CHAPTER

27 Radio spectrometry

Spectrometry or spectral analysis is the statistical characterization of random (*stochastic*) signals such as the IF voltage in a radio astronomy receiver, as described in Chapter 26. The spectrometers discussed in this chapter are all *multiplex* spectrometers, meaning that they measure N points on the spectrum simultaneously. This is to be distinguished from swept-frequency spectrum analyzers, which measure spectral points sequentially. Multiplex spectrometers are used when long integration times are needed to pull a signal out of the noise, as in radio astronomy. They are also used for low-frequency spectrum analysis, where narrow channel bandwidths require long measurement times to process a sufficient number of independent samples. Most often the signal is Gaussian; if a fine-grained histogram of samples of the signal's amplitude is scaled to make the area below the curve equal to unity, the average curve will be the Gaussian probability density function, $f(V) = (2\pi\sigma^2)^{-1/2} \exp(-V^2/2\sigma^2)$. You can verify (Problem 27.1) that σ is the rms value of V, i.e., σ^2 is the power, $\langle V^2 \rangle$. But total power does not completely characterize the signal. A complete description is contained in the *power spectral density function*, $S(\omega)$, the distribution of power vs. frequency. A set of bandpass filters and power meters, i.e., a set of radiometers, serves to measure points on the spectral density function (usually called the PSD or simply the power spectrum). However, the simplest mathematical definition of the power spectrum uses the signal's *autocorrelation function*, $R(\tau)$, a function of time delay. The value of the autocorrelation function for a given time delay, τ, is defined as the average value of the "lagged product"

$$R(\tau) = \langle V(t) \cdot V(t+\tau) \rangle. \tag{27.1}$$

For a signal whose characteristics are unchanging (a *stationary process*), $R(\tau) = R(-\tau)$. We will see that the Fourier transform of the autocorrelation function is the power spectral density so, directly or indirectly, we describe the signal in terms of sine-wave basis functions, i.e., we find the magnitudes (but here not the phases) of a set of sine waves whose superposition has the same power spectrum as the original signal. Why are sine waves preferred to other sets of basis functions? Often the process under study is a spectral line which shows up

clearly in a few adjacent Fourier coefficients (frequency bins) or a Doppler shift which just displaces the spectrum. The frequency spectrum is clearly a natural way to represent such signals. Of course we deal with sine waves all the time as the characteristic functions of linear systems; a sinusoidal signal of a given frequency remains a sinusoid with the same frequency after passing through any arbitrary chain or network of linear elements such as amplifiers, filters, and transmission lines. A variety of instruments have been developed for spectrometry. In approximate historical order, these include analog filterbanks, autocorrelators, chirp-z spectrometers, Fourier transform spectrometers, and acousto-optical spectrometers.

27.1 Filters and filterbanks

We have already seen the analog filterbank in Chapter 26, as part of a radiometer bank. Usually the signal is converted down to a convenient intermediate frequency, e.g., tens of MHz, and the filters are the various *LC* bandpass filters discussed earlier or crystal filters or digital filters. A traditional disadvantage of an analog filterbank spectrometer has been that it is not "elastic," i.e., the filters have a fixed width; if narrower or wider filters are needed, another filterbank must be built. With high-speed digital-to-analog conversion and fast memory, this restriction can be lifted by using an interesting spectral expansion technique. Sequences of data are read into the memory at whatever slower rate is needed and then played back at a higher rate into a single wideband filterbank.

27.2 Autocorrelation spectrometry

Autocorrelation spectrometers estimate the power spectrum by taking pairs of samples with a given time separation, multiplying them, and averaging these "lagged products." (In Fourier transform spectral analysis the samples are multiplied by sine waves but the resulting coefficients are squared before averaging. In both cases powers, rather than voltages, are averaged.) When the time separation is zero, the average lagged product is just the total power. But if we measure other lagged products as well (where the time separation is not zero), we can indirectly find the power spectrum because the autocorrelation function (ACF) and the power spectrum form a Fourier transform pair:

$$S(\omega) = \int_{-\infty}^{\infty} R(\tau) \mathrm{e}^{-\mathrm{j}\omega\tau} \mathrm{d}\tau \tag{27.2}$$

$$R(\tau) = \frac{1}{2\pi} \int_{-\infty}^{\infty} S(\omega) \mathrm{e}^{\mathrm{j}\omega\tau} \mathrm{d}\omega. \tag{27.3}$$

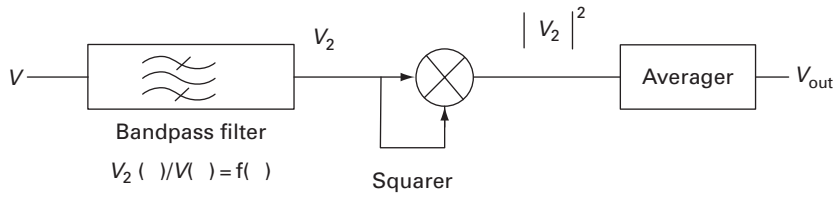

Figure 27.1. Single-channel radiometer.

This relation, the Wiener–Khinchin theorem, is often used to define the power spectrum. To see that S(ω) as defined by Equation (27.2) agrees with our intuitive ideas of the power spectrum, consider the single-channel radiometer shown in Figure 27.1. The (in situ) voltage transfer function of the bandpass filter is defined in the frequency domain by $H(\omega) = V_2(\omega)/V(\omega)$. The power response function or filter shape is therefore given by $|H(\omega)|^2$, so we expect that the time average of $(V_2)^2$ will be given by

$$\left\langle |V_2(t)|^2 \right\rangle = \int_{-\infty}^{\infty} S(\omega)|H(\omega)|^2 (2\pi)^{-1} \mathrm{d}\omega \tag{27.4}$$

where $S(\omega)$ is the spectral power density function. Note that $(2\pi)^{-1}\mathrm{d}\omega = \mathrm{d}f$; the power spectral density has units of watts/Hz.

Let us now show that the output of this radiometer is indeed given by Equation (27.4), and that $S(\omega)$ is given by Equation (27.2). We can express V_2 in the time domain as

$$V_2(t) = \int_{-\infty}^{t} V(t')h(t-t')\mathrm{d}t', \tag{27.5}$$

where $h(\tau)$, the impulse response function, is the inverse Fourier transform of $H(\omega)$.[1] The output of the squarer is then given by

$$(V_2(t))^2 = \int_{-\infty}^{t}\int_{-\infty}^{t} V(t')V(t'')h(t-t')h(t-t'')\mathrm{d}t'\mathrm{d}t''. \tag{27.6}$$

Since we assume $V_{\mathrm{in}}(t)$ is a stationary process, the value of t is arbitrary. Let us pick $t = 0$, which gives

$$(V_2(t))^2 = \int_{-\infty}^{0}\int_{-\infty}^{0} V(t')V(t'')h(-t')h(-t'')\mathrm{d}t'\mathrm{d}t''. \tag{27.7}$$

[1] The impulse function, $h(t)$, is the filter output voltage at t, in response to a delta function input at $t = 0$. Note that $h(t)$ is real and that $h(t) = 0$ for $t < 0$.

From causality, $h(t) = 0$ for $|t| < 0$, so the upper limits of integration can be changed from 0 to infinity. Doing this, making a change of variable, $t'' = t' + \tau$, and taking the time average, we find

$$\left\langle (V_2(t))^2 \right\rangle = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \langle V(t')V(t'+\tau)\rangle h(-t')\, h(-t'-\tau)\mathrm{d}t'\mathrm{d}\tau. \qquad (27.8)$$

Using the definition of the ACF, $R(\tau) = \langle V(t)V(t+\tau)\rangle$, we find

$$\left\langle (V_2(t))^2 \right\rangle = \int_{-\infty}^{\infty} R(\tau) \left[\int_{-\infty}^{\infty} h(-t')h(-t'-\tau)dt' \right] \mathrm{d}\tau. \qquad (27.9)$$

Applying the convolution theorem,[2] the Fourier transform of the term in square brackets is $|H(\omega)|^2$. Therefore, the bracketed term is the inverse Fourier transform of $|H(\omega)|^2$ or

$$\left[\int_{-\infty}^{\infty} h(-t')h(-t'-\tau)\mathrm{d}t' \right] = \frac{1}{2\pi} \int_{-\infty}^{\infty} \mathrm{e}^{\mathrm{j}\omega\tau} |H(\omega)|^2 \mathrm{d}\tau. \qquad (27.10)$$

Substituting this expression into Equation (27.9) produces

$$\left\langle (V_2(t))^2 \right\rangle = \int_{-\infty}^{\infty} R(\tau) \left[\frac{1}{2\pi} \int_{-\infty}^{\infty} \mathrm{e}^{\mathrm{j}\omega\tau} |H(\omega)|^2 \mathrm{d}\omega \right] \mathrm{d}\tau. \qquad (27.11)$$

Finally, interchanging the order of integration gives us

$$\left\langle (V_2(t))^2 \right\rangle = \int_{-\infty}^{\infty} |H(\omega)|^2 \left[\int_{-\infty}^{\infty} \mathrm{e}^{\mathrm{j}\omega\tau} R(\tau)\mathrm{d}\tau \right] (2\pi)^{-1}\mathrm{d}\omega. \qquad (27.12)$$

After all this effort, we can now compare Equation (27.12) with Equation (27.4) and confirm that the power spectrum defined by Equation (27.2) agrees with what we would expect to measure with a bandpass filter followed by a squarer and an averager. (Don't worry about the sign discrepancy in the $\mathrm{e}^{\mathrm{j}\omega t}$ term in Equation (27.12); since $V_2(t)$ and $R(\tau)$ are real, we can replace this term by $\mathrm{e}^{-\mathrm{j}\omega t}$.)

27.2.1 Hardware autocorrelators

A hardware autocorrelator is a special-purpose parallel signal processor that calculates an averaged ACF, usually in real time. A single Fourier transform operation in a computer then turns the ACF into the power spectrum with the same number of frequency bins (points on the spectrum) as the number of points

[2] If $H(\omega)$ is the Fourier transform of $h(\tau)$, then $|H(\omega)|^2$ is the Fourier transform of $\int h(t)h(t+\tau)\mathrm{d}t$.

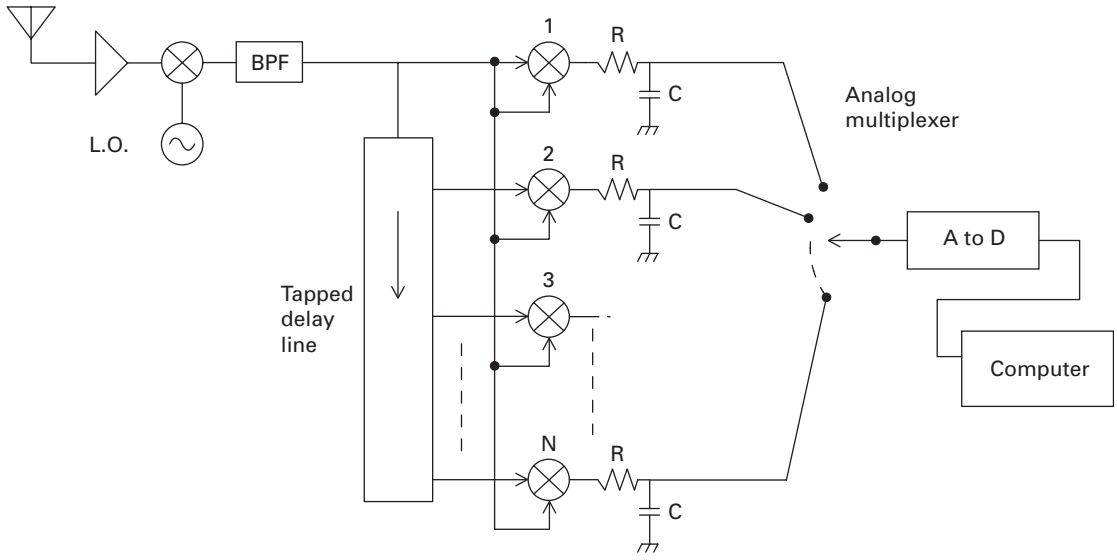

Figure 27.2. Analog autocorrelator.

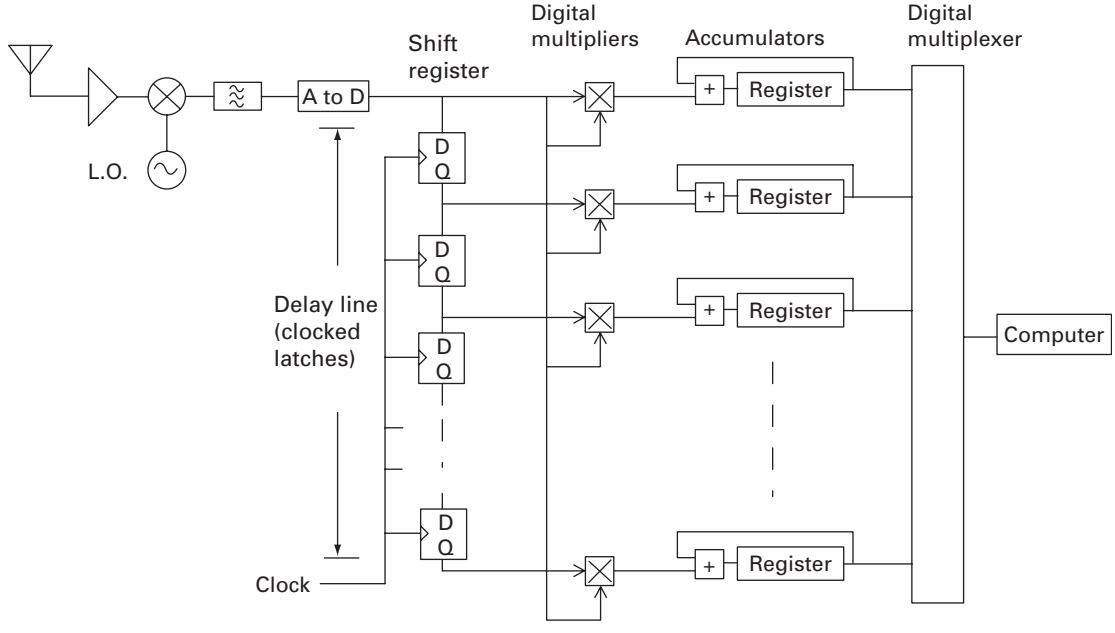

Figure 27.3. Digital autocorrelator.

on the ACF. Autocorrelators use a simple expandable architecture. Figure 27.2 shows an analog version of the autocorrelator.

Figure 27.3 shows a digital autocorrelator. The analog delay line is replaced with a digital shift register. The analog multipliers and averagers are replaced by digital multipliers and accumulators.

The autocorrelators of Figures 27.2 and 27.3 produce values only for discrete points on the ACF. But, as long as the input signal is band-limited (usually by a bandpass "anti-aliasing" filter), the spectrum can be represented as a Fourier *series*, so sampled ACF points are sufficient to compute the continuous power spectrum. In practice, however, we cannot measure an infinite number of points on the ACF. The spectral estimate obtained by transforming ACF points that extend only to τ_{max}, i.e., by transforming an ACF "windowed" with a rectangular function that is unity for $\tau_{max} < \tau < \tau_{max}$ and zero outside this range, produces a function which is the true spectrum convolved by the function $\sin(\omega \tau_{max})/(\omega \tau_{max})$. A sharp line in the spectrum therefore appears as a broadened line, surrounded by sidelobes. It is common to apply a smoother windowing function such as $\cos(\frac{1}{2}\pi \tau/\tau_{max})$ to reduce the sidelobes even though the main load will be broadened – a trade-off between spectral resolution and leakage between channels.

27.2.2 One-bit autocorrelation

One-bit autocorrelation is a technique that greatly reduces the complexity of the digital circuitry. The input analog voltage is fed to a comparator whose output (one bit) indicates the sign, i.e., the polarity, of the voltage. The continuous range of analog voltages is compressed down to just two values: plus one and minus one. (In the digital hardware, a digital "1" indicates a value of +1 and a digital "0" indicates a value of −1.) The correlator's delay line shift register needs to be only one bit wide and the one-bit × one-bit multipliers can be exclusive NOR gates, since the output of an exclusive NOR gate is + 1 when the inputs are the same and −1 when they are different. The integrators that follow the multipliers are simple counters. Although the output of a one-bit correlator is a distorted version of the actual correlation function, the distortion can be entirely corrected. Because the input signal has Gaussian amplitude statistics, the output of the one-bit correlator turns out to be:

$$\rho_{1-\text{bit}}(\tau) = (2/\pi)\sin^{-1}[\rho(\tau)] \tag{27.13}$$

where ρ is the normalized autocorrelation function: $\rho(\tau) = R(\tau)/R(0) = R(\tau)/\sigma^2$. This functional relation was derived during World War II by J. H. van Vleck while studying the spectrum of clipped noise. (Jammers can have much better power efficiency if jamming with clipped noise works as well as jamming with noise having Gaussian amplitude statistics.)

The averaged values from a one-bit autocorrelator can, therefore, be inverted to get the true correlation function.

$$\rho(\tau) = \sin[(\pi/2)\rho_{1-\text{bit}}(\tau)]. \tag{27.14}$$

A straightforward proof of the van Vleck relation (Problem 27.4) proceeds from the Gaussian bivariate probability density,

$$f(x_1, x_2) = \frac{1}{2\pi\sigma_1\sigma_2(1-\rho^2)^{1/2}} \exp\left(\frac{-1}{2(1-\rho^2)}\left(\frac{{x_1}^2}{{\sigma_1}^2} - 2\rho\frac{x_1x_2}{\sigma_1\sigma_2} + \frac{{x_2}^2}{{\sigma_2}^2}\right)\right). \tag{27.15}$$

Here x_1 and x_2 are the voltages V(t) and V(t+τ), so $\sigma_1 = \sigma_2$. One gives up some radiometric sensitivity for the simplification of one-bit processing; the integration time has to be increased by a factor of $\pi^2/4$ compared to a correlator using multi-bit quantization.

Several generations of digital correlators have been used in radio astronomy. Custom CMOS correlator chips have thousands of lag channels and clock rates above 100 MHZ. A variety of serial and parallel multiplexing techniques are used to combine these chips to increase the number of lag channels and to increase input data rates to several hundred MHZ.

27.3 Fourier transform spectrometry

Fourier transform methods are more efficient than autocorrelation spectrometry simply because the FFT (Fast Fourier Transform algorithm) is an efficient way to compute the DFT (Discrete Fourier Transform). Processing a sequence of N data points requires on the order of N log(N) operations for an FFT vs. N^2 operations for either autocorrelation or a brute force DFT.

27.3.1 DFT spectrometer

The Fourier transform is, by itself, a spectrometer. The DFT of a sequence of voltages, V_n, is given by

$$u_k = \sum_{n=0}^{N-1} V_n e^{-j2\pi nk/N}. \tag{27.16}$$

We can write this DFT in terms of time and frequency as

$$u_{\omega_k} = \sum_{n=0}^{N-1} V_{t_n} e^{-j\omega_k t_n}, \tag{27.17}$$

where $t_n = n/f_{\text{sample}}$ and $\omega_k = 2\pi k f_{\text{sample}}/N$. Note that the transform is equivalent to a bank of N pairs of frequency converters with sine and cosine L.O.s. The k-th pair multiplies the input signal by cos($\omega_k t$) and sin($\omega_k t$), thus shifting the input band to place ω_k at zero frequency. For every converter, N consecutive multiplier output values are summed so one output value is produced every N/f_{sample} seconds. This summation is equivalent to lowpass filtering and decimation (reducing the sample rate). Successive values of u_k are, therefore, equivalent to (complex) voltages from the signal at and around ω_k, after being shifted to zero frequency and then lowpass filtered. When the input voltage is a random

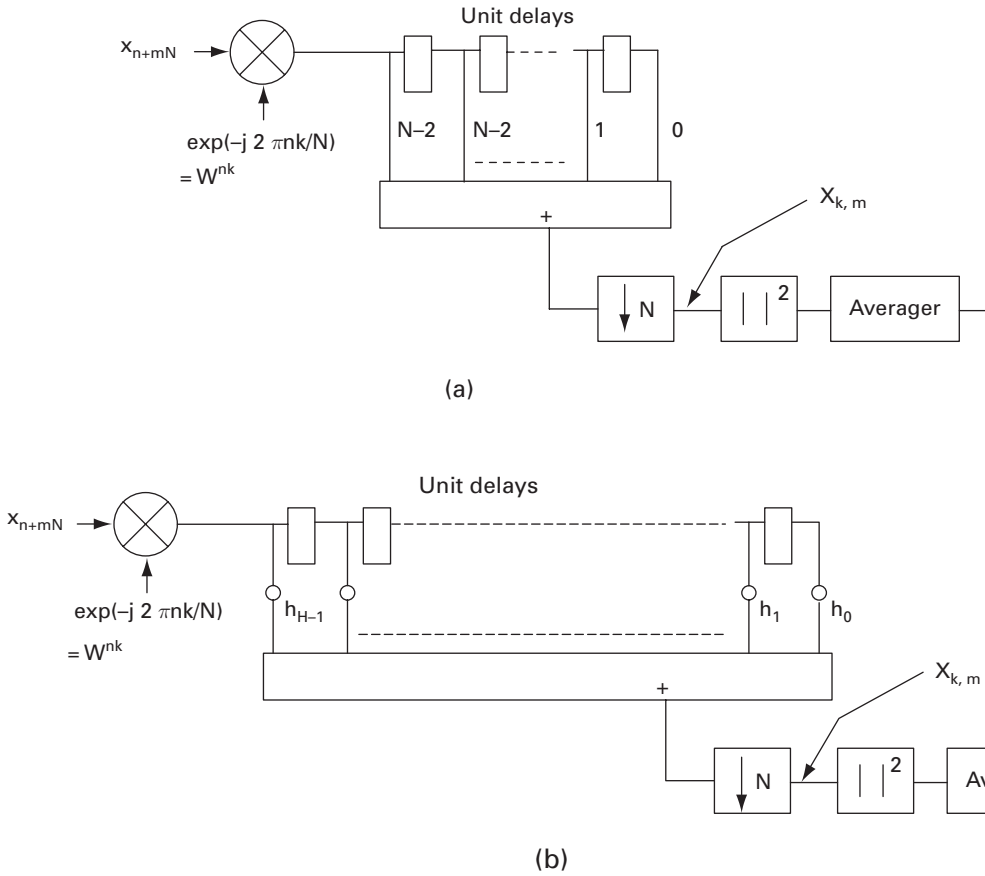

Figure 27.4. (a) One channel of a straight DFT spectrometer (periodogram averager); (b) filter channel using a longer sum and with arbitrary weights on the inputs to the summer.

process, u_k is also a random process. Its mean is zero. The squared magnitude of u_k represents power, and the set of points, $\{|u_k|^2\}$, is known as a *periodogram*. To estimate the power spectrum of a random process, many periodograms must be averaged together. Figure 27.4(a) shows one channel of a straight periodogram spectrometer.

A shortcoming of the DFT spectrometer (periodogram averager) is that the equivalent filter shape of each channel is not rectangular but, instead, has a $\sin(x)/x$ form, due to the uniform boxcar lowpass filter weights inherent in the DFT. It therefore suffers "leakage" from adjacent and even not-so-adjacent channels. We saw the same problem with the autocorrelation spectrometer. There are obvious remedies for this problem. The application of a weighting function to the input data blocks to give relatively low weight to the data points near the beginning and end of the blocks is equivalent to tapering the weights of the low-pass filter and therefore reduces the amplitude of the filter sidelobes. Such weighting, however, broadens the main lobe and increases the variance of the spectral estimates, since the spectrum is estimated mostly from the data points around the center of the data blocks. This loss of sensitivity can be remedied by

computing periodograms from data blocks that overlap each other and including these periodograms in the average. (Even with uniform data weighting, there is a loss of sensitivity when periodograms are computed only from nonoverlapping blocks of data.)

27.3.2 Combined FIR filter/DFT spectrometer

An "FIR/Fourier transform" spectrometer design that fixes the problems of the DFT spectrometer begins with the filter channel shown in Figure 27.4(b). Like the DFT filter, this filter channel will produce one output value for every N input samples. But now, the mixer is followed by a lowpass FIR (finite impulse response) filter which forms a running weighted sum of H successive mixer output values by summing the numbers, weighted by factors, h_i, at the taps of a digital delay line.[3] This would be the same as the DFT filter of the periodogram if all the weights, h_i, were equal to unity and the number of taps, H, were equal to N. Here, however, the number of taps will be an integral multiple of N, $H = RN$. This can be quite a large number, e.g., $H = 4 \times 1024 = 4096$, and the FIR filter can therefore produce a nearly ideal rectangular channel response. As with the straight DFT spectrometer, the bandwidth reduction from lowpass filtering lets us decimate the output of the summer, keeping only one sum per N-point block of input data. Again, the magnitudes of the output of the summer must be squared and averaged to form an estimate of the power spectrum. Note that, because the filter window, H, is longer than the block length, N, the blocks processed by the filter are overlapped; while one output value is produced for every N input samples, the output values are produced from H input samples, where $H > N$. We will designate the input voltage samples as x_n. They enter in blocks of N samples. The output of the summer is $X_{k,m}$, where m is the block number and k, which runs from zero to N, designates the point on the spectrum. Defining W as $W = \exp(-2\pi\, \mathrm{j}\, /N)$, inspection of Figure 27.4 allows us to write

$$X_{k,m} = \sum_{n=0}^{H-1} h_n W^{-k(mN+n)} x_{mN+n}. \tag{27.18}$$

Now, since $W^{-kmM} = 1$, we can rewrite this equation as

$$X_{k,m} = \sum_{n=0}^{H-1} h_n W^{-kn} x_{mN+n}. \tag{27.19}$$

Let us next express n as $rN + \rho$ where the index r goes from 0 to $R-1$, where $R = H/M$, and the index ρ goes from 0 to $M-1$. Using this notation, we have

[3] The finite (duration) impulse response is obvious; when an input sample makes its way down the line and falls off at the end, it no longer contributes the output. A FIR filter can be given a desired frequency response function by selecting the tap weights to form the desired impulse function, i.e., the Fourier transform of the desired frequency response.

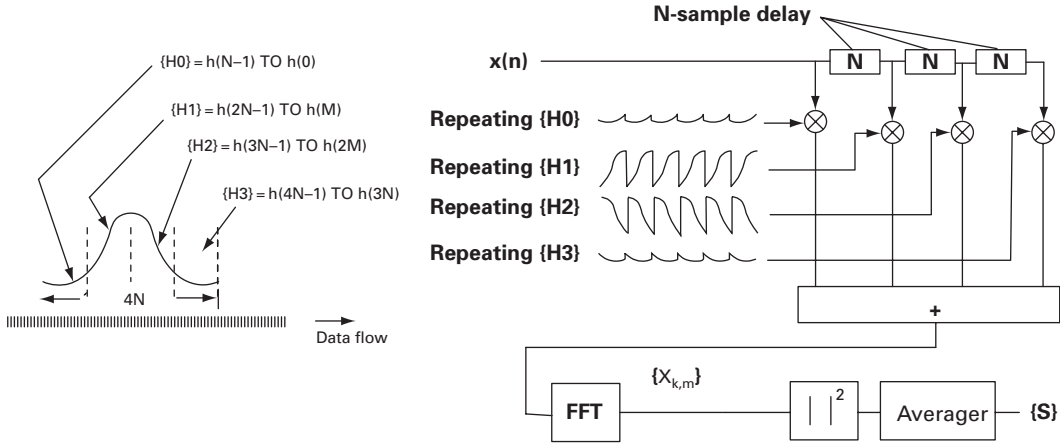

Figure 27.5. Weighted overlap-add filterbank.

$$X_{k,m} = \sum_{\rho=0}^{N-1} \left(\sum_{r=0}^{H/N-1} h_{rN+\rho} x_{mN+n} \right) W^{-k\rho}. \tag{27.20}$$

Note that Equation (27.20) is just an N-point DFT of the term in brackets, and that this term is independent of k, the frequency index. Consider this term for the case where $H = 4N$. The bracketed term is given by

$$( ) = h_\rho X_{Nm+\rho} + h_{N+\rho} X_{N(m+1)+\rho} + h_{2N+\rho} X_{N(m+2)+\rho} + h_{3N+\rho} X_{N(m+3)+\rho}. \tag{27.21}$$

In the first term on the right-hand side, successive samples are multiplied by h_0 through h_{N-1}. In the second term, delayed samples are multiplied by h_N through h_{2n-1}, and so on. An architecture that forms this bracketed term is shown in Figure 27.5. Four multipliers are used, together with three delay lines, each of length N. The set of H FIR coefficients is split into four sections, each of which feeds one of the four multipliers in a continuous cycle. An adder sums the products flowing from the multipliers. This preprocessor section feeds a standard FFT block. This architecture is called a *weighted overlap-add* (WOLA) filterbank. Another way of producing the bracketed term yields the *polyphase filter bank* [3].

27.4 *I* and *Q* mixing

In the discussions above, it was implicit that the spectrum of the band-limited signal voltage, $V(t)$, extends from zero to some cutoff frequency, B. This baseband signal can be produced from an IF signal that has been bandpass filtered and then down-converted to put the lower edge of the passband at zero frequency. For digital spectrum analysis, the sampling frequency must be at least $2B$ to avoid aliasing (the stroboscope effect). In the case of the DFT spectrometer and the improved WOLA version, the N channels extend in frequency from

0 to the sampling rate, i.e., from 0 to $2B$. Therefore, only the first $N/2$ output channels can be used, since those channels correspond to the frequency range 0 to B. (The second $N/2$ channels will contain the same information; channel $N–k$ will have the same data as channel k.) This effective waste of half the DFT channels can be avoided with a technique known as *complex sampling*. In this technique, the IF signal of bandwidth B is down-converted into baseband "I" and "Q" (in-phase and quadrature) channels by using two mixers having L.O. signals $\cos(\omega_c t)$ and $\sin(\omega_c t)$, where ω_c is the center of the IF band. We can express the original IF signal in terms of the mixer outputs, $I(t)$ and $Q(t)$, as

$$V_{\mathrm{IF}}(t) = I(t)\cos(\omega_c t) + Q(t)\sin(\omega_c t). \tag{27.22}$$

To find the spectral value at the point $\omega_c + \Delta$, we must multiply (mix) the IF voltage by $e^{-j(\omega c + \Delta)t}$. Carrying out this multiplication and ignoring terms of frequency $2\omega_c$, we find

$$V_{\mathrm{IF}}(t)e^{-j(\omega_c+\Delta)t} = [I(t) - jQ(t)]e^{-j\Delta t}. \tag{27.23}$$

The term on the right-hand side is equivalent to mixing a complex-valued signal, $I(t)–jQ(t)$, so that the spectral point at $\omega_c + \rho$ is shifted to zero frequency. The beauty of this is that both $I(t)$ and $Q(t)$ have bandwidth $B/2$, since $B/2$ is the greatest separation between ω_c and any point in the IF band. Therefore we can feed $I(t) - jQ(t)$ as a sequence of complex numbers to the straight DFT spectrometer or the improved WOLA version and the N output values will span the entire bandwidth B. The spectrum from $\Delta = 0$ to $B/2$ will appear in the first $N/2$ channels, while the spectrum from $\Delta = -B/2$ to 0 will appear in the second $N/2$ channels. When I and Q mixing is used in autocorrelation spectrometry, both autocorrelation functions and cross-correlation functions are computed from the I and Q signals.

27.5 Acousto-optical spectrometry

The acousto-optical spectrometer (AOS) uses an optical crystal which becomes a diffraction grating by virtue of mechanical waves (sound waves at radio frequencies) propagating through it. These waves produce corrugations in the refractive index along the length of the crystal. The waves are launched into this *Bragg cell* by an electrical-to-acoustic transducer (usually a piezoelectric crystal), driven by the IF signal to be analyzed. Figure 27.6 shows how the crystal is illuminated by a laser and how a multi-element CCD array detects (and averages) the diffraction pattern.

The linear CCD array, like those used in supermarket checkout scanners, accumulates charge along its length at a rate given by the incident light intensity. For read-out, the charge packets are clocked down the length of the CCD in a bucket brigade fashion. The voltage at the end of the array, which is proportional to the charge at that position, is digitized and made available to the data-taking computer.

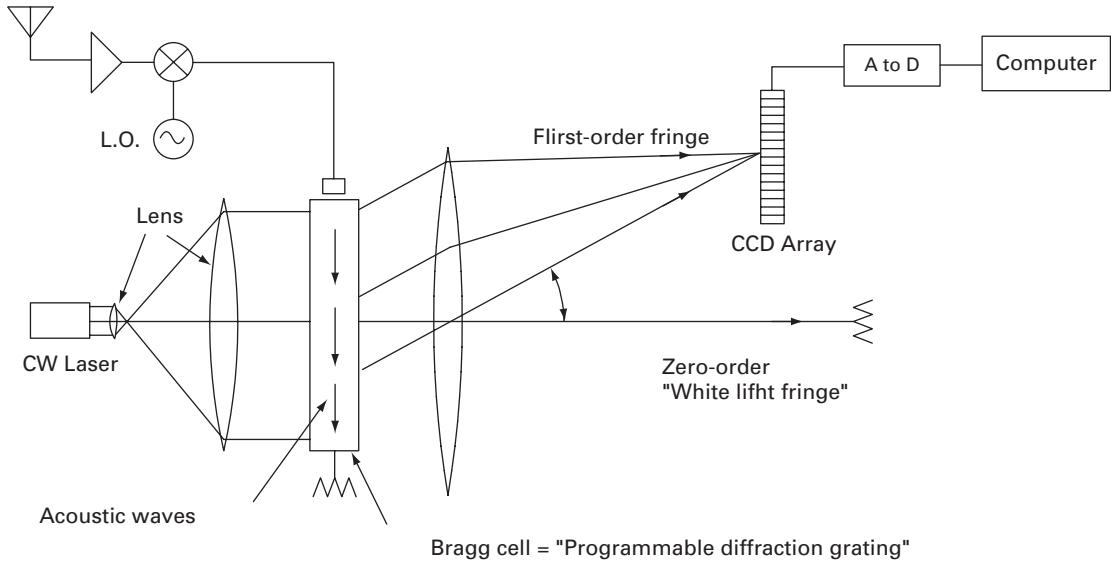

Figure 27.6. Acousto-optical spectrometer (AOS).

A figure of merit for any spectrometer is its number of frequency bins, i.e., its analysis range divided by its resolution. For the AOS, the effective number of channels is given by the time–bandwidth product of the Bragg cell. To see this we note that the angular size (in radians) of the light incident on the CCD is given by $\lambda_{\text{laser}}/L_{\text{cell}}$, the ratio of the laser wavelength to the length of the cell, just as the beam size for an aperture antenna with diameter D is given by λ/D. The first-order diffraction condition (so that all the rays arrive in phase at the CCD) is given by $\lambda_{\text{grating}}\sin\theta = 1 \times \lambda_{\text{laser}}$ where λ_{grating} is the corrugation wavelength in the cell. Now λ_{grating} can be written as v_{cell}/f where v_{cell} is the sound velocity and f is the input RF frequency. Approximating sin θ by θ, we find the total angular range over the total frequency range given by $\Delta\theta/\Delta f = \lambda_{\text{laser}}/v_{\text{cell}}$. The number of channels is then given by the range divided by resolution,

$$N = \frac{\lambda_{\text{laser}}\Delta f L_{\text{cell}}}{v_{\text{cell}}\lambda_{\text{laser}}} = \frac{L_{\text{cell}}}{v_{\text{cell}}}\Delta f = T\Delta f, \tag{27.24}$$

where T is the propagation time through the length of the cell. The quantity $T\,\Delta f$ is known as the *time–bandwidth product* of the cell. Materials used for Bragg cells include quartz, lithium niobate, and glass. A spectrometer for radio astronomy might use cells with a 200-MHz bandwidth and a time–bandwidth product of one thousand.

27.6 Chirp-z spectrometry

The chirp-z spectrometer is shown in Figure 27.7. Also known as the microscan or compressive receiver, it uses a dispersive filter and a swept L.O. While it appears to be a swept spectrum analyzer, it is actually a multiplex analyzer. The

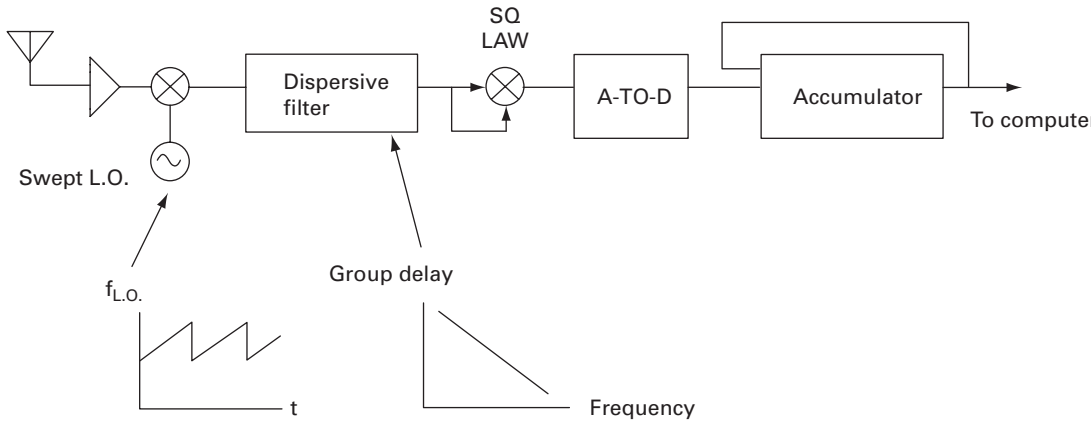

Figure 27.7. Chirp-z spectrometer.

filter has a group delay characteristic that is linear with respect to frequency. In Figure 27.7, the delay is greatest for frequencies at the low end of the filter band. The local oscillator frequency is given a linear sweep or *chirp*, repeated in a sawtooth fashion with period *T*. The ramp rate of the L.O. (MHz/sec) is made to be the reciprocal of the group delay slope (seconds per megahertz) of the filter.

Consider a continuous wave (cw) input signal. Since the L.O. is swept, the mixer output will also sweep in frequency. In the case of a high-side mixer (L.O. frequency above the input signal frequency), the mixer output sweeps up in frequency at the same rate as the L.O. sweep. As the L.O. sweeps upward, the first signals entering the filter are at the filter's low-frequency end. These signals, as they travel through the filter, are delayed the most. Frequencies entering the filter farther up the sweep are delayed less. But all frequencies arrive simultaneously at the filter output, producing a single sharp output pulse. If the input signal consists of two cw tones, there will be two output pulses per sweep. The position of the pulses in time will indicate the frequencies of the cw input signals. Of course this linear system can handle any combination, i.e., spectrum, of input signals. The output data rate from the squarer is equal to the input data rate so additional circuitry (the digital accumulator in Figure 27.7) is required for signal averaging.

The principal application of chirp-z spectrometers has been in electronic warfare – as surveillance receivers to sense radar or communication signals over a wide spectrum and with a high probability of intercept. To find the effective number of channels in this spectrometer we can first consider the output pulse produced by a single cw input signal. This output pulse is produced by a coherent superposition of frequencies over the entire bandwidth, B, of the filter. The width of this pulse is therefore given by $1/B$. Dividing the total output time T by the pulse width, the effective number of channels is given by TB, the time–bandwidth product, just as with the Bragg cell spectrometer.

Dispersive filters used for chirp-z spectrometry are usually surface acoustic wave devices. Piezoelectric transducers convert the input signal into mechanical

surface waves and then back again to an electrical output signal. In one design, diffraction gratings on the surface of the crystal are arranged to make the input-to-output path length longer for low-frequency waves than for higher frequency waves. Surface-wave dispersive filters can have bandwidths of hundreds of MHZ and time–bandwidth products of several hundred. Dispersive filters are also made from charge-coupled devices (CCDs). Generically these filters are allpass networks and the first dispersive filters were cascades of many second-order allpass sections. (The number of sections required is equal to the time–bandwidth product.) Note: the filters are not just allpass networks; they are usually given some amplitude "taper," i.e., the amplitude response is made to roll off at the band edges, in order to eliminate the "sidelobes" of the instrumental function that result from uniform amplitude response. The same is done with the AOS; the transmission of the Bragg cell is reduced near the ends.

Problems

Problem 27.1. For the Gaussian probability density, $f(V) = (2\pi\sigma^2)^{-1/2} \exp(-V^2/2\sigma^2)$, verify that

$$\int_{-\infty}^{\infty} f(V)\mathrm{d}V = 1 \quad \text{and} \quad \int_{-\infty}^{\infty} f(V)V^2\mathrm{d}V = \sigma^2.$$

Problem 27.2. (a) The noise-like signals observed in radio astronomy can be modeled as a comb of delta functions in frequency space – a picket fence of sine waves. Use your computer to generate 50 sine waves spanning the angular frequency range 1.00–1.10. Let these frequencies be somewhat random, e.g., $\omega_n = 2\pi(1+0.002n + 0.05\ \mathrm{rnd}(1))$ where rnd(1) is a random variable in the range zero to one. Give each sine wave a random starting phase, $\phi_n = 10\ \mathrm{rnd}(1)$. Plot the sum of the sine waves as a function of time for $t = 0$ to $t = 20$ using a time interval of 0.1. This narrowband spectrum should look quite sinusoidal but with a slowly varying phase and amplitude.

(b) Repeat this exercise but with the angular frequency range extending from 1 to 2, i.e., $\omega_n = 2\pi(1 + 0.05n + 0.05\ \mathrm{rnd}(1))$. This wideband spectrum should look like random noise.

(c) In either case, (a) or (b), the resulting voltage (sum) behaves like a random variable with Gaussian statistics. Plot a histogram of voltage samples to verify this.

Problem 27.3. Suppose a random signal is produced by applying an ideal lowpass filter to white noise. The resulting power spectrum is flat from dc to the filter cutoff frequency ω_c. Show that the normalized autocorrelation function is given by $\rho(\tau) = \sin(\omega_c\tau)/(\omega_c\tau)$.

Problem 27.4. Derive the VanVleck relation, Equation (27.13), that gives the one-bit autocorrelation function in terms of the (normalized) autocorrelation function. Use the bivariate Gaussian probability density function, Equation (27.15), where x_1 and x_2 represent the voltages $V(t)$ and $V(t+\tau)$, and $\sigma_1 = \sigma_2$. Since the product $\mathrm{sign}(x_1)\ \mathrm{sign}(x_2)$ is equal to +1 in the first and third quadrants and −1 in the second and fourth quadrants,

the expectation of sign(x_1) sign(x_2) is the integral of $f(x_1,x_2)$ over x_1 and x_2, weighted by +1 or −1 according to the quadrant. Hint: change to polar coordinates and keep a table of integrals handy.

Problem 27.5. If the dispersive filter used in the chirp-z spectrometer or pulse compression radar is given an impulse (delta function in time), what sort of output waveform does it produce?

References

[1] Blackman R. B. and Tukey, J. W., *The Measurement of Power Spectra*, New York: Dover, 1959.
[2] Childers, D. B. ed., *Modern Spectrum Analysis*, New York: IEEE Press, 1978.
[3] Crochiere, R. E. and Rabiner, L. R., *Multirate Digital Signal Processing*, Englewood Cliffs: Prentice-Hall, 1983.
[4] Jack, M. A., Grant P. M. and Collins, J. H., The theory, design, and applications of surface acoustic wave Fourier-transform processors. *Proc. IEEE*, vol. 68, pp. 450–468, 1980.
[5] Kesler, S. B. ed. *Modern Spectrum Analysis, II*, New York: IEEE Press, 1986.
[6] Thomas, J. B., *An Introduction to Statistical Communication Theory*, New York: John Wiley, 1969.
[7] Thompson, A. R., Moran, S. M. and Swenson, Jr. G. W., *Interferometry and Synthesis in Radio Astronomy*, Malabar, Florida: Krieger Publishing Company, 1991.