 is centered at position x_0 at that time, as shown in Figure 17-11a. Now consider the motion of the platform over one PRI. The aft aperture will move vT m in T seconds and therefore be in the same position as the fore aperture one pulse earlier if the aperture spacing $\Delta x = vT$, as suggested by Figure 17-11b. More generally, the aft aperture will be in the same position as was the fore aperture M_s pulses earlier if

$$vM_sT = \Delta x \tag{17.29}$$

M_s is the "time slip" in pulses; M_sT is the time slip in seconds. Equation (17.29) is known as the *DPCA condition*.

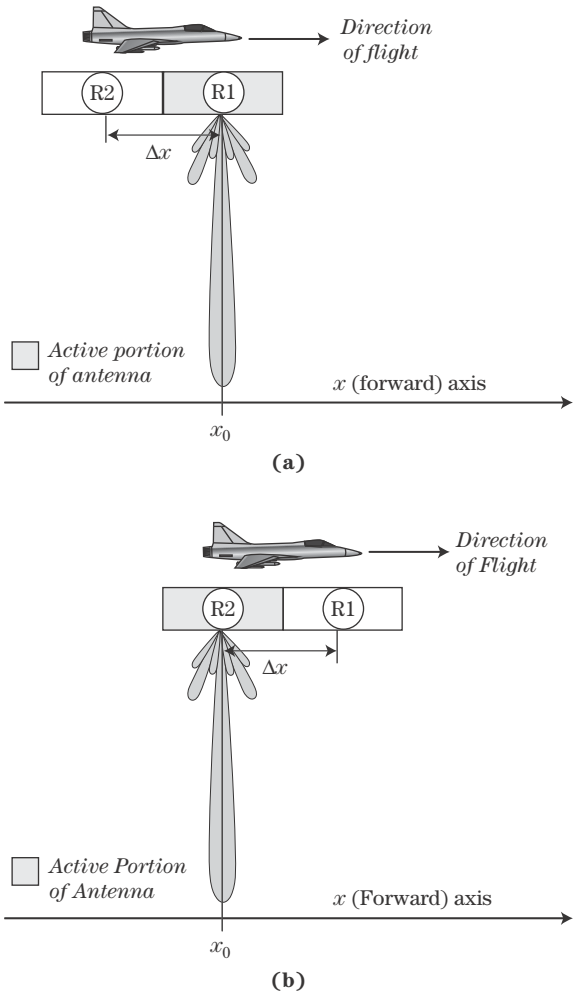

FIGURE 17-11 ■ Relationship of transmit and receive aperture phase centers in DPCA processing.

The significance of the DPCA condition is that if it is satisfied, the data stream received on the aft receive aperture is geometrically equivalent to the data stream received on the forward receive aperture M_s pulses earlier. Consequently, two-pulse cancellation can be implemented by taking each sample in a given range bin from the R1 data stream and subtracting the sample from the same range bin in the R2 data stream taken M_s pulses later, as illustrated in Figure 17-12 for $M_s = 3$. Even though these data samples were collected

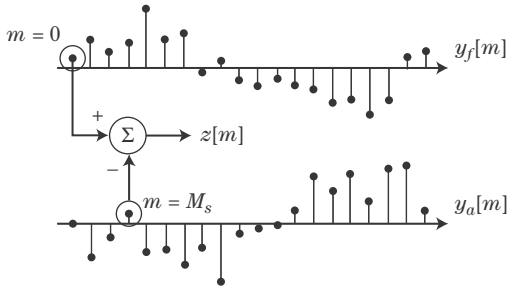

FIGURE 17-12 ■ Illustration of combining of data to achieve two-pulse cancellation across two received data streams in DPCA with a time slip $M_s = 3$ pulses.

on different receive apertures and different pulses, their effective transmit/receive phase centers are the same, so they appear equivalent to successive pulses from a *stationary* antenna. This phase center stationarity results in greater correlation of the R2 and delayed R1 data streams, providing better two-pulse cancellation and improving the detection of slow-moving ground targets.

The radar and platform operators can vary the platform velocity, PRI, and time slip M_s to attempt to satisfy the DPCA condition. However, this may be difficult to do in general. The platform velocity, generally the cruise velocity of an aircraft or orbital speed of a spacecraft, may be reasonably variable over only a small range. The PRI is heavily constrained by range and Doppler ambiguity concerns. Thus, M_s will often not be integer. For example, if $\Delta x = 3$ m, $v = 225$ m/s, and $T = 2$ ms, then $M_s = 6.67$ pulses. While conventional band-limited interpolation could be used to implement fractional-PRI timing adjustments, in practice there will also be amplitude and gain mismatches between channels that will make it impossible to achieve high cancellation ratios even if the time alignment is perfect. A typical DPCA implementation will therefore round M_s to the nearest integer for coarse alignment of the two data streams and then will modify the basic two-pulse canceller weights to minimize the clutter residue at the processor output and therefore to maximize the improvement factor. The coefficients used to combine the fore and aft data streams are computed from the data itself, making the system an *adaptive* DPCA (A-DPCA) processor. In addition, it is common to perform pulse-Doppler processing on both the fore and aft slow-time signals first and then to apply A-DPCA separately to each DFT bin in the clutter region of the spectrum. (There is no advantage to applying it in the clear region, where noise is the dominant interference.) Separate adaptive weights can then be computed for each clutter-region DFT bin, further improving the clutter cancellation. Details are given in [1,18].

In practice, the full antenna is often used on transmit to maximize the gain and directivity. Consequently, the phase center of the antenna on transmit is in the middle of the full antenna, halfway between the two receive phase centers R1 and R2. The effective phase center for the transmission of a pulse and its reception at R1 is then halfway between the center of the full antenna and R1; a similar result holds for reception on the aft aperture. The effective spacing of the phase centers for the fore and aft data stream is then only $\Delta x/2$, making the DCPA condition

$$vM_sT = \frac{\Delta x}{2} \quad \Rightarrow \quad 2vM_sT = \Delta x \tag{17.30}$$

In the example just given, the required time slip would now be 3.33 pulses.

17.5 PULSE-DOPPLER PROCESSING

Pulse-Doppler processing is the second major class of Doppler processing. Recall that in MTI processing, the fast-time/slow-time data matrix is high-pass filtered in the slow-time dimension, yielding a new fast-time/slow-time data sequence in which the clutter components have been attenuated. Pulse-Doppler processing differs in that filtering in the slow-time domain is replaced by explicit spectral analysis of the slow-time data for each range bin. Thus, the result of pulse-Doppler processing is a range/Doppler data matrix in which the dimensions are fast time and Doppler frequency. In the range/Doppler data, the energy from a moving target is separated from that of the clutter and competes only with the noise in the target's Doppler bin. In addition, an estimate of the Doppler shift (and thus

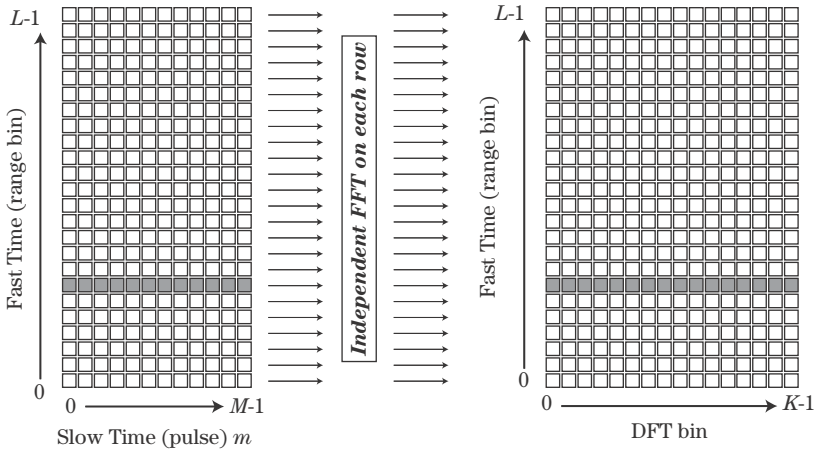

FIGURE 17-13 ■ Conversion of the fast-time/slow-time data matrix to a range-Doppler matrix by applying a DFT to each slow-time row.

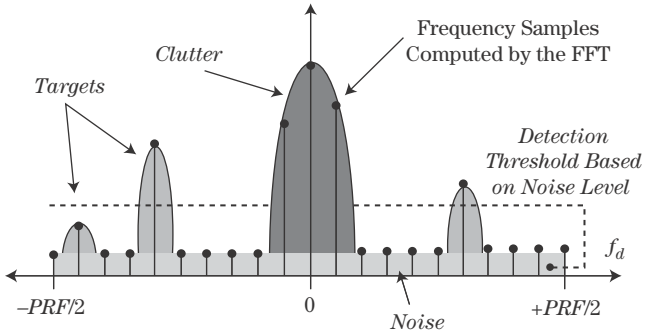

FIGURE 17-14 ■ The concept of pulse-Doppler processing for detection of moving targets.

radial velocity) of detected targets can be estimated based on the Doppler bin in which the detection occurs.

The spectral analysis is most commonly performed by computing the DFT of each slow-time row of the data using an FFT algorithm as shown in Figure 17-13, but other techniques can also be used [19]. The DFT size K can be equal to or greater than M.[5] Good general discussions of pulse-Doppler processing are contained in [20–22]. The DFT bins 0 through K-1 correspond to Doppler frequencies $-PRF/2$ to $+PRF/2$ as described in Chapter 14.

Figure 17-14 illustrates the concept of pulse-Doppler processing for moving target detection using a notional pulse-Doppler spectrum for a single range bin, similar to the one in Figure 17-2. The function in the background is the DTFT of the slow-time data, whereas the circles represent the K samples of the DTFT computed by the DFT. White thermal noise is present in every DFT sample at, on average, equal power. Assuming the clutter has been centered at zero Doppler, those spectral samples at or near 0 Hz will consist of both clutter and noise. In many cases of practical interest, the clutter dominates the noise in this clutter region. A target, if present, may appear anywhere in the spectrum, as appropriate to its Doppler shift. If the target is in the clear region, away from the clutter energy, only thermal noise will interfere with its detection. Each Doppler spectrum sample in the clear

[5] It is also possible to have $K < M$ in a sensible way, a process called *data turning*. This is used infrequently today because of increased computing power, but the technique is described in [1].