CHAPTER 13

Signal Processing

One man's noise is another man's data.

—Anonymous

13.1 INTRODUCTION

Most of the time, the end result of a measurement is a digital data set. Preparing the squishy analog stuff coming from the front end amplifier to be digitized is the job of the *analog signal processing subsystem*,which is the subject of this chapter. The objective is to pull the desired signal out of whatever noise and spurious signals it is imbedded in, amplify and filter it appropriately, and digitize it without introducing artifacts.

The basics of designing circuits to do this sort of thing are covered ably and engagingly in Horowitz and Hill, and some of the fine points are covered in Chapters 15 and 18; here we look mainly at the systems engineering and theoretical aspects. Often a high level design is all you have to do, since there may be off-the-shelf components that you can just connect together to make a complete subsystem suitable for demos or for lab use.

The approach advocated here is top–down and iterative; figure out the broad requirements, budget the system specifications for noise, bandwidth, and so on, then trade off the difficulties brought to light. We'll go through some examples along the way to make the process clear. The other common technique is to do it unsystematically; put together a prototype optical system, then string a bunch of electronic modules and lock-ins together with coax cables to do the job. Both of these are valid ways of working, but they get you to different places: the systematic approach gets you a design with optimal performance that can be produced inexpensively, and the unsystematic one gets you data in a hurry from one setup.

Even an experimental apparatus that is going to be built mostly from commercial gear needs a noise budget and a high level system design, because that will help it work well and not have to be taken apart and redone. These are similar to the photon budgets we did in Chapter 1, but there are fewer first principles to work from here, and more idiosyncrasies of components to worry about. This chapter assumes familiarity with basic circuits, including resistance and reactance (and their antipodes, *conductance* and *susceptance*).

Building Electro-Optical Systems, Making It All Work, Second Edition, By Philip C. D. Hobbs

Aside: Fourier Transform Sign Conventions. Positive frequencies are easier to handle than negative ones. In the first part of the book, we were doing a lot of spatial transforms, and accordingly took as our plane wave eigenfunction $\exp(i(kz - \omega t))$, whose time dependence is a negative exponential. This convention is standard in physics books and leads to the forward Fourier transform in time having the kernel $\exp(+i2\pi ft)$, and a phase lag (or delay) being $\exp(i\phi)$. (A delay means that the measured phase corresponds to that emitted at a previous time, $t \rightarrow t - \Delta t$.) In electronic signal processing, we do mostly temporal transforms, so the convention is that a positive frequency has a positive exponential time dependence. Thus the forward Fourier transform kernel is $\exp(-j2\pi ft)$ and a delay is $\exp(-j\phi)$. Electrical engineers use j instead of i for $\sqrt{-1}$, to avoid confusion with i for small-signal current, and it's a good reminder here. Since one rarely does signal processing and wave propagation in the same calculation, the opportunity for confusion is slight. Accordingly, we bow to convention and use the $-j$ transform convention in our electronic discussions. This mild form of schizophrenia will enable you to talk to both physicists and engineers without becoming perplexed (by sign conventions, at least).

13.2 ANALOG SIGNAL PROCESSING THEORY

The foundation of signal processing is the theory of linear, time-invariant systems. This theory is so workable and intuitive that nonlinearities or time variations are usually treated as perturbations of linear time-invariant systems. This is similar to the way we treated optical calculations in Chapter 1 (the theorems and definitions are in Section 1.3.8). It's worth having a copy of Bracewell at your elbow while doing this stuff the first few times.

13.2.1 Two-Port Black Box

A sufficiently simple signal processing system can be modeled as a deterministic *two-port network*, that is, a black box with two pairs of wires coming out of it. For any input voltage $g_1(t)$ applied to one port, we get a corresponding output voltage $g_2(t)$ on the other one, which will be uniquely determined by $g_1(t)$ (apart from noise and drift, which we'll treat separately). We assume that the network response is causal, so $g_2(t)$ depends only on current and past values of g_1, so that for some operator M,

$$g_2(t) = M(g_1(t')), \qquad t' \leq t. \tag{13.1}$$

13.2.2 Linearity and Superposition

A network is said to be *linear* if it obeys two rules for all input functions g:

1. $M(\alpha g(t)) = \alpha M(g(t))$, *scale invariance*, and
2. $M(g_1(t) + g_2(t)) = M(g_1(t)) + M(g_2(t))$, *superposition*.

In algebra, these conditions define a homogeneous linear system, one with no offset term. The inevitable offsets of real networks are put in by hand afterwards, like noise and drift.

The most general causal relationship that satisfies these two conditions is the integral equation

$$g_2(t) = \int_{-\infty}^{t} m(t, t')g_1(t')U(t - t')dt', \tag{13.2}$$

where $U(t)$ is the Heaviside unit step, $U(t) = 0$ for $t < 0$ and 1 for $t > 0$.

13.2.3 Time Invariance

Most of the time we're dealing with networks whose behavior is not a strong function of time—fixed-tuned filters, amplifiers made with fixed-value components and run at constant bias, that sort of thing. For these networks, we can assume *shift invariance*: if $g_1(t)$ produces $g_2(t)$, then $g_1(t - \tau)$ will produce $g_2(t - \tau)$; that is, there's nothing special about the time origin. This is equivalent to the condition that the kernel function $m(t, t')$ depends only on $t - t'$. Replacing $m(t, t')$ with $h(t - t')$, requiring that $h(\tau) = 0$ for $\tau < 0$, we get the convolution

$$g_2(t) = \int_{-\infty}^{\infty} g_1(t')h(t - t')dt'. \tag{13.3}$$

13.2.4 Fourier Space Representation

As we saw in Section 1.3.8, a convolution in real space corresponds to a multiplication in the frequency domain. This means that the response of a linear time-invariant network can be completely described by a multiplication of the input spectrum $G_1(f)$ by a complex transfer function $H(f) = \mathcal{F}h(t)$. One important consequence of this is that if the input is an exponential function (a δ-function in **k**-space), the output is exponential as well, with the same time constant. This applies for both real and complex exponents, although of course complex ones must appear in conjugate pairs, because the voltages g_1 and g_2 are real numbers.[†]

Because g_1 and g_2 are real, $h(t)$ must also be real, since otherwise some choice of g_1 would produce a complex-valued output from a real-valued input. This means that H is *Hermitian*—Re$\{H\}$ is an even function, and Im$\{H\}$ is odd. In polar form,

$$\begin{aligned} H(f) &= |H(f)|e^{j\phi(f)}, \\ H(-f) &= |H(f)|e^{-j\phi(f)}. \end{aligned} \tag{13.4}$$

Another way to say this is that in order for $h * g$ to be real, h must apply equal gain and equal and opposite phase shifts to $\exp(j2\pi ft)$ and $\exp(-j2\pi ft)$.

If we put a function, say, $\cos(2\pi f_0 t)$, into the input, then we get

$$\begin{aligned} g_2(t) &= \frac{|H(f_0)|}{2}(e^{j\phi(f_0)}e^{j2\pi f_0 t} + e^{-j\phi(f_0)}e^{-j2\pi f_0 t}) \\ &= |H(f_0)|\cos(2\pi f_0 t + \phi(f_0)). \end{aligned} \tag{13.5}$$

[†]A circuit with a 1 ms time constant can in principle reproduce a 1 μs real exponential, but the tiny exponential response is masked by the turn-on transient in such cases.

That is, the effect of a linear, time-invariant network on a sinusoidal input is to multiply it by a constant and shift its phase.

Aside: Poles and Zeros. Except in microwave work, we build filters out of lumped elements—individual resistors, capacitors, and inductors, whose impedances go as powers of f. Combining them in series or parallel produces composite impedances and transfer functions that are rational functions of the element impedances, and a rational function of a rational function is another rational function. A rational function is completely specified by a single gain and the roots of its numerator and denominator polynomials, which are the zeros and poles of the function. Thus the behavior of all our circuits and models is determined by the poles and zeros of their transfer functions. So far, so classical. The confusing thing about it for newcomers is that the poles never lie on the real frequency axis, where all our measurements take place.† RC poles (such as an ordinary RC decoupling network) lie on the (positive) imaginary frequency axis, and LC poles lie in pairs, distributed symmetrically above the real frequency axis so that there's no phase shift at DC. Analog engineers are so used to this that they say a 10 kΩ–15 nF lowpass RC network "has a pole at 1 kHz," pointing at the −3 dB (half-power) frequency of a limp-looking smooth rolloff, and not at an infinite singularity. It's just mild sloppiness—the −3 dB point is numerically equal to the position of the real pole on the + imaginary axis, because the transfer function is $H = 1/(1 + j2\pi f RC)$—the pole is really at $f = j/(2\pi RC)$.

13.2.5 Analytic Signals

The bookkeeping necessary to carry along the explicitly real-valued nature of the signals in the network more than doubles the amount of algebra, which of course increases the number of blunders we make by at least four times. Furthermore, in more advanced network theory we need to use residue calculus a lot, which makes it even more inconvenient, since contours in the complex plane usually have to be chosen differently depending on the sign of the exponent. Accordingly, we make the following formal rules.

1. Set all the components at negative frequency to 0.
2. Double the positive ones to keep the real part the same.
3. Leave the DC alone.
4. Whenever any measurement or nonlinear operation (e.g., computing the power) is to be done, discard the imaginary part beforehand.
   4a. If you can be sure that the positive and negative frequency lobes of your signal never overlap significantly or intermodulate with each other, you can use the analytic signal in nonlinear operations (e.g., frequency modulation).

This keeps the real part $g(t)$ the same and adds an imaginary function, $-j\mathcal{H}g(t)$, producing the *analytic signal* $\hat{g}(t)$, as shown in Figure 13.1. The analytic signal is purely a calculating convenience, but as the convenience is considerable, and it is in general use, it's worth understanding and using it.

†Zeros can lie on the real axis, but usually don't.

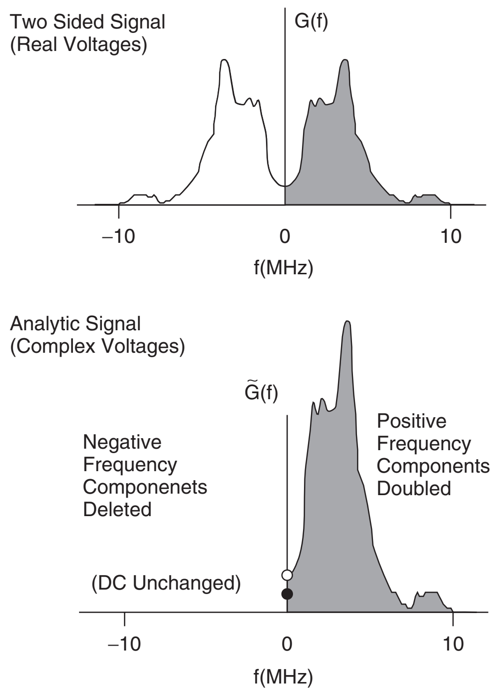

Figure 13.1. Analytic signal definition. Note that DC is a special case.

The operator $\mathcal{H}$ here is the Hilbert transform, which is an application of Cauchy's integral formula,

$$\mathcal{H}g(t) = g(t) * \frac{-1}{\pi t} = \frac{1}{\pi}\mathbf{P}\int_{-\infty}^{\infty} \frac{g(t')}{t' - t} dt' \tag{13.6}$$

where **P** denotes the Cauchy principal value. This is just an organized way of applying rules 1–4 to a given function.

The Hilbert transform phase-shifts every positive frequency component by $\pi/2$ radians, and every negative frequency component by $-\pi/2$. Multiplying by j and subtracting it from the original function cancels the components at $f < 0$ and doubles those at $f > 0$. The principal value makes sure that the component at DC stays the same, which is a good thing since you can't phase-shift DC in real life. The term *analytic signal* is slightly too restrictive, since the procedure also applies to functions with singularities, such as delta functions. Here are some examples:

True Signal	Hilbert transform	Analytic Signalignal
1	0	1
$\cos(2\pi f t)$	$-\sin(2\pi f t)$	$e^{j2\pi\mathcal{H}} \int\int\int$
$\sin(2\pi f t)$	$\cos(2\pi f t)$	$-je^{j2\pi\mathcal{H}} \int\int\int$
$\delta(t)$	$\mathbf{P}\{-1/(\pi t)^{\dagger}$	$j/(\pi t)$
$g(t)$	$\mathcal{H}g(t)$	$g(t) - j\mathcal{H}g(t)$

†Delta functions have meaning only inside integrals, and the same is true of their Hilbert transforms. The **P** indicates that this operator applies the Cauchy principal value to any integral containing it.

Like the real-space propagators in Sections 1.3.2 and 9.3.2, you rarely use the Hilbert transform in real space; it is much more common to take the transform, chop off the negative frequencies, double the positive ones, and transform back.

The Hilbert transform has practical applications in image reject mixers, which we'll see in Section 13.8.7, and theoretical ones in the synthesis of stable networks and optimal filtering, where the condition that $h(t) = 0$ for $t < 0$ forces a similar analyticity constraint on the frequency domain (see Section 13.8.4).

Analytic signals crop up all the time in signal processing literature and lead to confusing factors of 2 in bandwidths and power spectral densities; for example, the Fourier transform of a 1 second averaging window is 1 Hz wide (−0.5 Hz to +0.5 Hz) if negative frequencies are allowed, but only 0.5 Hz on an analytic signal basis. It's obvious that the Hilbert transform has the same squared modulus as the AC part of g, so the AC power of the analytic signal is twice that of g, which is a common source of blunders if rules 1–4a are not kept in mind.

Although amplitude and frequency ought to be local (i.e., instantaneous) quantities, neither is easy to define for a real-valued modulated wave; normally we have to take some average over one or more cycles. The analytic signal allows us to define them uniquely. If the analytic signal is $\hat{g} = A(t)\exp(j\phi(t))$, the instantaneous amplitude is $A(t)$, and the instantaneous frequency is

$$f(t) = \frac{1}{2\pi}\frac{d\phi}{dt}. \tag{13.7}$$

13.3 MODULATION AND DEMODULATION

A sinusoidal signal whose amplitude and phase are not intentionally made to vary with time is referred to as *continuous wave* (CW). A CW signal contains no useful information, except perhaps as a frequency or amplitude standard. Its spectrum is a single narrow spike.

The impressing of information on a CW *carrier* by changing its amplitude, phase, or both is called *modulation*. Modulation can be artificially imposed, as in voice or data communications, but in optical instruments, it is more commonly the result of the physical process we're measuring. Modulation modifies the spectrum of the carrier by reducing the strength of the central spike and putting *sidebands* on it, as shown in Figure 13.2. Sidebands appearing on the high and low frequency sides of the carrier frequency f_C are called upper and lower sidebands, respectively.† Some of the more common types of modulation appear in Figure 13.3 and their spectra in Figure 13.4.

Aside: Continuous Wave. Like many other analog signal processing terms, this one has its origins in early radio, where the first transmitters used an energetic spark in a weakly selective network to generate their output directly from a Morse key. Since a spark is a good engineering approximation to a delta function, this was an inefficient use of spectrum, to put it mildly. The first CW transmitters, where a CW oscillator was turned on and off by the key, were a revolutionary improvement, allowing lower power sets to reach much further, and more stations to fit in a smaller geographical area without interfering with each other.

†This usage is not commonly applied to FM type sidebands, which are not separately intelligible.

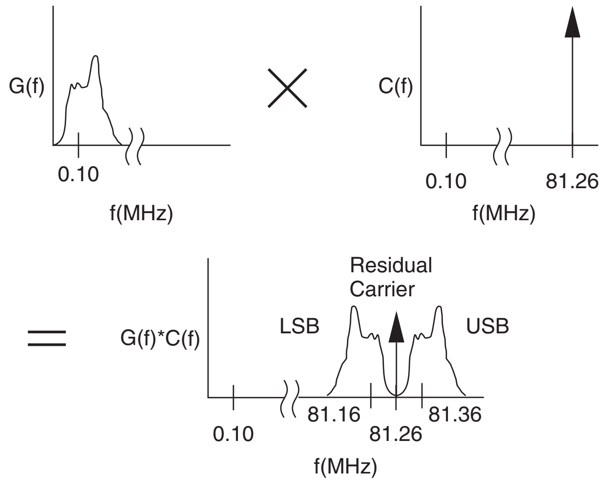

Figure 13.2. Modulating a carrier signal puts modulation sidebands on it. The structure of the sidebands, their phase relationships with each other, and how much of the carrier survives depend on the modulation scheme.

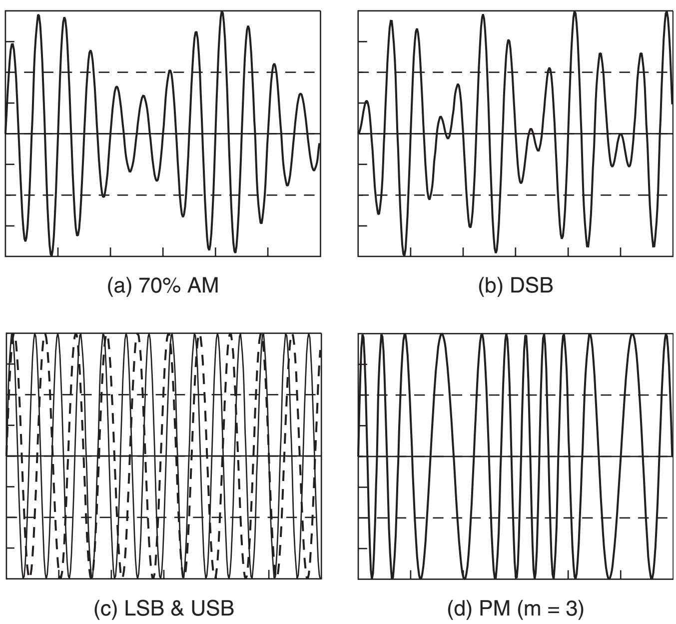

Figure 13.3. Time-domain traces of different modulation types. The modulation waveform in each case is a sine wave at 0.15 f_C. Note the phase reversals in the DSB modulated signal.

13.3.1 Terms

The discussion of modulation requires a few terms, which you must clearly understand, or the whole subject (which is a pretty one, full of good physical insight) will become opaque very rapidly (Figure 13.4). We've already encountered analytic signals, carriers, and envelopes. Here are a few more.

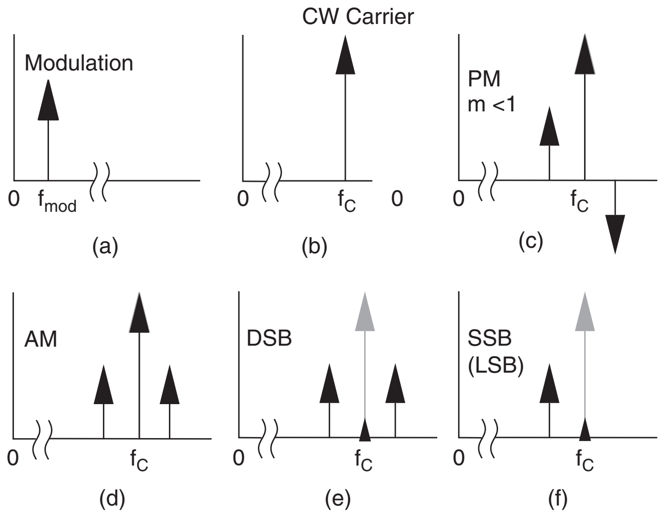

Figure 13.4. Common modulation schemes, in phasor view: (a) baseband modulating signal, (b) CW carrier, (c) narrowband PM, (d) AM, (e) double sideband (DSB), and (f) single sideband (SSB).

Baseband: Signals such as speech or video, which occupy a frequency band extending from near DC to some much higher frequency, are known as *baseband* signals. Typically what emerges from a demodulator is referred to as baseband, and nearly always the signal to a digitizer is baseband (stroboscopic digitizers such as digital receivers are exceptions). It's a fuzzy but useful concept.

Bandwidth: The output spectrum of our nice linear, time-invariant system is the product of the input spectrum with the system's transfer function. Because the signal we desire is always corrupted by additive noise before we get our hands on it, we want to choose the transfer function so as to preserve the signal while rejecting as much of the noise as we can. The wider it is, the faster it will respond but the more noise it will let through. The range of frequencies that a system will respond to is its *bandwidth*. A bandpass filter whose transfer function is down by a factor of $1/\sqrt{2}$ (0.5 in power, or 3 dB) at 99 and 101 MHz has a *3 dB bandwidth B* of 2 MHz. This is the most commonly quoted bandwidth, but there are special definitions of bandwidth for different situations. Of these, the most important is the *noise bandwidth*, which is the total noise power divided by the peak value of the output noise spectral density; that is, it's the equivalent width of the noise power spectrum. The idea is most useful when the spectrum is flat in the middle. The noise bandwidth of an RC lowpass is $1/(4RC)$, which is $\pi/2$ times wider than the 3 dB bandwidth.

Spectral Density: Radio people have very concrete minds, which helps in keeping things straight. Unfiltered shot noise and thermal noise have flat frequency spectra. Other kinds of noise, such as $1/f$ noise and filtered thermal noise, have pronounced frequency variations. What we mean by those statements is that if we connect a noise source to a filter with a constant and very narrow bandwidth B, and tune it around, the average electrical power coming out of the filter will vary with tuning or it won't, and that if B is known accurately, we can normalize it out. For a sufficiently narrow filter, doubling its bandwidth will double the power at its output. The *power spectral density* dP/df in watts per hertz is defined by

$$\frac{dP}{df} = \lim_{B\to 0} \frac{P_{\text{out}}}{B}. \tag{13.8}$$

Because of the concrete mindset alluded to above, this is almost always called the *1 Hz noise power*, because 1 Hz is usually narrow enough and the units come out right. People also talk about *voltage noise spectral density* or *current noise spectral density*, measured in volts or amps per square root hertz. Since it's power that is linear in bandwidth, these aren't really densities, but there's no getting rid of them at this point. When you have to talk about them, for example, for the transimpedance amplifier designs of Chapter 18, you can call them the 1 hertz voltage and current noise. (Just remember to scale them by $\sqrt{B}$.)

Furthermore, because the size of the quantities is frequently very small, we often quote them in a mongrel unit, dBm/Hz. The thermal noise is kT per hertz, 4.1×10^{-21} W/Hz or -173.8 dBm/Hz at 300 K. If you multiply this by the bandwidth, you'll get some pretty strange numbers, so either convert back to watts first or stay in decibels and add $\log_{10} B$ instead of multiplying.

Envelope: Modulation is expressed mathematically as multiplying a carrier by a time-varying complex envelope $\Phi(t)$, a widely useful move that we used in deriving the group velocity in Section 1.2.2, the paraxial wave equation in Section 1.3.2, and the eikonal in Section 9.2.3 A modulated signal $g(t)$ has an envelope $\Phi(t)$ defined by

$$g(t) = \Phi(t)e^{i2\pi f_C t}, \tag{13.9}$$

where f_C is the carrier frequency. The modulus of the envelope is the amplitude modulation and multiplies the carrier amplitude. The envelope phase adds to the carrier phase, and so is the phase modulation. If $f_C \gg f_{\text{mod}}$, the carrier peaks trace out a curve of $|G(t)|$ (as on a slow oscilloscope trace), but the envelope phase is invisible unless we use a phase-sensitive detector.

Modulation Frequency: Modulation frequency is different from regular frequency. Modulation frequency is the frequency of the modulation, which is related to but not identical with frequency offset from the carrier; for upper sideband (USB) modulation, $f_{\text{mod}} = f - f_C$, but in LSB, $f_{\text{mod}} = -(f - f_C)$, in AM and DSB each f_{mod} corresponds to two frequency offsets ($\pm f_{\text{mod}}$), and in FM to an infinite number, spaced at harmonics of f_{mod} from f_C. Fundamentally, the modulation frequency is the frequency that goes into the modulator or comes out of the demodulator.

Example 13.1: 5 Hz FM Can Be 1 MHz Wide. A frequency-modulated signal centered at 100 MHz, which sweeps sinusoidally back and forth by ± 0.5 MHz at a sweep frequency of 5 Hz, has a spectrum that looks like a solid blob a megahertz wide, but is really a forest of spikes at 5 Hz intervals (assuming the oscillator is quiet enough that its noise doesn't fill in the spaces completely). The modulation frequency here is 5 Hz, not 1 MHz or anything in between; the baseband signal fed to the modulator was a 5 Hz sine wave.

Sideband Folding: The distinction between frequency and modulation frequency carries over into the time domain as well. The modulating signal in Figure 13.2 looks as though it has a bandwidth of about 0.25 MHz, but the amplitude-modulated signal is twice as wide. However, the temporal response hasn't changed—putting your signal on a carrier doesn't make its envelope vary twice as fast. This is one of those cases where in principle we ought to transform back to the real-valued representation first, but it's easy to clear up the mystery. The one-sided modulation spectrum we plot is actually the analytic signal representation of the modulation, where we ignore the negative frequencies. Those negative frequencies are still there from the system's point of view, so modulation puts both the positive and negative frequency lobes on top of the carrier. The analytic signal math still applies—there is a mirror image of the two-lobed signal around $-f_C$, way off to the left of the axis. (This mild wart is worth tolerating—otherwise you'd have four products to worry about instead of two.)

AC Versus DC Measurement: It may be a bit confusing when we say that an AC measurement has twice the noise bandwidth of an equivalent DC measurement. Why does the bandwidth suddenly go from B to $2B$ when we move 1 Hz away from DC? The reason is that if we move from DC at all, we have to move further than B due to sideband folding. Consider an analog TV signal 6 MHz wide coming in on a laser beam. How would you like to demodulate it? IFs between 0 and 6 MHz are out, because the sidebands will fold over and the resulting interference will destroy information—your picture will be completely lost. DC works because the two sidebands become one sideband, and a high IF works because you can filter out the other sideband. The two schemes will differ by 3 dB in SNR, and there isn't any middle ground.

Noise Gain: The *noise gain* of a system is the partial derivative of its output with respect to a small input perturbation. It may be specified as an average over the total noise bandwidth ($q.v.$) or as a function of frequency. For an op amp, the noise gain is equal to the closed-loop noninverting gain (see Section 18.4.2).

I and Q: In Section 13.9, the envelope Φ of a modulated signal is a complex function, which can be represented in polar coordinates as amplitude and phase, or in rectangular coordinates as real and imaginary parts. In signal processing language, these are called I and Q, for *in phase* and *quadrature*. Quadrature is an old word for taking an integral, which also survives in the term *quadratic* for a second-order polynomial, and integrating a sinusoid shifts its phase by 90°. (It doesn't refer to a quarter of a cycle, although that's not a bad way of remembering it.) The most familiar sort of *I/Q* demodulator in the lab is a two-phase lock-in amplifier.

13.3.2 Phasors

Besides vocabulary, we also need to be able to visualize what's going on in a modulation process. An intuitive connection between the structure of the sidebands and what they do to the carrier is vital, and the best tool for this is the *phasor*. Phasors are vectors in the complex plane, representing the instantaneous value of each component of the analytic signal, and we visualize them spinning around the origin, as shown in Figures 13.5 and 13.6. They are often plotted in a composite 3D coordinate system, in which the

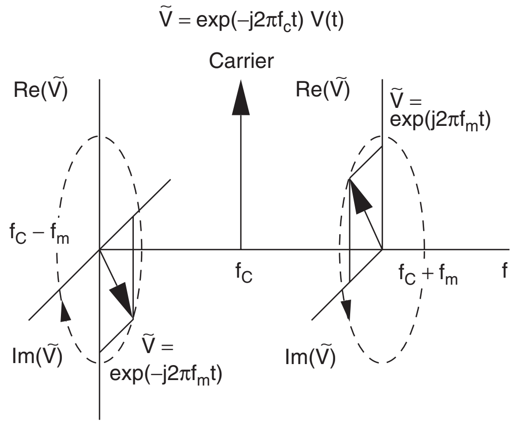

Figure 13.5. Rotating phasor plot of a narrowband PM signal, plotted in coordinates that corotate with the carrier. The oppositely phased sidebands make the amplitude stay constant when the phasors are added vectorially.

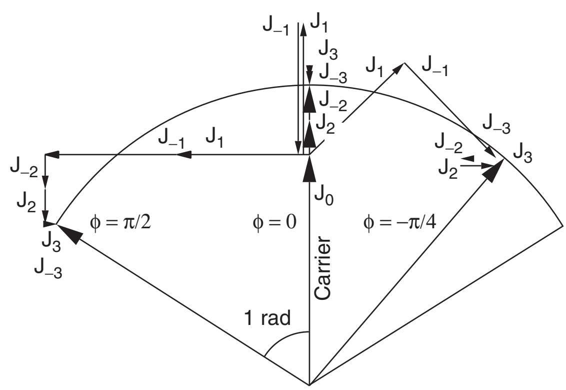

Figure 13.6. Two-dimensional phasor construction for narrowband angle modulation (FM and PM) with a modulation index $m = 1$ radian, showing how the sideband phasors add up to a constant radius.

third dimension is frequency, a sort of complex-valued version of a spectrum analyzer display. Since each frequency component spins around at a different speed, visualization is possible only with discrete frequency components (i.e., one or a few sinusoids), but this is enough. A second simplification is to transform to corotating coordinates, so that the carrier phasor is constant and the upper and lower sidebands rotate in opposite directions. Vector addition of the phasors then allows us to see easily that the in-phase sidebands of AM cause the instantaneous amplitude (length) to change but keeps the phase (direction) constant, and the 180° out of phase sidebands of FM always sum to zero length change but cause the carrier phase to wiggle back and forth. (See Figures 13.4 and 13.5)

13.3.3 Frequency Mixing

One of the most useful operations in signal processing is *frequency mixing*. The term *mixing* is familiar from audio, where it describes adding different sound sources together before feeding them into a recorder or amplifier. In signal processing, mixing refers to multiplying signals together rather than adding them, and the two usages must be kept quite distinct. The typical mixing operation is to multiply a signal at the RF port by a sinusoidal *local oscillator* signal at the LO port:

$$V_{\mathrm{IF}} = kV_{\mathrm{RF}} \cos(2\pi f_{\mathrm{LO}} t), \tag{13.10}$$

so a component of V_{RF} at f produces

$$\begin{aligned} V_{\mathrm{IF}} &= k \cos(2\pi f_{\mathrm{LO}} t) \cos(2\pi f t + \phi) \\ &= k[\cos(2\pi (f_{\mathrm{LO}} + f)t + \phi) - \cos(2\pi (f_{\mathrm{LO}} - f)t - \phi)]. \end{aligned} \tag{13.11}$$

The result is that both the RF and LO frequency components disappear, being replaced by new components at $|f_{\mathrm{LO}} \pm f|$. There are a couple of things to keep in mind here. One is that the frequencies and phases add and subtract, they don't multiply or divide, and that it's phase and not time delay that matters. If you mix signals at 10 GHz $\pm$ 1 kHz, you'll get outputs at 20 GHz and at 2 kHz. If you delay one of the 10 GHz signals by 10 picoseconds, which is 1/10 cycle, you will phase shift both the 20 GHz and the 2 kHz signals by 1/10 cycle—equivalent to a time delay of 5 ps for the one and 50 μs for the other. (Note that no causality violation is involved here—we'll talk about what happens with more complicated signals later.) All the derivatives of the phase also add and subtract, of course, so that modulation frequency, modulation index, and frequency offsets stay the same—time-domain multiplication by an exponential is a frequency-domain convolution with a shifted δ-function. The second thing is that in an ideal linear multiplier, the amplitudes of the sum and difference go as the product of the amplitudes of the LO and RF signals, but most practical mixers are run with the LO saturated, so that LO amplitude fluctuations are not impressed on the mixing products.

13.3.4 Amplitude Modulation (AM)

The simplest form of modulation is to change the amplitude of the carrier, by multiplying it by a real, nonnegative envelope function. This is naturally called *amplitude modulation* (AM). Achieving pure AM is not always trivial, since many processes producing it will modulate the phase as well (e.g., changing the gain of an amplifier usually makes it faster or slower). The *Modulation depth* is the ratio of the peak–valley envelope excursion to the peak height (the same definition as fringe visibility). A sinusoidal carrier modulated to a depth d by a sinusoidal signal at frequency f_{mod} is

$$\begin{aligned} v(t) &= \frac{1 + d \cos 2\pi f_{\mathrm{mod}} t}{1 + d} \cos 2\pi f_C t \\ &= \frac{1}{1+d} (\underset{\text{Carrier}}{\cos 2\pi f_C t} + \underset{\text{USB}}{\frac{d}{2} \cos 2\pi (f_C + f_{\mathrm{mod}})t} + \underset{\text{LSB}}{\frac{d}{2} \cos 2\pi (f_C - f_{\mathrm{mod}})t}) \,. \end{aligned} \tag{13.12}$$

The carrier has grown sidebands of equal strength, spaced above and below the carrier frequency by f_{mod}. A characteristic of AM is that the power in each sideband cannot be more than a quarter of the carrier power, which doesn't help the SNR one bit.

13.3.5 Double Sideband (DSB)

If our modulator can merely turn the amplitude of the sine wave up and down, we're limited to the AM case. If we have a more capable modulator, such as a double balanced mixer (see below), our modulation envelope can go negative as well, becoming

$$\begin{aligned} v(t) &= (\cos 2\pi f_{\text{mod}} t)\cos 2\pi f_C t \\ &= \tfrac{1}{2}(\cos 2\pi(f_C + f_{\text{mod}})t + \cos 2\pi(f_C - f_{\text{mod}})t). \end{aligned} \tag{13.13}$$

The carrier goes away, and only the AM-type sidebands are left. This sort of modulation is pedantically called *double sideband, suppressed carrier*.

13.3.6 Single Sideband (SSB)

If we use a sharp filter to chop off one of the sidebands in a DSB signal, we get single sideband modulation. SSB is just the shifting of the modulation by the carrier frequency, so that a modulation component at f_{mod} gets moved to $|f_C \pm f_{\text{mod}}|$. Choosing the plus sign leads to *upper sideband* (USB) modulation, the minus sign to *lower sideband* (LSB).† The two are not equivalent, since an LSB modulated signal has had its low modulation frequencies flipped around to higher signal frequencies. A common error in analog signal processing is to use an odd number of LSB conversions, so this flipping is not corrected. The symmetry of the sidebands in AM, DSB, and FM make this much less of a concern there.

13.3.7 Phase Modulation (PM)

Phase modulation is just what it says: the modulation signal goes in the argument rather than as a real number multiplier outside. From an analytic signal point of view, sinusoidally phase modulated waveform is multiplied by $\exp[jm\cos(2\pi f_{\text{mod}} t)]$ for some value of m—the envelope magnitude remains 1 but the phase varies. This can be done, for example, with a Pockels cell whose optic axis is aligned with the **E** field of our beam (so there are no polarization funnies), or for RF signals with a transmission line built on a ferroelectric ceramic such as barium titanate; it can also be done with variable all-pass networks, carefully linearized (see Section15.12.2). A PM signal has a more complicated spectrum than an AM one, since it has an infinite number of harmonics:

$$\begin{aligned} v(t) &= e^{j(2\pi f_C t + m\cos 2\pi f_m t)} \\ &= \sum_{n=-\infty}^{\infty} j^n J_n(m) e^{j2\pi(f_C + n f_m)t}, \end{aligned} \tag{13.14}$$

where J_n is as usual the nth order Bessel function of the first kind, and the constant m is the *modulation index*, which is the peak phase deviation in radians. The Bessel function $J_n(m)$ has significant amplitude only for $n \lesssim m$, falling off exponentially with n for $n > m$. Negative order Bessel functions obey

$$J_{-n}(x) = (-1)^n J_n(x). \tag{13.15}$$

†This assumes that $f_{\text{mod}} < f_C$. (Why?)

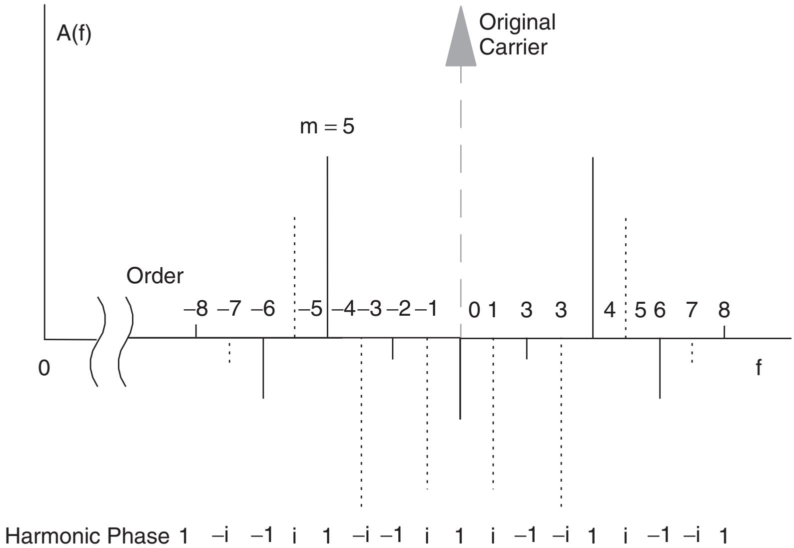

Figure 13.7. Detail of a wideband PM signal with $m = 5.0$, showing the phase relationships of the sidebands. The lines shown represent phasors, with the imaginary ones shown as dotted lines.

The instantaneous frequency is

$$\begin{aligned} f(t) &= \frac{1}{2\pi}\frac{d}{dt}(2\pi f_C t + m\cos 2\pi f_m t) \\ &= f_C + m f_m \sin 2\pi f_m t, \end{aligned} \tag{13.16}$$

so (as expected) the frequency swings between the limits $f_C \pm m f_m$, which are roughly the points where the Bessel functions cut off. Figure 13.7 is a phasor construction showing how the FM sidebands add up to produce phase modulation with no amplitude variation. Figure 13.7 shows the complex amplitudes of the carrier and modulation sidebands for a phase modulated signal with $m = 5.0$. Note that the carrier's phase is inverted,† and that the phases of the sidebands are not trivial to predict without using (13.14).

There are three relevant regimes of m: $m \ll 1$, narrowband PM/FM, where only the ± 1 order sidebands are significant; $1 \lesssim m \ll f_C/f_m$, quasistatic wideband PM/FM, where many sidebands contribute; and $m \gtrsim f_C/f_m$, the nonquasistatic case. The first two are the usual situations, where the variation of the envelope phase is slow compared with the carrier period, so that the modulation sidebands die out well before they get to DC, making the analytic signal expression (13.14) valid. The nonquasistatic case, where the sidebands are significant in amplitude down to DC, is a complicated mess—for example, modulating a 100 MHz carrier at 1 GHz f_m, or deviating it ± 6 radians ($m = 6$) at a 20 MHz f_m. For modulation like this, the analytic signal representation has to be discarded, and the phase calculated explicitly as the time integral of the frequency.

†The carrier goes to 0 at the roots of J_0, which is how you calibrate phase and frequency modulators. The first Bessel null is at $m \approx 2.405$.

Messiness is typical of situations where positive and negative frequency lobes start to overlap. Nonquasistatic PM/FM is one such situation; others are aliasing in gratings (Section 7.3.1) and in sampled-data systems (Section 17.4.4).

13.3.8 Frequency Modulation (FM)

Frequency is the time derivative of phase in cycles, so frequency modulation (FM) is a close relative of PM. The two are often called *angle modulation* for this reason. A given modulation voltage moves the carrier by a specific phase offset in PM and a specific frequency shift in FM. In principle, the modulation signal could be integrated and applied to a phase modulator to produce FM. This is not a practical approach, however; a DC term in the modulation waveform causes the average frequency of the signal to change, so that the phase shift grows secularly with time, eventually exceeding the range of any phase shifter. Because of this, FM is in fact generated by putting the modulation on the control voltage pin of a voltage controlled oscillator (VCO) and is demodulated differently as well. The integrating action of the VCO reduces the modulation index of high frequency components by a factor of f_m, so that the total signal bandwidth (approximately mf_m) is roughly independent of f_m, a desirable characteristic when f_m varies widely, as in music broadcasting. FM demodulators are relatively insensitive to oscillator phase noise, since the same integrating action boosts signal power at low modulation frequencies by a factor of f_m^{-2}, enough to overwhelm the low frequency noise (see Section 13.6.10).

Aside: FM Preemphasis. In order to maximize the received SNR for a given signal bandwidth, the spectrum of the baseband signal should be white, that is, flat up to its maximum frequency. The spectrum of music is concentrated at low to mid audio frequencies, so FM broadcasting uses *preemphasis*: an RC lead–lag network boosts the high audio frequencies before modulation, and a matching *deemphasis* network after the receiver's demodulator attenuates them again, *along with the high frequency noise*. This is an example of *whitening*, a widely applicable method of improving SNR when the noise is white and the signal isn't—see Section 13.8.10.

13.4 AMPLIFIERS

Amplifiers, as we all know, are devices for making signals bigger. From a signal processing point of view, we want them to do that while remaining nice linear, time-invariant networks, and perhaps having other attributes. Design details are found in Section 14.6.7 and most of Chapter 18. Here we'll stick with their behavior.

13.5 DEPARTURES FROM LINEARITY

Most of the time, the nonlinearities of our networks will be small, so that they can be modeled as perturbations to the behavior of an ideal linear network. The amplitude response of the network can then be written as a perturbation power series in V (valid near $V = 0$),

$$V_{\text{out}} \approx \sum_{n=0}^{\infty} d_n V_{\text{in}}^n, \tag{13.17}$$

which is called the *distortion polynomial*, probably because nobody ever uses more than five terms. It must be underlined that for a real network, the d_n are not fixed in value but depend on the signal frequencies and other things, and that for sufficiently large V, the network will become nonlinear and time varying enough that this series is invalid. The value of the distortion polynomial is that it predicts the amplitude dependence of different distortion products. These products are called *spurs*, a fortunate abbreviation of "spurious" that also connotes small, sharp protuberances, which is what they look like on a spectrum analyzer. Spurs due to distorted modulation (e.g., driving an amplitude modulator too hard) are collectively known as *splatter*, which is what they look like on the analyzer when you shout into the transmitter microphone.

The best thing about RF signal processing is that spurious mixing products land at frequencies away from the desired output, so that they can be removed almost completely by filtering. This luxury does not exist at baseband, for example, in the output of a phase-sensitive detector. (Compression and cross-modulation can still spoil the party—read on.)

13.5.1 Harmonics

The best known consequence of nonlinearity in a network is harmonic generation. If the input is $A\exp(j2\pi ft)$, then the output will be

$$V_{\text{out}} \approx \sum_{n=0}^{\infty} d_n A^n e^{j2n\pi ft}, \tag{13.18}$$

so that the nth harmonic amplitude will go as A^n. Accordingly, a 3 dB change in A will produce a 6 dB change in the second harmonic, 9 dB in the third, and so on. Do remember that this only works in the weak distortion limit. This is because we have broken our rule about going back to the purely real representation before applying nonlinear operations. A trigonometric function raised to the nth power contains components at nf, $(n-2)f$, $(n-4)r$, ... down to f or 0, depending on whether n is odd or even. Equivalently, the component at nf will depend on A^n, A^{n+2}, A^{n+4}, and so on. Still, since the d_n are not rapidly increasing, the small-signal leading behavior of the nth harmonic will be controlled by the term in A^n.

This is made even more complicated by biasing effects, which make the response asymmetric between positive and negative going peaks. To reflect this, there should notionally be a DC term added to the input signal before all the nonlinear operations, which will mean that the term in A^n will really contain contributions from order n and all higher orders, not just $n+2$, $n+4$, ... This is part of the reason that the distortion polynomial is mainly a conceptual aid.

How prone a device is to produce harmonics is usually specified by its nth order *intercept points*, which are where the extrapolated power law curve for the fundamental ($n = 1$) and nth harmonic cross, as shown in Figure 13.8 (this isn't the most commonly quoted intercept point—see the next section).

13.5.2 Frequency Multipliers

Sometimes harmonics are useful, for example, to produce a reasonably stable high frequency signal from a very stable low frequency one. Devices that do this on purpose

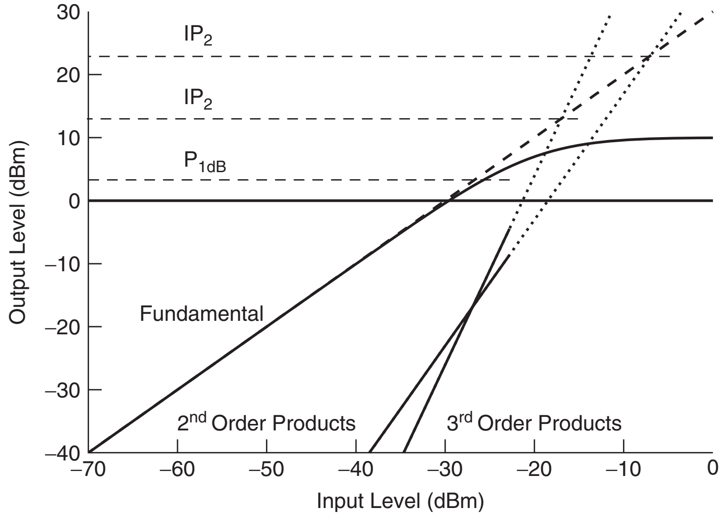

Figure 13.8. Polynomial behavior of distortion products of a weakly nonlinear network. IP_2 and IP_3 are the second- and third-order output intercept points.

are called *frequency multipliers*. Any modulation on the signal being multiplied will get changed. The multiplier is strongly nonlinear, so it isn't easy to say in general what happens to amplitude modulation, except that it's nothing pretty. The behavior of any PM is much easier to predict: since the frequency is multiplied, the phase is too, and so a frequency n-tupler multiplies the modulation index by n but *leaves the modulation frequency unchanged*. For weak phase modulation (e.g., oscillator phase noise), the sideband power goes as n^2. Frequency multipliers may exhibit AM–PM conversion (see Section 13.5.6), in which case the phase noise may be worse than this. A phase-locked loop (PLL) with a frequency divider between the VCO and phase detector is another sort of frequency multiplier. In frequency dividers, amplitude information is lost but the modulation index goes down by a factor of N. A divider introduces an N-fold phase ambiguity in its output since we can never be sure which of N states it wakes up in. This causes no problems in straight PLL multipliers (why?) but it does in M/N synthesizers, where both the reference and VCO frequencies are divided before the PD. Thus *direct synthesizers*, using multipliers and mixers, is preferred to PLL synthesizers whenever the absolute phase is important (see Section 13.9.4 for more).

13.5.3 Intermodulation

If V_{in} is not a pure tone but contains several frequency components, the distortion will include mixing products between all combinations of these components, collectively known as *intermodulation distortion* (IMD). Second-order intermodulation is not usually much of a problem, since these products are either near DC or near twice the signal frequency. As with harmonics, unless we're working at baseband, these are easily filtered out. On the other hand, third-order IM products at $2f_1 - f_2$ tend to land in your passband, so they're much more obnoxious.

Intermodulation performance is specified with intercept points, as harmonics are, but since there are more degrees of freedom here, the test conditions have to be carefully specified. The two-tone test is the most common. Two closely spaced sinusoidal signals of equal amplitude are fed in to the input. The strengths of the nth order mixing products go as V^n, as before. The third-order intermodulation intercept point is where the extrapolated $2f_1 - f_2$ amplitude hits that of one of the desired signals (this is what is most commonly meant by IP_3). When there is a forest of these spurs, the power law dependence is very helpful in identifying the intermodulation order of each one. Exact identification, at least of the lower lying orders, is possible by varying the tone amplitudes separately over a narrow range.

Aside: Optical IMD. You even see this with lasers. If you shine a multiple longitudinal mode HeNe laser on a photodiode and look at the resulting spectrum, you will usually find sharp spurs that sweep aimlessly back and forth from near DC to a megahertz or so. These are fourth-order mixing products, caused by the interference of one mode with a third-order IM product of two others, all four signals being near 632.8 nm (474 THz). The frequency offset and instability result from slight dispersion in the plasma and the mirrors. (Optics people call this *four-wave mixing*.)

13.5.4 Saturation

As the drive level is increased, eventually the ability of the device to drive its output will be exceeded. In high frequency devices, this limit is quite gradual; they don't crash into the supply rail and stop the way op amps do. This behavior is usually summarized by the 1 dB compression point $P_{1\text{ dB}}$, which is the output power at which the fundamental gain G has dropped by 1 dB. This is typically about 10 dB below IP_3. The compression behavior is governed by the next term in the distortion polynomial that gives rise to a signal at the fundamental, which (remembering the DC bias) is the A^2 term. At a level 10 dB below the 1 dB compression point, the compression is typically 0.1 dB, as we'd expect from a quadratic term. Below $P_{1\text{ dB}}$, where no serious waveform distortion occurs, the compression error δ is

$$\delta \equiv 1 - \frac{P_{\text{out}}}{GP_{\text{in}}} \approx \frac{P_{\text{out}} \ln 10}{10 P_{1\text{ dB}}}, \tag{13.19}$$

and the compression factors $(1 - \delta)$ due to cascaded devices multiply. To the same order, the deltas add, so we can define the *compressibility coefficient* CC of a given element in a system as

$$\text{CC} = \frac{G_{\text{tot}}}{P_{1\text{ dB}|\text{out}}(\text{mW})}, \tag{13.20}$$

where G_{tot} is the total small-signal power gain (in W/W) from the input of the system up to the output of that element. The overall compression error is then

$$\delta_{\text{tot}} \approx 0.23 P_{\text{in}} \sum_i \text{CC}_i \tag{13.21}$$

and the input 1 dB compression point of the system is

$$P_{1\text{ dB}|\text{in}} = \frac{1}{\sum_i \text{CC}_i}. \tag{13.22}$$

(For devices such as mixers, where the compression point is specified at the input, use the total gain up to the input instead.) If the input level is turned up high enough, the output level will stop increasing altogether, and the amplifier will become a limiter (see Section 14.7.16). Most don't limit well unless they're specifically designed to.

13.5.5 Cross-Modulation

As with all mixing products, the phases of the spurious signals are additive (since the frequencies add, the phases must too), and the amplitudes are multiplicative. This means that if the input consists of different frequency components, the modulation of the stronger ones will be transferred to the other components.

This most obnoxious behavior is called *cross-modulation*. An amplitude modulated signal mixing with itself produces a baseband product that approximately reproduces the modulation waveform.[†] If this product in turn mixes with a component at f_2, the result will be an amplitude modulated spurious product with a carrier at f_2, right on top of the desired component. Thus the desired component will acquire the AM of the interfering signal. This is physically obvious, since if the peaks of one signal are strong enough to reduce the gain of the device, the resulting compression will affect other signals in the input as well.

13.5.6 AM–PM Conversion

One kind of distortion that is not an obvious consequence of the distortion polynomial is AM–PM conversion. Saturating a network changes the time-averaged device parameters (e.g., the phase of the current gain), so that the time delay through the network changes too. Amplitude modulation of the signal will cause the delay to vary with the modulation waveform, thus causing phase modulation. This effect is particularly strong in limiting amplifiers, used for FM and PM IF strips, where it is frequently the largest source of error. Since a limiter is intended to be operated in a strongly nonlinear region, the obvious solution of reducing the signal level to reduce the AM–PM conversion is not applicable. You just have to pick amplifier designs that don't slow down so much when they go into limiting, such as FET differential amplifiers with low impedance drain loads, and reduce the frequency at which the limiting occurs to as low a value as practicable, so that a delay τ produces a smaller phase shift. This is one instance in which a double-conversion system makes good sense.

13.5.7 Distortion in Angle Modulated Systems

Because the information in angle modulated systems is encoded in the phase, it is highly insensitive to amplitude distortion, allowing the use of amplitude limiting in the receiver to suppress AM noise, interference, and signal fading. This doesn't qualify as a free lunch; the information is encoded in phase changes, so phase nonlinearity causes the same sorts of trouble as amplitude nonlinearity in AM systems. (This is obvious from a time-domain point of view.) In the frequency domain, phase nonlinearity causes the carrier and sideband phasors not to sum to a constant length. After limiting, the resulting AM information is lost, so the undistorted phase information cannot be recovered by any

[†]It isn't perfect—signal $\times$ carrier reproduces the modulation, but signal $\times$ signal is a quadratic nonlinearity.

equalizing filter. Even a linear-phase filter can cause distortion by rejecting sidebands. FM sidebands depend nonlinearly on the modulating signal, so each one by itself contains intermodulation products between different baseband frequencies that only cancel out when all sidebands contribute to the detected signal.

Aside: Signal Detection in Noise.

If your experiment needs statistics, you ought to have done a better experiment.

—Ernest, Lord Rutherford

Rutherford was a great enthusiast for clear, cheap, clever experiments. Following this piece of advice would render many sorts of experiments impossible even in principle, but it holds an important point: don't passively accept poor quality data. Spend lots of effort on improving the quality of the data, with the aim of achieving a beautiful, clean result. When that is done, *then* worry about statistical analysis. A properly designed statistical signal processing system can often yield an extra few decibels in sensitivity, which is well worth having, but only after the last decibel has been wrung out of the front end. In any case, don't muck about with parameters whose physical significance you can't explain in less than a paragraph.

13.6 NOISE AND INTERFERENCE

We live in a fallen world, so the signals we process are never free of noise, distortion, and extraneous interfering signals. We must be forearmed against them by knowing what they do to our measurements, and how.

13.6.1 White Noise and $1/f$ Noise

Noise whose power spectral density (PSD) is constant with frequency is said to be *white*, by analogy with light. If you look at the baseband output of some system on a spectrum analyzer when no input signal is present, most of the time you'll see the noise rise to a big peak at DC. Above there, usually, there's a white noise region known as the *flatband*, followed by a rolloff toward high frequency. There may also be a noise peak between the flatband and the rolloff, as we'll see in Section 18.4.2.

The low frequency peak is partly due to the finite bandwidth of the spectrum analyzer smearing out the DC offset, but it's partly real. A system with any DC offset whatever notionally has a component of $\delta(f)$ in its spectrum, and often exhibits a very steep rise in noise near 0. The two major components of this near the low edge of the flatband are drift, which has a characteristic $1/f^2$ dependence in its power spectrum, and a motley collection of pops, clicks, and bangs called $1/f$ noise from its frequency dependence. (Since the real frequency dependence is often far from $1/f$, it's also known as *flicker noise* or, less descriptively, *excess noise*.) Both are statistically nonstationary and generally very badly behaved. Dialing up the time constant on your lock-in amplifier and taking longer to do the measurement will not help these sorts of noise, because for a fixed number of data points, the low frequency cutoff scales with the high frequency one; in that case the variance of $1/f$ noise is independent of the upper frequency limit, and for anything steeper, the variance gets *worse* as the bandwidth goes down. (This is obvious in the case of drift—averaging a linear ramp for a longer time just takes you further from home.) Many a graduate student has come to grief by not knowing this.

Flicker noise comes from sources like contamination migrating around on the surfaces of ICs, conductance and carrier density fluctuations in carbon resistors, electromigration of metal atoms in conductors or interstitial atoms in silicon, and a million more. Because these are conductance variations, $1/f$ noise only shows up when a current is flowing through the noisy part, and the noise power density is proportional to i^2, just as signal power is, rather than i as with shot noise. Migrating contamination, often due to a cracked IC package or a manufacturing defect, shows up as small jumps in offset voltage, known as *popcorn noise* or *telegraph noise*. Flicker noise is a bear to deal with, since the system will sometimes just sit there being well behaved, only to get noisy again minutes or hours later. It is a strong sign of an unreliable component.†

13.6.2 Johnson (Thermal) Noise

A beautiful result of classical thermodynamics is the *fluctuation–dissipation theorem*. It states that any process that can dissipate energy exhibits thermal fluctuations, and tells how to calculate them. What that means to us is that circuit elements that can dissipate electrical power as heat (e.g., resistors and lossy capacitors) are also noise sources. A resistor R at temperature T, connected to a matched load (i.e., another resistor whose value is also R) will transfer noise whose average power in a 1 Hz bandwidth is

$$p_N = kT. \tag{13.23}$$

As usual, we can calculate the total noise power P_N by integrating the noise power per hertz p_N times the squared modulus of the transfer function of the network. For flat-topped frequency responses, this basically amounts to multiplying p_N by the noise bandwidth in hertz.‡

Since the resistor is a linear device, the noise can be modeled with a Thévenin or Norton model,§ that is, as a noise voltage source v_N in series with a noiseless R, or as a noise current source i_N in parallel with a noiseless R, as shown in Figure 13.9. If you open-circuit a Thévenin model, the rms open circuit noise voltage v_N is twice that measured under matched conditions:

$$e_N = \sqrt{4kTR}. \tag{13.24}$$

Similarly, if you short-circuit a Norton model, the short-circuit current is twice the matched value:

$$i_N = \sqrt{\frac{4kT}{R}}. \tag{13.25}$$

†If you find a part that produces unusually bad flicker noise, *throw it away*, and seriously consider pounding it into dust with a ball-peen hammer first, so it can never fool you again. (If you find a whole lot of parts like that, save them to pound the manufacturer with.)

‡This is in the analytic signal picture: classical equipartition predicts that in the two-sided spectrum, the uncertainty in a 1 second measurement in a 1 Hz bandwidth is $kT/2$, just as in every other classical degree of freedom—the factor of 2 comes from the analytic signal conversion.

§Horowitz and Hill have more on Thévenin and Norton models. If you're reading this chapter because you need to design or understand the guts of how low noise circuits work, and you don't own a copy of *The Art of Electronics*, then do not walk but run to get a copy.

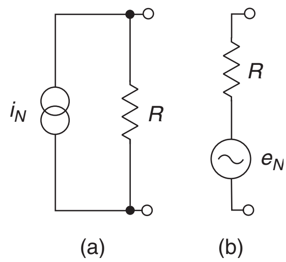

Figure 13.9. Noise models of a resistor: (a) Norton and (b) Thévenin. The two are exactly equivalent, so pick whichever one makes the algebra easier.

In a reactive network, the reactances are noiseless, so the open-circuit noise is that of the resistive part of the total impedance,

$$V_N^2 = 4kT \int_0^\infty R \, df. \tag{13.26}$$

Equivalently, shorting out a parallel network makes the reactive currents zero, so the short-circuit noise current is that of the conductance,

$$I_N^2 = 4kT \int_0^\infty G(f) \, df. \tag{13.27}$$

Example 13.2: Noise of Resistors in Series and Parallel. Two resistors R_1 and R_2 wired in series behave electrically as a single resistor of value $R_1 + R_2$. We expect their combined noise properties to be the same as well. Choosing the Thévenin equivalent noise circuits, the series resistance is indeed $R_1 + R_2$, and the two noise voltages add in rms fashion (power); thus the total noise is

$$\begin{aligned} v_{N\text{tot}} &= \sqrt{v_{N1}^2 + v_{N2}^2} \\ &= \sqrt{4kT(R_1 + R_2)}. \end{aligned} \tag{13.28}$$

Similarly, using the Norton model for resistances in parallel, the noise currents add in rms fashion, so that the total is

$$\begin{aligned} i_{N\text{tot}} &= \sqrt{i_{N1}^2 + i_{N2}^2} \\ &= \sqrt{4kT(1/R_1 + 1/R_2)}, \end{aligned} \tag{13.29}$$

both of which are just what we expect. A note to those fond of negative impedance converters: negative resistances don't have imaginary noise.

Example 13.3: kTC Noise. This is an alternative derivation of Johnson noise and is helpful in nonlinear situations where R may be changing in unknown ways with time (e.g., due to switching). In a parallel RC network, the electrical connection couples the capacitor's degree of freedom (i.e., its stored energy) to the temperature reservoir, so we know from classical equipartition that the ensemble average energy on the capacitor is $\frac{1}{2}kT$. Since the average stored energy $W = \frac{1}{2}CV_N^2$, the rms noise voltage V_N must be $(kT/C)^{1/2}$, or equivalently the rms charge noise $\Delta Q_N = (kTC)^{1/2}$. This drops nicely out of (13.24), since the resistance of a parallel RC is $R/[1 + (\omega RC)^2]$:

$$\langle V_N^2 \rangle = 4kT \int_0^\infty \frac{R}{1 + \omega^2 R^2 C^2} \frac{d\omega}{2\pi} = \frac{2kT}{\pi C} \int_0^\infty \frac{du}{1 + u^2} = \frac{kT}{C}. \tag{13.30}$$

Since R and C can be anything, and we can take sums and differences of different bandwidths, this applies in any bandpass; thus the noise must be flat. Since an RC lowpass has a noise bandwidth of $1/4RC$, the spectral density contributed by the resistor is

$$v_N = \sqrt{(kT/c)(4RC)} = \sqrt{4kTR}. \tag{13.31}$$

This uncertainty applies whenever we're resetting the charge on a capacitor with a feedback-free network such as a switch or CMOS gate; since amplifiers need not have noise temperatures near their physical temperature, a feedback network can do better in general. This *kTC* noise dominates the readout noise of CCDs, because the charge on their readout stages is dumped on each cycle, a problem solved by *correlated double sampling* (see Section 3.9.4).

13.6.3 Shot Noise in Circuits

Most resistors don't exhibit shot noise; metal film resistors especially are known to be very quiet (see Section 3.10.2). Currents in which carriers arrive more or less independently (e.g., via tunneling or photoionization) have full shot noise. These include base, emitter, and collector currents of BJTs, where the conduction is by minority carriers. FETs generally do not have full shot noise, but their other noise sources generally dominate, making FETs noisier than BJTs except in high impedance circuits.

13.6.4 Other Circuit Noise

There are lots of other sources of junk in signals, for example, $1/f$ noise due to conduction fluctuations in carbon and thick-film resistors; microphonics due to high-k ceramic capacitors becoming piezoelectric; jiggling cables with DC bias on them; capacitive and inductive pickup; popcorn noise from electrolytic capacitors; and wideband noise and spikes on power supplies and voltage references. Those ones you just have to foresee and design out. When you're learning, you usually find them during debug and remember them for next time.

13.6.5 Noise Figure, Noise Temperature, and All That

At low frequencies, it is not difficult to measure voltage and current separately. At RF, on the other hand, everything becomes inductive or capacitive. High impedances are difficult to maintain; a good rule is that in a lumped-constant (i.e., not transmission line) circuit, life begins to get difficult at impedance levels of about 100 kΩ/f(MHz) (corresponding to a node capacitance of 1.6 pF). The outputs of things talk to their inputs, and so forth. It's just a more complicated regime.

Reactive mismatches make the noise behavior of high frequency circuits more complicated. The noise of an RF amplifier is usually quoted as a *noise figure*, which is simply how much more noise (in dB) this amplifier puts out versus a noiseless amplifier, if both of their inputs are terminated with a 50 Ω resistor at $T = 300$ K. (It's made more confusing by the fact that although the maximum power is coupled to the amplifier when the source is impedance matched, this does not result in the best SNR because the amplifier's noise usually is not a simple Thévenin model characterized by a single temperature.)

Consider an amplifier with power gain G. If we notionally hook up a noiseless 50 ohm resistor to its input, it will produce a 1 Hz output noise power $p_{N\text{out}}$. This allows us to define its intrinsic 1 Hz input noise power with a matched load, $p_{N\text{in}}$, as

$$p_{N\text{in}} = \frac{p_{N\text{out}}}{G}. \tag{13.32}$$

A resistor at temperature T has Johnson noise power of $p_{N\text{th}} = kT$ in 1 Hz, which allows us to define the noise figure NF:

$$\text{NF(dB)} = 10\log\left(1 + \frac{p_{N\text{in}}}{kT}\right)\bigg|_{T=300\text{K}}. \tag{13.33}$$

We can't actually get a noiseless resistor, but we can do pretty well by doing measurements with the input resistor at different temperatures and using curve fitting to get G and $p_{N\text{in}}$. Other techniques using calibrated noise sources are available, so it isn't always necessary to use liquid nitrogen. See Section 2.5.4 for more on practical noise measurements.

The noise figure of a generic amplifier is around 3–4 dB, and a really good one is below 1 dB.[†] Such low noise figures are usually only needed when the source is in fact not a room temperature resistor. Low temperature sources include cryogenically cooled IR detectors, and antennas pointing away from the Earth and Sun. Photodiode networks are usually badly mismatched (otherwise we'd be wasting half our signal power inside the network), so they don't look like a room temperature resistor either.

For these low noise, oddball applications, the noise figure isn't too meaningful because it's always near 0 dB. Another more useful parameter is the noise temperature,

$$T_N = \frac{P_{N\text{in}}}{k}, \tag{13.34}$$

which is how far you'd have to cool your resistor before its noise equaled the amplifier's. Good low noise amplifiers have noise temperatures of 100 K or below, even though their physical temperature is 300 K. (In case you're worried about the fluctuation–dissipation theorem here, don't be—these devices are that quiet only when they're turned on; the fluctuation–dissipation theorem's assumption is a closed, isothermal system, which breaks down as soon as power is applied.)

Table 13.1 is a handy conversion chart of noise figure to noise temperature.

[†]Even vacuum tube amplifiers got below 2 dB, which is pretty good for something running at 1000 K. Check out the Miteq catalog for amps with guaranteed NFs of 0.35 dB ($T_N \approx 25$ K).

TABLE 13.1. Conversion from Noise Temperature to Noise Figure

T_N (K)	NF(dB)	T_N (K)	NF(dB)	T_N (K)	NF(dB)
650	5.0	100	1.25	25	0.35
500	4.3	78	1.00	15	0.21
400	3.7	75	0.96	10	0.142
300	3.0	50	0.67	5	0.072
200	2.2	37	0.50	1	0.0144

13.6.6 Noise Models of Amplifiers

The origins of current and voltage noise are physically diverse and may have widely different noise temperatures in a single device. These temperatures may or may not be close to junction temperature T_j. For instance, the base current shot noise of a BJT looks like the Johnson noise of a resistance of β/g_M at 150 K, whereas its extrinsic base resistance is a real resistance at T_j, so its voltage and current noise are predicted by the Johnson noise formula. Similarly, the collector current shot noise becomes an input-referred voltage noise contribution of $\sqrt{2eI_C}/(\beta g_m)$ (see Example 18.2). Since base and collector shot noise are uncorrelated, all of these contributions add in power. On the other hand, the base shot noise current times the extrinsic base resistance produces a noise voltage perfectly correlated to that part of the current noise. Most of the time these correlations are small, and when they're not small enough to ignore, a simple physical model beginning with uncorrelated noise contributions is perfectly adequate. Dostal gives the full analysis for difficult cases.

The effective noise temperature of an amplifier in a circuit may be nowhere close to ambient. For instance, an OP-27 op amp has $v_n = 3$ nV/Hz$^{1/2}$, $R_{\text{in}} = 6$ MΩ, so $T_N = v_n^2/(4kR) = 22$ mK for voltage noise, and $i_N = 0.4$ pA/Hz$^{1/2}$, so $T_N = i_N^2 R/(4k) = 17{,}400$ K for current noise. Noise temperature is thus a sensitive function of termination resistance and feedback components, and not a simple Thévenin or Norton model as with combinations of resistors. In addition, noise current and noise voltage behave differently in the circuit; noise current is a real current that will flow out of the amplifier terminals and will affect other circuit elements, as when we parallel amplifiers to get independent noise measurements in Section 17.11.6: the current noise contributions add but the voltage noise remains independent.

Example 13.4: Noise of an Inverting Amplifier. Let's do a concrete example: a CMOS op amp inverting amplifier, as shown in Figure 13.10. The LMC6032 is a jellybean 15 V CMOS dual op amp that typically has a 1.4 MHz gain–bandwidth product and 1.1 V/μs slew rate. Its specifications are not exactly ironclad; its data sheet has exactly one (1) guaranteed parameter: the slew rate is guaranteed to be at least 0.8 V/μs at 25 $^\circ$C and at least 0.4 V/μs from -40 $^\circ$C to $+85$ $^\circ$C. Not the first choice part for building space probes, but the price is right: under 40 cents each in quantity. At 1 kHz, the LMC6032 typically has 1-Hz voltage and current noise of 22 nV and 0.4 fA.

The noise contributions add as follows. The voltage noise is differential, so we can notionally put it in the noninverting input, where it obviously sees the total noninverting gain of the stage, which is $1 + |A_{V\text{cl}}|$. Since feedback is holding $V_{\text{in-}}$ at 0, there's ideally no voltage drop across R_I, so all current noise goes out through R_F, contributing $R_F i_N$. (All the current coming into the $-$ input gets summed and multiplied by R_F, which

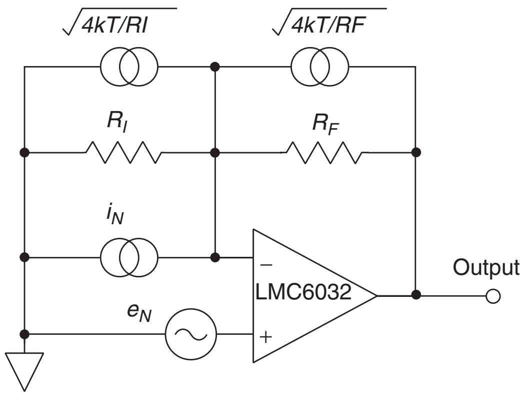

Figure 13.10. Noise model of an LMC6032 inverting amplifier.

is why it's called the *summing junction*.) The resistor noise contributions are calculated using the Norton model, so they contribute $\sqrt{4kT/R}$. The total current noise is multiplied by R_F and summed with the v_N contribution (all sums are RMS of course). Figure 13.11 shows the results, using the typical e_N and i_N at 1 kHz.

As expected, e_N dominates for small R_F and resistor noise for intermediate values. At very large R_F, the amplifier's current noise starts to dominate, but note that it's much less important at higher voltage gain, because e_N gets multiplied by $1 + |A_V|$ and the smaller $R_I (= R_F/|A_V|)$ contributes more noise current. The noise of the inverting ($A_V = -1$) amplifier is the same as that of a noninverting amplifier of $A_V = +2$, because for noise purposes the circuit is exactly the same: R_I and the + input are both grounded. Thus at $|A_V| = 1$, the available dynamic range is 6 dB less in the inverting configuration (see Section 15.11).

13.6.7 Combining Noise Contributions

It is important to remember that these v_N and i_N are RMS values of random fluctuations. Noise sources whose physical origins are diverse produce fluctuations that are completely statistically independent. They *never* cancel, unless derived from the same source at the same moment; when we add or subtract two voltages or currents, the noise contributions from each source always add in RMS fashion.

$$v_N^2 = v_{N1}^2 + v_{N2}^2. \tag{13.35}$$

Circuit designers and textbook writers always toss around noise sources like that, connecting purely notional noise voltage and current sources as though what they were doing had some physical significance, which it may not seem to. Beginners are understandably hesitant about doing this. While this caution is commendable, there are two reasons why the circuit designers' approach is the right one: first, all these noise contributions really do behave that way; and second, it isn't obvious how else to proceed. (Besides, as Figure 13.11 shows, usually one noise source dominates.)

The deeper reasons for the success of this approach are that the circuits we're interested in are operating in a highly linear regime, by design, and that the physical origins of these

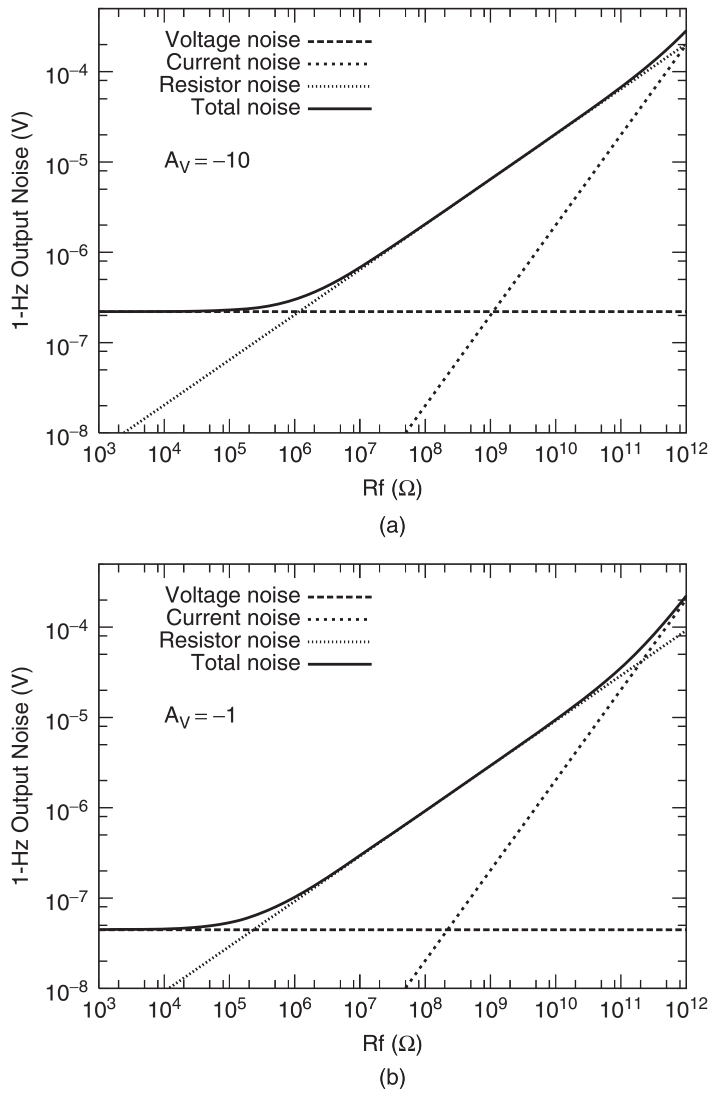

Figure 13.11. Noise performance of inverting amplifiers made from a jellybean LMC6032 op amp, calculated from typical specs: (a) $A_V = -10$ and (b) $A_V = -1$.

noise sources are diverse. Besides the naturally good linearity of op amps, transistors, and resistors, the small amplitude of the noise sources and the linearizing action of negative feedback ensure that the effect of all the noise sources combined will be precisely the rms sum of their individual effects. There are cases in which the correlations of the voltage and current noise are important, for example, when using huge resistors with bipolar op amps, so that the bias current noise is converted to a voltage, but the diversity of the fundamental noise sources means that a sufficiently low level circuit description will produce these correlations properly at the output. Since the output is almost always at a very low impedance, either in a 50 ohm system or an op amp output, it is unnecessary to separate voltage and current noise contributions there.

Aside: Noise Correlations. The fully general version of (13.35) includes a correlation term,

$$v_N^2 = v_{N1}^2 + v_{N2}^2 + 2Cv_{N1}v_{N2}, \tag{13.36}$$

where C is defined as

$$C = \langle v_{N1}(t)v_{N2}(t)\rangle / \sqrt{v_{N1}^2 v_{N2}^2} \tag{13.37}$$

and can have any value between -1 and 1 (we write $v(t)$ for instantaneous values and v for ensemble-averaged RMS values). Fully correlated noise ($C = 1$) could come from two cables connected to the same noise source, and nearly fully anticorrelated noise could come from opposite ends of a transformer winding whose center tap is grounded, or from the collectors of a differential BJT pair. Usually if you calculate assuming the noises are all uncorrelated (i.e., $C = 0$), and keep your eyes peeled for situations like these, where simple physics says they have to be correlated, you'll get the right answers.

13.6.8 Noise of Cascaded Stages

Calculating the noise of a cascade (several circuits in a row) is simple. Assuming the circuits are linear, the total noise is the RMS sum of the individual contributions. If there are M stages with noise voltage v_{Ni}, the output noise is

$$v_N^2 = \sum_{i=1}^{M} v_{Ni}^2 G_{ti}^2, \tag{13.38}$$

where G_{ti} is the total gain from the input of stage i to the output, that is, $G_{ti} = \prod_{j=i}^{M} G_j$, where G_j is the voltage gain of the jth stage.

Aside: RMS Noise Measurements. Not every instrument does a good job of measuring RMS noise. Many good ones, such as the HP 400EL AC voltmeter and your average lock-in, measure the average value of the rectified signal, and then (perhaps) apply a correction factor to get RMS. The problem is that the correction factor is correct for sine waves but wrong for noise (and everything else): the RMS/average ratio of a rectified sine is $\sqrt{2}$, whereas for Gaussian noise it's $\sqrt{\pi}$, which is 0.98 dB higher. So when doing noise measurements with an average-reading instrument, remember to add 1 dB. If you can get hold of a HP 3400A RMS voltmeter, it's the bee's knees for noise measurement: it has a 10:1 crest factor (i.e., maximum peak/RMS ratio). True RMS meters like the 3400A also get the right answer with more complicated signals, where the fudge factor technique fails.

13.6.9 Interference: What Does a Spur Do to My Measurement, Anyway?

An interfering signal at the same frequency as our signal causes beats,

$$A(\exp(j2\pi ft) + \epsilon \exp[j(2\pi ft + \theta)]) = A\exp(j2\pi ft)(1 + \epsilon\cos\theta + j\epsilon\sin\theta), \tag{13.39}$$

which shift its amplitude and phase,

$$\begin{aligned}\frac{\Delta A}{A} &= \sqrt{(1+\epsilon\cos\theta)^2 + (\epsilon\sin\theta)^2} - 1\\ &= \sqrt{1+\epsilon^2+2\epsilon\cos\theta} - 1\\ &\approx \epsilon\cos\theta \end{aligned} \tag{13.40}$$

and

$$\begin{aligned}\Delta\phi &= \arctan\left(\frac{\epsilon\sin\theta}{1+\epsilon\cos\theta}\right)\\ &\approx \epsilon\sin\theta. \end{aligned} \tag{13.41}$$

Thus a weak sinusoidal interference term, ϵ times as large as the carrier, causes a maximum amplitude change of ϵ times, and a maximum phase perturbation of ϵ radians. If the two signals are not at the same frequency, the phase θ is not constant, but the time-dependent amplitude and phase perturbations are given by the same formulas. These estimates allow spurious responses to be budgeted based on required system performance.

13.6.10 AM Noise and PM Noise

If the unwanted signal is not a pure sinusoid but a noise waveform, we cannot predict its exact time dependence without knowing that of the noise. We have to settle for predicting the statistical uncertainty of the amplitude and phase caused by the additive noise. This is greatly simplified by the linear dependence of ΔA and $\Delta\phi$ on the (small) relative amplitude ϵ of the interfering signal, which allows us to ignore the effects of different modulation components intermodulating with each other.

This means that for a given type of modulation, the performance of different demodulation schemes tends to be very similar when the SNR is large, and that the performance depends mainly on the strength and character of the noise and the bandwidth of the narrowest filter.

When the SNR becomes small, this convenient approximation is invalid; the phase and amplitude shifts are nonlinear and coupled. Furthermore, the behavior of different demodulation schemes can be very different then, so that extra thought and simulation are needed. For example, a phase-locked loop FM demodulator will lose lock, a limiter will start to suppress the signal in favor of the noise, and a diode AM detector will detect mainly the noise peaks. This is unfortunate, as the ultimate sensitivity of a measurement is defined by how well it handles the low SNR case.[†]

If the noise is filtered additive Gaussian noise, such as thermal or shot noise, then each frequency component is independent of all others, and also of the signal. The phase θ is therefore a uniformly distributed random variable, and the effects due to noise at different frequencies add in power. Thus we can get the statistics of amplitude and phase by taking the rms magnitudes of ΔA and $\Delta\phi$, using (13.40) and (13.41). Doing so, we get

$$\frac{\langle\Delta A\rangle}{A} = \frac{1}{\sqrt{2}}\sqrt{\frac{P_N}{P_C}} \tag{13.42}$$

[†]Hint: For AM and angle modulation of low m, use a synchronous detector; for wideband FM, use a PLL with wideband AGC and no limiter.

and

$$\langle \Delta\phi \rangle = \frac{1}{\sqrt{2}}\sqrt{\frac{P_N}{P_C}}. \tag{13.43}$$

We see that the quality of our measurement is determined by the quantity P_C/P_N, the power *carrier to noise ratio* (CNR). This is slightly different from the SNR, because what we consider signal may be the sidebands rather than the carrier. Because of the importance of this ratio, the noise and interference are usually specified in decibels with respect to the carrier (dBc). Noise and spurs at −40 dBc total will produce rms phase errors of 7 milliradians, and an rms amplitude uncertainty of ±0.7%.

Because these contributions enter linearly, the amplitude statistics of small AM and PM fluctuations are the same as those of the noise itself, up to a scaling. Thus weak additive Gaussian white noise generates amplitude and phase fluctuations whose statistics are also Gaussian. This is helpful in predicting the response of a detector to the noisy signal: if you put it into an AM detector, you see the AM noise; if you put it into a phase detector (running near null), you see the PM noise. PM is also the natural description for oscillators whose frequency is not constant, because weak white noise fed into a phase modulator produces flat-topped noise sidebands, like AM but unlike FM. Thus frequency fluctuation in oscillators is generically known as *phase noise*.

An FM measurement with fixed signal power and constant noise PSD gets quieter as m increases. The total phase shift (and hence the demodulator output voltage) grows as m, as does the required bandwidth, so that the noise power is proportional to m; this means that the total phase noise power goes as m. Since the signal power goes as m^2, the SNR at the demodulator output will go as m, at least until we leave the high-SNR limit.†

Aside: AC Versus DC Again. The assumption here is an AC measurement—once again, you can't phase-shift DC, so all the noise power winds up in the amplitude in a DC measurement. We already knew that, of course, because the noise power is precisely the RMS uncertainty in a measurement of the instantaneous signal power. This is another example of the effect of sideband folding, or alternatively of quoting noise in SSB or DSB terms. (See Section 13.1 for more on AC versus DC.)

13.6.11 Additive Versus Multiplicative Noise

The most basic distinction between noise sources is whether they depend on the level of the signal. Thermal noise remains the same (in W/Hz$^{1/2}$) whether there is a signal present or not; the noise waveform is just added to the signal. With constant DC photocurrent (e.g., in a heterodyne interferometer or strongly background limited system), shot noise is additive as well. Here the signal-to-noise ratio goes up linearly with signal power.

Multiplicative noise comes from noise modulating the signal source; examples are laser residual intensity noise (RIN) and the phase noise of an oscillator. In oscillators, the low frequency noise of the active device and the resonator, together with external effects such as power supply ripple, will modulate the amplitude and phase of the signal.‡

†After all our discussion of minimizing measurement bandwidths, this should strike you as odd. It is. (It's closely related to the SNR improvement from pulsed measurements in Johnson noise.)

‡While these are theoretically multiplicative, if their main effect is modulating the pedestal and not the signal, you can think of them as additive.

The effects of multiplicative noise get stronger (in W/Hz$^{1/2}$) as the signal gets stronger, so that the SNR is constant with signal level when multiplicative noise dominates. Multiplicative noise is frequently very obnoxious, since it is systematic in nature, so that the usual noise reduction methods based on averaging of one sort or another (e.g., narrow filters, lock-ins, and signal averagers) often fail. This is aggravated by the tendency of this sort of noise to be strongest at low modulation frequencies, as in baseband $1/f$ noise, which we've already encountered. Multiplicative noise is a form of unwanted modulation that puts noise sidebands on your signal (see Section 13.3.4), so it follows you wherever you go.

13.6.12 Noise Statistics

Noise whose average properties do not change with time is said to be *statistically stationary*. Since in practice most properties of noise waveforms are measured via time averaging, this statement may not seem to have much content—how could an average over all time change with time? The answer is that the properties of noise, such as its rms amplitude, instantaneous probability distribution function, autocorrelation, and so forth, are usually defined in terms of *ensemble averages*. That is, if we took a very large number of identical, noninteracting noise sources, we could measure the instantaneous noise voltage $V_N(t)$ at some time t from every one. From these measurements, we could construct a histogram of V for that t, and so on. This is not quite as silly as it sounds; there are many processes in which this could be done straightforwardly, such as imaging with a nearly uniformly illuminated CCD, where the signals from each pixel form an ensemble (after correcting for device nonuniformities of course).

Example 13.5: Ensemble Averages in Real Life. Two dozen solar telescopes pointing at the solar corona would have outputs consisting of the instantaneous coronal output at each point in the image, plus their own uncorrelated shot noise and technical noise. The coronal component would be highly correlated between receivers, in fact essentially identical apart from time delays depending on their position in space relative to the source. The technical noise would be statistically independent provided it did not depend on the signal level, variations in the common AC power line, local interfering signals, or those sorts of things. The outputs of the telescopes would be different functions of time, but drawn from the same *random process*. A random process is an ensemble of functions representing all possible outcomes of a given experiment.

The corona is far from time invariant, so that the statistics of the signal would be far from time invariant, but that doesn't bother the ensemble average.

More familiar noise sources, such as shot noise and Johnson noise, have statistics that depend on one or more parameters. Shot noise depends on the average current, and Johnson noise on the temperature and resistance. Time variations of these parameters will cause the noise to be statistically nonstationary. This is important in cases where the signal modulation is strong, for example, an interferometer whose beams are nearly equal in strength.

Noise whose time-averaged statistics are equal to its ensemble-averaged ones is said to be *ergodic*, and this class includes most fundamental noise sources, such as shot noise, Johnson noise, and so forth, providing that any parameters are time invariant. This is good, since it makes a lot of calculations and measurements easier, but bad in that it leads to confusing distinct but related concepts such as statistical autocorrelation and time autocorrelation.

Aside: Time and Statistical Autocorrelations. The time autocorrelation of any signal is

$$g \star g = \int_{-\infty}^{\infty} g(t')g(t'+t)dt'. \tag{13.44}$$

It is the inverse Fourier transform (in time) of the signal's power spectrum, a fact that can be proved from the convolution theorem in about three lines of algebra. This applies to any waveform whatever, assuming the integrals converge. There are different ways of normalizing it, but we'll ignore that fine point.

The statistical autocorrelation is a somewhat more slippery concept: it's still related to the average value of the function multiplied by a shifted replica of itself, but now it's an ensemble average over many different functions drawn from the same random process,

$$\mathrm{Corr}(g, g, t, \tau) = \langle g(t)g(t+\tau)\rangle. \tag{13.45}$$

Its Fourier transform is the ensemble-averaged power spectrum of the random process, but this is a much deeper fact, requiring the famous Wiener–Khinchin theorem, as well as more attention to normalization than we've given it.[†]

13.6.13 Gaussian Statistics

Noise sources such as shot noise and thermal noise arise from the random motions of very many charge carriers. In cases where their motions are not correlated with each other, the central limit theorem of statistics suggests that the instantaneous noise pulses contributed by each electron will sum up into a Gaussian probability density of voltage or current,

$$\frac{dP(V)}{dV} = \frac{1}{\sqrt{2\pi}\,V_{\mathrm{RMS}}} \exp\frac{-(V-\langle V\rangle)^2}{2V_{\mathrm{RMS}}^2}, \tag{13.46}$$

where V_{RMS} is the RMS voltage, considered as an ensemble average, and $\langle V\rangle$ is the mean voltage. Integrating $P(V)$ from V_0 to ∞ (10 standard deviations is infinite enough for almost everything) gives the probability of measuring a voltage greater than V_0.

$$\begin{aligned} P(>V_0) &= \frac{1}{\sqrt{2\pi}\,V_{\mathrm{RMS}}} \int_{V_0}^{\infty} e^{-(V-\langle V\rangle)^2/2V_{\mathrm{RMS}}^2} dV \\ &= \frac{1-\mathrm{erf}\left(\dfrac{V_0-\langle V\rangle}{\sqrt{2}V_{\mathrm{RMS}}}\right)}{2} \end{aligned} \tag{13.47}$$

For signal detection in noise, we're usually interested in finding all separate events larger than some given threshold level. The quantity $P(>V)$ is the relevant one in

[†]Even some of the big guys confuse the Wiener–Khinchin theorem and the autocorrelation theorem. This is a symptom of really not understanding the difference between deterministic and stochastic processes, so if you find some author doing it, keep your powder dry.

situations where you're sampling at predetermined times, and where your sample interval is large compared with the width of the impulse response of your narrowest filter (otherwise adjacent samples will not be uncorrelated).

If you're sampling more rapidly than that, or you have a continuously operating threshold circuit (e.g., a comparator), what you care about instead is the rate of upward going threshold crossings per second. This is a bit more complicated (see Papoulis for details) and depends on the details of the power spectrum of the noise, being proportional to the square root of the second derivative of the autocorrelation of the noise,

$$\lambda_{TC}(x) = \frac{1}{2\pi}\sqrt{\frac{-d^2R/dt^2}{R(0)}}e^{-x^2/2} \tag{13.48}$$

(where $x = V_0/V_{rms}$), which may not seem too useful since that parameter is not exactly at the tips of our tongues. (We could find it easily enough from the transfer function and the derivative formula for Fourier transforms).

The reason it's not a simple function of bandwidth is that a single noise event will tend to resemble the impulse response of the filter. A narrow bandpass filter has noise events that look like tone bursts, and so have many nearly identical peaks. The threshold crossing rate λ_{TC} is then many times higher than the true arrival rate of events. If we want only the first one of the burst, we're really asking for the upward going threshold crossing rate of the envelope of the analytic signal, which is a somewhat different matter. Since we typically need this number accurate only to the $\pm 20\%$ level, we can approximate it by assuming that our filter has a noise bandwidth B, and that the threshold crossing rate of the envelope is the same as if we were using a lowpass of noise bandwidth B. Using a brick-wall lowpass filter of bandwidth B, our upward going threshold crossing rate becomes[†]

$$\lambda_{TC}(x) = \frac{B}{\sqrt{3}}\exp(-x^2/2), \tag{13.49}$$

which is a very useful result.[‡]. Note that if we use a thresholding approach with tone burst signals, or with bandpass filtered noise, we'll need to put in a dead time of a few times the width of the filter impulse response, or we will get multiple triggers per event.

This function falls off astoundingly rapidly with V, as shown in Table 13.2. The last two rows show a decrease of nearly 1000 times for a 1.3 dB change in the threshold voltage. Because Gaussian noise falls off so rapidly at large multiples of V_{RMS}, at some point other noise sources that are less well behaved are going to dominate. Knowing where your nice Gaussian histogram stops falling off and starts being dominated by outliers is a matter of some concern in thresholding applications.

The combined statistics of signal + noise are *not* Gaussian. This is another case where we assume the problem doesn't exist, and fix it by hand afterwards—the power and generality of the Gaussian noise analysis are not to be lightly put aside.

[†]For bipolar thresholds, where a positive going crossing of V_T and a negative going crossing of $-V_T$ are both accepted as events, this rate has to be doubled.

[‡]S. O. Rice, Mathematical analysis of random noise. *Bell Syst. Tech. J.*, **23**, 3 (July 1944) and **24**, 1 (January 1945)

TABLE 13.2. False Alarm Rate Versus Threshold to RMS Ratio for Pure Lowpass Gaussian Noise

Normalized Excursion $(V - \langle V \rangle)/V_{rms}$	Normalized Probability Density $V_{rms} \cdot dP/dV$	Probability of Being Above Threshold $P(>V)$	Upward Threshold Crossing Rate λ_{TC} ($B = 1$ Hz)
0.0	0.399	0.50	0.577
0.5	0.352	0.309	0.510
1.0	0.242	0.159	0.350
1.5	0.130	0.067	0.187
2.0	0.0540	0.0228	0.0781
2.5	0.0175	0.00621	0.0254
3.0	0.00443	0.00135	0.00641
4.0	1.34×10^{-4}	3.17×10^{-5}	1.94×10^{-4}
5.0	1.49×10^{-6}	2.87×10^{-7}	2.16×10^{-6}
6.0	6.08×10^{-9}	9.87×10^{-10}	8.79×10^{-9}
7.0	9.13×10^{-12}	1.28×10^{-12}	1.32×10^{-11}

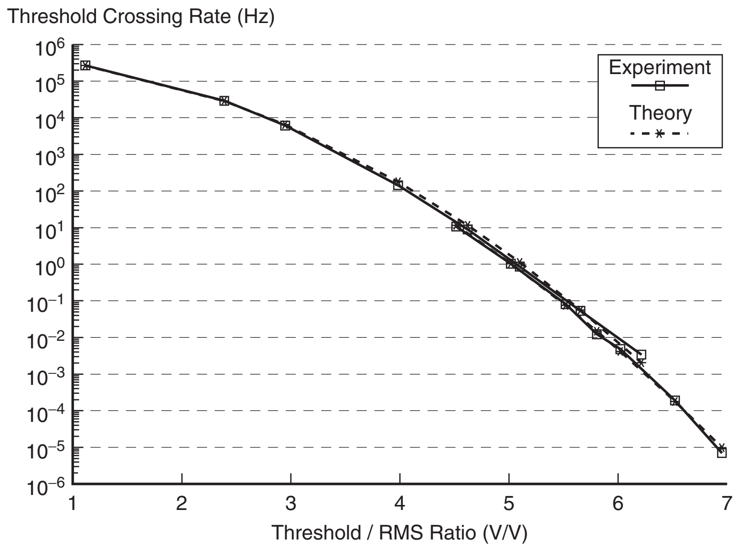

Figure 13.12. Shot noise statistics. False alarm rate of the ISICL sensor of Example 1.12, as a function of the bipolar threshold level, compared with theory.

13.6.14 Shot Noise Statistics

At high arrival rates, shot noise really is Gaussian. Figure 13.12 shows some false count rate data taken with the ISICL sensor of 1 in a 1.1 MHz bandwidth, compared with threshold crossing theory (13.49); note the close agreement over 10 decades of vertical scale. When translated into an equivalent amplitude shift by plugging the observed false alarm rate into (13.48) and solving for x, the imputed error is within 0.1 dB over the full range.

13.6.15 Thresholding

Thresholds cannot be set intelligently without knowledge of the noise statistics. We've already covered the Gaussian case, but there are lots of times when it doesn't apply, for example, if the dominant noise source is a fluctuating background. Unless you have reason for confidence (i.e., lots of measurements, not just statistical theory) to support your claim that your noise statistics are Gaussian, proceed with great caution. Don't jack up your thresholds unnecessarily out of fear, though. You're much better off compensating for effects such as source power variations in your threshold setting circuit. The RCA *Electro-Optics Handbook* has a useful discussion of thresholding with matched filters in noise, and Alexander has more depth.

13.6.16 Photon Counting Detection

If there are not a great many photons per second, the counting statistics begin to deviate seriously from the Gaussian case. A well-designed photon counting system with an average dark count rate λ (standard momenclature, unfortunately) and a quantum efficiency η has dark pulses that arrive in a Poisson process. Shot noise is also Poissonian, but since the average current consists of very many electrons in practical measurements, the shot noise is almost always Gaussian to extremely high accuracy. Photon counting systems, on the other hand, have noise that is dominated by large dark pulses resulting from single photoelectrons emitted by the photocathode. The noise variance in a time t from such events is $R = \lambda t$ and the probability of observing exactly M events in a time t is

$$P(M) = \frac{(\lambda t)^M e^{-\lambda t}}{M!}. \tag{13.50}$$

The probability of observing at least M events in a time t is

$$\begin{aligned} P(\geq M) &= \sum_{n=M}^{\infty} \frac{e^{-\lambda t}(\lambda t)^n}{n!} \\ &\approx \frac{(\lambda t)^M}{M! e^{\lambda t}} \left(\frac{1}{1 - \dfrac{\lambda t}{M} e^{1/2M}} \right) \left(\frac{1}{1 + \dfrac{0.4M}{(M - \lambda t)^2}} \right). \end{aligned} \tag{13.51}$$

The approximation in the second line of (13.51), which is valid for $\lambda t > 0.5$, has a relative error of less than 11% for all M for which $P(\geq M) < 0.1$, which covers all relevant cases for signal detection. For $\lambda t < 0.5$, direct evaluation of the sum becomes easy, since a two-term truncation of the sum is accurate within 4% for any value of M. Thus we can set thresholds with some confidence.

At low to moderate count rates, Poisson statistics have different asymptotic behavior than Gaussian ones; the exponential falloff of (13.50) for the $M = 0$ case is much slower than the Gaussian, so that to achieve a very high probability of detection, we need more photons per pulse for the Poissonian case than for the Gaussian one. On the other hand, assuming that our noise count rate is sufficiently small, the photon counting system can often detect an event with only a couple of photoelectrons. This is especially interesting since the information content is identical (remember the Rule of One). For detection schemes other than straight thresholding (e.g., coherent detection, QPSK, and so on), see Alexander.

13.7 FREQUENCY CONVERSION

Most high frequency measurement systems employ frequency conversion, which is the transferring of the desired modulation from one carrier frequency to another, perhaps moving an entire block of frequencies at once. It is used because signal detection, filtering, and digitizing are easiest at low frequency. Frequency conversion is exactly the same as SSB modulation, but with a different aim: to shift a (possibly modulated) signal from f_1 to f_2, by mixing with another signal at $f_3 = |f_1 \pm f_2|$. Often we want f_2 to be constant even though f_1 varies, so that subsequent stages can be better optimized. From the early days of radio, the frequencies involved have been known as the radio frequency (RF), intermediate frequency (IF), and local oscillator (LO), respectively.

Frequency conversion is also known as *heterodyning*. The term *superheterodyne* was early radio ad-speak for a receiver that had a tuned RF amplifier before its mixer stage, but it stuck. A superhet radio uses fixed-tuned IF filters, which are easily made with good skirt selectivity and passband accuracy. To cover a broad range of frequencies, the RF section of the radio is a frequency converter (a tuned RF amplifier followed by a mixer), which subtracts the tunable LO frequency from the RF frequency to get to the fixed IF. The choice of IF is nontrivial. (Figure 13.13 has some symbols for these devices.)

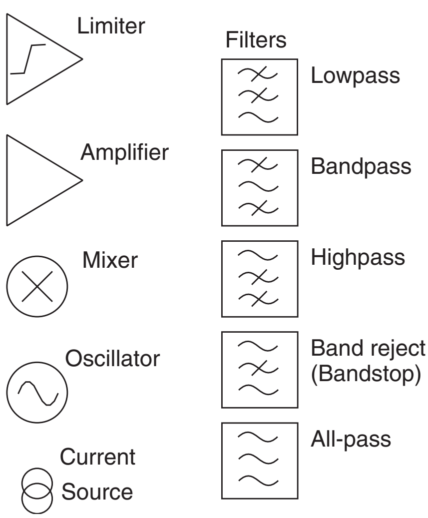

Figure 13.13. Some signal processing block diagram symbols.

13.7.1 Mixers

We encountered mixers in Sections 13.3.3 and 13.3.5 in the context of DSB modulation. Ideally a mixer is a device that algebraically multiplies (i.e., intermodulates) the signal voltage present on the RF port by the local oscillator signal at the LO port, to produce a signal at the IF port.[†] Mixers come in various types, which we'll see in more detail in Section 14.7.6. Any nonlinear circuit will produce mixing products—a mixer just does a better job, particularly a balanced mixer.

13.7.2 Choosing an IF

Frequency mixers never really produce just one output frequency. There are always strong components at $|f_{\mathrm{RF}} \pm f_{\mathrm{LO}}|$, one of which is the IF, and the other one is typically unwanted. A corollary is that a receiver with a given LO and IF can receive signals at more than one frequency. If the RF input contains many frequencies, then we must be concerned with another frequency, $f_{\mathrm{image}} = |f_{\mathrm{RF}} \mp 2f_{\mathrm{IF}}|$, called the *image* frequency, which will also produce a first-order signal at the IF. If the IF is too low, the front end filter will not roll off enough to reject it adequately.

The subtlety of choosing an IF lies in getting good enough image rejection, and spreading out the spurious signals enough that they don't hit your IF passband (which favor a high IF) while maintaining adequate IF selectivity, gain, and phase stability (which favor a low IF). Furthermore, unless you intend to build all your own filters or cost is no object, the commercial availability of good, cheap filters at frequencies such as 455 kHz and 10.7 MHz is an important consideration. It is usually best to start out with a filter catalog in your hand.

There are also the weaker spurious signals in the mixer's output to consider: LO feedthrough, RF feedthrough, and spurious mixing products[‡] at $f_{\mathrm{spur}} = |mf_{\mathrm{RF}} + nf_{\mathrm{LO}}|$, for all integers m,n. Their amplitude falls off toward larger $|m|$ and $|n|$, generally much faster with m. The LO harmonics are worse because mixers work much better if we drive the LO port hard: conversion loss decreases and overload characteristics improve.

Ideally, all these unwanted products will miss your IF passband, over the whole tuning range of the instrument. In practice, unless your measurement is narrowband, this may be difficult to accomplish. In radio applications, the RF input will probably have unwanted signals many times stronger than the one you want, but in electro-optical instruments the signal sources are much better controlled. In any case, make very sure that the strongest spurs (especially the LO and RF frequencies) don't cross your IF passband.

13.7.3 Image Rejection

If your RF signal varies from f_{low} to f_{high}, then in order that the image frequency never be in the passband, the IF must be greater than $(f_{\mathrm{high}} - f_{\mathrm{low}})/2$, plus some extra to allow the RF and IF filters to reject the image signal; this is why FM entertainment radios ($f_{\mathrm{RF}} = 88$–108 MHz) have a 10.7 MHz IF. This assumes that $f_{\mathrm{LO}} > f_{\mathrm{IF}}$; for the more

[†]This is not the same as an audio recording mixer, which sums instead of multiplying.

[‡]These spurious products are closely analogous to unwanted diffraction orders in grating spectrometers and have to be avoided in analogous ways.

general case, the image rejection condition is

$$\min(f_{\mathrm{LO}}, f_{\mathrm{IF}}) > \frac{f_{\mathrm{high}} - f_{\mathrm{low}} + B_{\mathrm{IF}}}{2}. \tag{13.52}$$

It was once customary to use double conversion, with a high first IF for image and spur rejection, and a lower second IF for high gain and good selectivity. Double conversion really works to reject the image but is often not as successful at rejecting spurs, because the two conversions generate a veritable forest of extra mixing products of their own.

Nowadays, high gain, high frequency amplifiers are not nearly so expensive as they once were, and good high selectivity filters are readily available at high frequency. Thus single conversion with a fairly high IF is usually best. The IF can be several times higher than the RF frequency in some cases.

Aside: Image Frequency Noise. It might seem that image rejection is not important in an instrument that doesn't have to cope with the vagaries of on-the-air reception, but that is not the case. Even if it doesn't contain interfering signals, the image contains shot and Johnson noise, which is uncorrelated with the desired signal's noise. All the noise contributions add in power at the IF, so failing to filter out the image noise before the mixer (or use an image reject [SSB] mixer) will reduce your SNR by 3 dB for no good reason. (This is the same effect as the noise difference between AC and DC measurements, which can also be avoided by SSB mixing—see Section 10.3.8.) If your devices are rolling off with frequency, the image noise may be even stronger than the signal's noise, making the SNR loss correspondingly worse.

Example 13.6: The AM Radio. Reducing noise and interference is not a new requirement. Your AM radio does an excellent job of rejecting out-of-band signals while pulling in the traffic report giving you the bad news about your commute. Similar techniques are very useful in electro-optical instruments.

Figure 13.14 shows how it works. The first stage is a tunable RF amplifier (i.e., it contains a tunable filter). This relatively crude filter rejects the worst out-of-band signals, reducing the opportunity for intermodulation, and also helps reject the image frequency. The gain helps overcome the loss and noise of later stages. The IF is 455 kHz, less than half the width of the AM band (540–1630 kHz). It is thus too low for the IF filter to adequately reject the image, which is one reason stations cluster near midband.

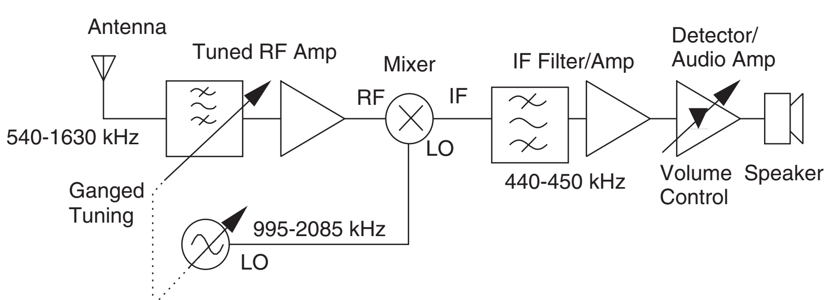

Figure 13.14. Block diagram of a superheterodyne radio, an archetypal noise-and-interference rejecting signal processing system.

A tunable *LC* local oscillator at 995–2085 kHz tracks the RF stage tuning, either by ganging the tuning capacitors together, or by adjusting the tuning voltage ranges if the stages are electronically tuned.

The mixer stage is usually a single bipolar transistor running in a strongly nonlinear regime. This is a poor mixer, but it's adequate to the purpose and is very cheap. AM radio ICs have better mixers, usually Gilbert cell multipliers.

A fixed IF filter, usually two solid ceramic resonator devices in cascade, provides about a 10 kHz bandwidth, with good skirt selectivity. A high gain IF amplifier with slow automatic gain control (AGC) presents a nearly constant average power level to the detector, which produces low level audio output. The AGC stage ensures that there is enough power delivered by the IF to the detector, without distorting the envelope by overdriving the amplifiers. The RF amplifier is also occasionally AGC controlled. Discrete detectors are usually half-wave rectifiers followed by *RC* filters, but integrated ones are almost always multipliers, which multiply the signal by a clipped replica of itself. An adjustable gain audio amplifier and speaker complete the signal path.

The ability to use good fixed-tuned filters means that the filter characteristics do not vary with tuning, as they inevitably do with tunable stages. Even more than that, though, the fixed IF frequency relaxes the requirements on the signal and AGC detectors, and allows the IF strip and back end to be tested easily as a separate subsystem, which is very important in system development and production.

13.7.4 High Side Versus Low Side LO

A frequency plan is more than just choosing an IF; there's also the LO to consider. Unless your IF is at DC, there are always two possible choices of LO: $f_{LO} = |f_{IF} \pm f_{RF}|$. In a wideband measurement involving down conversion, the choice is usually obvious. For example, in an AM radio, where $f_{RF} = 540$ to 1630 kHz and the $f_{IF} = 455$ kHz, a high side LO would be 995 to 2085 kHz, whereas a low side LO would be 85 to 1175 kHz.

A low side LO would be disastrous for two reasons. First, spurious responses: the LO and its second through fifth harmonics would cross the IF, leading to huge spurs from LO feedthrough. Near the low band edge, the LO harmonics would lead to a huge number of spurious responses to other stations throughout the band. Second, circuit requirements: it is much more difficult to make a good stable oscillator tunable over a 14:1 range than a 2:1 range. We couldn't use an *LC* oscillator, because the capacitance would have to change nearly 200:1. Fortunately, in an AM system, the symmetric sidebands allow LSB conversion with no frequency-inversion problems (see Section 13.3.6).

13.7.5 Direct Conversion

Synchronous receivers such as lock-in amplifiers generally work with the LO at the center of the RF passband—that is, the IF is at DC. Mathematically, this is equivalent to multiplying $A\cos(\omega t + \phi)$ by $\cos\omega t$—the envelope is preserved, but the quadrature component of the modulation is lost due to sideband folding. To preserve it, we can use SSB mixing (see Sections 10.3.8 and 13.8.7), use I/Q demodulation, which is what a two-phase lock-in does†, or phase-lock to the carrier (Section 13.9.3).

†See also Section 13.9.6.

13.7.6 Effects of LO Noise

It is often said that mixers are insensitive to LO noise, but this is somewhat misleading. A diode bridge mixer, for example, is usually operated in a switching mode, with a LO signal large enough to turn the diodes completely on and off each cycle. A small variation in LO power will not significantly change the diodes' resistances, so it causes a much smaller than linear change in the IF output amplitude. LO phase noise is of course passed unattenuated; the instantaneous phase of the IF output is the sum or difference of the instantaneous phases of the RF and LO signals. This will degrade FM/PM and SSB measurements. Phase noise in the LO does not greatly affect AM measurements, because the envelope amplitude is independent of the carrier frequency. The only time you get into trouble with phase noise in AM is if it's so large that it extends to the skirts of the IF filter. In that case the filter will turn the PM into AM, so you have to increase your IF bandwidth to accommodate it. It's the same in optics—a LED has horrible phase noise but excellent amplitude stability.

13.7.7 Gain Distribution

For stability and low spurious generation, it is important to distribute the gain between the RF, IF, and baseband stages wisely. The mixer usually has the lowest dynamic range of any stage, so put just enough RF gain in to override its noise. Putting too much gain in the IF is liable to lead to oscillation, especially if the layout is too tight. Because any pickup from the baseband stages is at a different frequency, it is rejected by the interstage filters and does not cause oscillation. (It's also at lower frequency, so it's much less likely to couple where you don't want it.) There are power issues as well: there's no point in pumping a 1 kW signal through a 100 dB loss—unless you're building a radar, you can get better SNR at much lower cost and lower power by distributing the gain and loss a bit more evenly.

13.8 FILTERING

Just about any transfer function describes a *filter*, loosely defined as something that treats different frequencies differently. Filters are very important, because they allow us to choose the frequency components we want, and to modify their amplitude and phase as required by the signal processing scheme. The details of the different kinds of filters, their characteristics, and some design information are covered in Section 13.8.9. Here we stick to the systems design and theoretical issues.

Most of the time, what we want is a filter that passes what we care about and rejects everything else. Typical filters of this type are lowpass, which reject everything above the cutoff frequency f_c, and highpass, which reject everything below f_c. Bandpass and bandstop (band reject) filters similarly pass or reject everything between f_1 and f_2. We'd like to reject the bad stuff (the stopband) completely, while keeping all of the good stuff (the *passband*), and have a vertical cliff separating the two regions, but are willing to allow a *transition band*, as shown in Figure 13.15. The ratio of output power to input power is called the *insertion gain* (its reciprocal is the more commonly heard *insertion loss*), and the phase shift between input and output is the *insertion phase*. The flatness of the passband is the ratio of the peak-to-average insertion gain ratio, or alternatively the square root of the peak–valley ratio, and is normally specified in decibels.

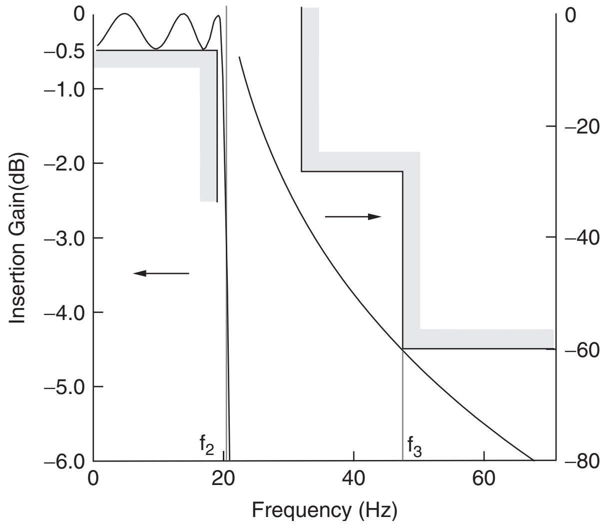

Figure 13.15. Filter specification, showing passband (left), stopband (right), and skirt selectivity requirements, together with a Chebyshev filter design that just meets the specification.

In signal processing, we often want at least one edge of the filter to fall off very steeply. The trailing-off edges of the filter's passband are called the *skirts*, and the skirt selectivity is expressed by the *shape factor*, the ratio of the filter's bandwidths measured at 60 dB down and 3 dB down (f_3/f_2 in Figure 13.15). A shape factor of less than 3 is pretty sharp, but fancy filters can sometimes achieve 1.1. (The time-domain response of such a filter is not pretty to look at.) A one-pole RC has a shape factor of 1000.

13.8.1 Cascading Filters

Most signal processing systems involve at least a few stages of filtering. Remember that the transfer functions of cascaded filters multiply each other. In the process, the passband gets narrower, and any ripple usually gets worse. Consider two identical bandpass filters, whose 3 dB bandwidth is 10 to 13 kHz. When the two are cascaded, these frequencies are down by 6 dB instead of 3. If the top of the filter is roughly Lorentzian in shape, like a single *LC* section, the 3 dB bandwidth will have decreased by a factor of $1/\sqrt{2}$. Filters with more nearly rectangular transfer functions will narrow less; pointy ones will narrow more.[†]

If the filters are passive *RLC* units, whose transfer functions depend on the source and load impedance, then cascading two or more of them will change their shapes, because a passive filter usually looks like a nearly pure reactance in its stopband, which effectively changes the *LC* values of the end sections of the filter. Putting a pad (attenuator) or a well-matched amplifier in between will help.

Deep in the stopbands, the rejection may not be as good as the product of the two filters. Eventually you're going to be dominated by other signal paths: capacitive and

[†]Examples are skin-depth attenuation in coaxial cables, which goes as $f^{-1/2}$, so that doubling the length of a cable will reduce its bandwidth by *four times*.

inductive pickup, power supply coupling, slightly wiggly grounds, that sort of thing. Don't be surprised if cascading two filters with 70 dB rejection results in 92 dB rejection rather than 140 dB, at least at first. A useful rule of thumb is that in a filter you can get 60 dB of attenuation per inch of separation on the board, but that anything more is hard, requiring shields and extra care.

13.8.2 Impulse Response

We specify filters most often in the frequency domain. Since the transfer function multiplies the input spectrum, the frequency response H is the transform of the output signal produced by an input whose transform G is 1 for all f. Equivalently, h is the output signal in the time domain produced by an input g whose transform is 1—that is, an impulse function, $\delta(t)$. For this reason, $h(t)$ is called the *impulse response* of the filter.†

13.8.3 Step Response

In many applications, the response of the system to a step change in the input is the most important parameter. After a sudden change, the filters and amplifiers must slew to the new value and settle there to high accuracy before a measurement can be made with confidence. The figures of merit are the rise and fall times, usually quoted between 10% and 90% of the total change, and the settling time, which must elapse following the input step before a good measurement can be made. Rise and fall times are set almost entirely by the filter bandwidth, and the settling time by a combination of the bandwidth and group delay characteristics. (This assumes that the network is highly linear; nonlinear settling is considered in Section 14.7.15.) Lowpass filters settle roughly twice as fast as bandpass filters with the same bandwidth. (As usual, this assumes that the low frequency cutoff of the bandpass filter isn't too close to 0—the bandpass is even slower if it is.) Normally we don't worry a great deal about the settling performance of a highpass filter, because we want the highpass response to be constant over our measurement time (i.e., slower is usually better).

13.8.4 Causality

The causality condition says that effects do not precede their causes, so $h(t) \equiv 0$ for $t < 0$. Enforcing this condition in the Fourier domain is exactly analogous to what we did in the time domain to form the analytic signal—chopping off all negative frequencies—except with the opposite sign of j in the imaginary part due to flipping the t axis. (See Section 13.2.5.) Thus the real and imaginary parts of H must be a Hilbert transform pair:

$$\Im\{H(f)\} = \mathcal{H}\Re\{H(f)\}. \tag{13.53}$$

This constraint is important in the design of filters, since it enables us to take a frequency-domain specification for $|H|^2$ and turn it into a possible filter design. A less palatable consequence is that it specifies a minimum *delay* for a given filter response.

†For the mathematically inclined, it is the Green's function for the differential equation expressed by the transfer function of the filter, with the initial condition on the derivatives that $h^n(0^-) = 0$ for all n.

It is reasonable that a filter which can distinguish sharply between frequencies f and $f + \Delta f$ should require a time delay on the order of $1/\Delta f$ seconds to do it. More subtly, a filter with a cutoff frequency of 1 kHz should be able to distinguish between DC and 2 kHz faster than it can between 900 and 1100 Hz. Thus the time delay through a filter conceptually need not be exactly the same for all frequencies, and in fact generally is not.

Aside: Stable and Unstable Transfer Functions. One of the most basic constraints on the impulse response of a filter is that it should not grow without bound as $t \to \infty$—otherwise we've inadvertently built an oscillator. This happens all the time with active circuits, but would be strange to say the least in a pure *RLC* filter. The Fourier transform of a rational function can be found by residue calculus (don't worry, no branch cuts here—just isolated poles). It's worth doing this, because it shows the deep connection between stability and causality.

Consider a rational transfer function $H(f)$, which has been decomposed into partial fractions:

$$H(f) = \sum_{k=1}^{N} \frac{r_k}{f - f_k}, \tag{13.54}$$

where r_k is a complex constant. (This form assumes that the order of the numerator is not greater than that of the denominator, which must always be true for passive networks. Why?) The inverse Fourier transform of this is

$$h(t) = \sum_{k=1}^{N} \int_{-\infty}^{\infty} \frac{r_k}{f - f_k} e^{j2\pi f t}\, df, \tag{13.55}$$

which we propose to evaluate by residues, using a contour consisting of the real axis plus a semicircle of large radius R, which we will let go to infinity. (Note that if $h(t)$ does not die away at large times, the integral along the real axis will not exist, so we can't use this particular contour with unstable transfer functions.) As always, poles inside the contour of integration contribute to the integral, whereas those outside do not. Poles above the real f axis will contribute residues with exponentially decaying time dependence, so their contributions are stable as $t \to +\infty$. Poles below the axis will contribute an exponentially growing time dependence, so stable transfer functions must not have poles in the lower half plane.

When we go to close the contour, using Jordan's lemma to let us neglect the contribution from the large semicircle, we find that the upper half-circle contributes 0 for positive times, but that the integrand is unbounded there for negative times. We therefore close the contour above for positive times, and below for negative times.

Requiring causal behavior means that there can be no poles inside the contour at negative times, so causality forbids any poles whatsoever to be in the lower half-plane. Thus our requirements on stability and causality lead to the same condition on the poles of H, a remarkable result first proved by Titchmarsh.†

†E. C. Titchmarsh, *Introduction to the Theory of Fourier Integrals*, 2nd ed. Oxford University Press, Oxford, UK, 1948.

13.8.5 Filter Design

If you don't care much about the time-domain response of your filter, or if it is much wider than the bandwidth-setting filter (usually the second IF filter), you can just pick the power spectrum $HH^*(f)$, factor the denominator into complex-conjugate poles, discard the ones in the lower half-plane, and you have a causal filter function. Then you have to synthesize a network with that transfer function, which is the hard part. We'll talk more about the details in Section 15.8, but first we need to understand the problem.

13.8.6 Group Delay

In Section 1.2.2, we encountered dispersion, where different colors of light traveled at different speeds, and group velocity, where the modulation traveled at a different speed than the phase (always more slowly). All causal, frequency-selective systems exhibit the same sort of behavior. In optics, the causality condition is called the Kramers–Kronig relation, but it's just the same idea.

If the phase shift of a sinusoidal signal passing through the filter is $\phi(f)$, then the phase delay is just $\phi/(-2\pi f)$, but the group delay, which is what the modulation sees, is $(-1/2\pi)\partial\phi/\partial f$.

A pure time delay of τ phase shifts each component by $-2\pi f\tau$ radians; its graph is a straight line with no offset. This delays the modulation but does not distort its shape. Networks whose delay has this characteristic, at least approximately, are called *linear phase networks*. Coaxial cable and Bessel filters are two examples. Variations in the group delay across the passband distort the relative phases of the Fourier components of the modulation envelope. Filters whose group delay varies all over the place in the passband have weird time-domain behavior, typically surfacing as severe overshoot and ringing on their step responses. Group delay in the stopband is not of much consequence, since signals out there don't appear in the output, but the transition band matters nearly as much as the passband.

As we might expect, there is a trade-off between sharp cutoff and good group delay performance. For a given filter order (number of poles + number of zeros), sharper filters have larger group delay in general, with delay being greater in the passband than in the stopband. Sharp filters also exhibit tall, narrow group delay peaks near the band edges, which is reasonable, as we mentioned earlier. Elliptic filters are the worst, followed by Chebyshev, Butterworth, Bessel, Gaussian, and the champion, the equiripple group delay filter (see Zverev).

Linear phase filters whose group delay is as small as causality permits (the so-called minimum phase condition) all have a rather gentle slope in their transition band and significant bowing in the passband, which makes them far from ideal for rejecting spurious signals and especially for cascading. This property is not improved by adding more poles to the filter, because when you do this the filter approaches a Gaussian response, which is no flatter.

All is not lost, however; it is possible to synthesize networks that have varying group delay but whose transfer function has unit magnitude. These are generically called *all-pass filters*, in this case specifically *group delay equalizers*. Putting in an equalizer will at least double the number of elements in the filter but will yield sharp filters with flat group delay (albeit several times longer than the peak delay of the original minimum phase filter).[†] The required order of the equalizer increases more rapidly than the filter order, and especially with Cauer filters, it becomes prohibitive quickly as the filter becomes sharper. Have a look at Zverev or *Reference Data for Engineers* for lots more design information. You can buy filters like this, so you don't need to become a filter design whiz to make your measurement work. If group delay is really killing you, consider using a FIR digital filter, which can have exactly constant delay (see Section 17.6).

Aside: Ladder Networks. Most filters built from discrete *RLC* sections are *ladder filters*—they are connected with series and shunt sections, and their schematic looks like a ladder. All passive ladder networks are automatically minimum phase,[‡] so group delay equalizers and other non-minimum-phase filters must have more complicated connections (e.g., the lattice filter of Section 15.20).

13.8.7 Hilbert Transform Filters

It is frequently very useful to construct the envelope of a function, e.g. a tone burst of several cycles; this makes an especially nice AM detector and is good for detecting the amplitude of pulses. Other times, we want to build image reject mixers (single sideband mixers), where we use the phase relationships between upper and lower sidebands to reject one and pass the other. SSB mixers are useful in rejecting image frequency noise, which is sometimes good for a 3 dB improvement in measurement sensitivity. Supporting both of these worthy aims is the *Hilbert transform filter*. An ideal Hilbert transformer, remember, would apply a 90° phase shift to all positive frequencies and reject all negative ones. Thus we want to apply a phase shift of exactly 90° to all signals we care about, without changing their amplitudes. Unfortunately, it's impossible to do this. What we actually do is make two all-pass filters, whose phase shift differs by 90°, within some accuracy, over some given bandwidth. This is much easier, especially for relatively narrow bandwidths (an octave or less). The two outputs, notionally 0° and 90°, are known as I and Q, respectively, for in-phase and quadrature. The basic scheme is called *vector modulation*.

Provided the modulation sidebands don't go down too close to DC, the analytic signal is $I + jQ$, and its magnitude is

$$|\Phi|^2 = I^2 + Q^2. \tag{13.56}$$

This can be constructed by several means. Section 15.5.4 has a cute trick used in radar to make a good approximation of $|\Phi|$, very fast and without massively increasing the required dynamic range by squaring the signal.

Example 13.7: Building an SSB Mixer. Single sideband mixers can be made from two power splitters, two mixers, and a power combiner (which is a splitter run in reverse). The LO and RF splitter/combiners are quadrature devices,[§] which contain the I-Q filter

[†] Since we can't go backwards in time, equalizers have to work by adding group delay to the quick parts, not taking it away from the slow ones. Thus the total delay through the equalized filter will be greater than the maximum delay of the original one.

[‡] H. W. Bode, *Network Analysis and Feedback Amplifier Design*. Van Nostrand, New York, 1945, p. 243.

[§] The 90° power splitters are also called quadrature hybrids.

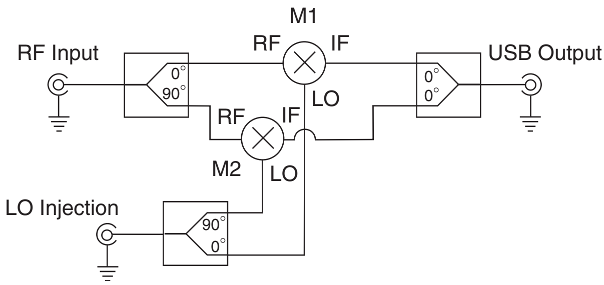

Figure 13.16. Single sideband (SSB) mixer: due to the opposite phase dependence of upper and lower sidebands, making two DSB signals with out of phase sidebands allows the unwanted sideband to be canceled, producing an SSB output.

pair inside them. The block diagram is shown in Figure 13.16. Assuming the RF input is $\sqrt{2}\sin\omega t$, the LO injection is $\sqrt{2}\sin\Omega t$, and ignoring for the moment the common-mode phase shift in the quadrature splitters, the outputs of M1 and M2 and the resulting summed IF voltage are

$$\begin{aligned} V_{\mathrm{M1}} &= \cos\omega t\cos\Omega t \\ &= \tfrac{1}{2}\cos(\omega+\Omega)t + \tfrac{1}{2}\cos(\omega-\Omega)t \\ V_{\mathrm{M2}} &= \sin\omega t\sin\Omega t \\ &= \tfrac{1}{2}\cos(\omega+\Omega)t - \tfrac{1}{2}\cos(\omega-\Omega)t \\ \hline V_{\mathrm{IF}} &= \cos(\omega+\Omega)t \end{aligned} \tag{13.57}$$

which is a pure USB output. Switching the outputs of one of the quadrature splitters will produce an LSB output instead. Equivalently, you can subtract the two IF signals instead of adding them, and even have both an adder and a subtracter, producing both sidebands separately. The unwanted sideband rejection obviously depends on excellent phase and amplitude matching between the two signal paths. The degree of matching required can be determined by the kind of simple analysis we used in Section 13.6.9: 40 dB sideband rejection requires an amplitude error of <1% , a phase error of <0.01 radian, or some combination of smaller errors in each. Such close matching of signal processing components will require one or two trims in general.†

The 0° and 90° phase LO signals can also be made by dividing a signal at $4f_{\mathrm{LO}}$ by 4 in a synchronous ÷4 walking ring counter (Johnson counter), or at UHF by putting a quarter-wave piece of transmission line in one arm of a 0° splitter.

13.8.8 Linear Phase Bandpass and Highpass Filters

As we'll see in Chapter 15, the usual procedure for designing a bandpass or bandstop filter with a 3 dB bandwidth from f_1 to f_2 involves taking a prototype (i.e., canned)

†The difficulty of this is one reason for the current trend toward DSP radios, in which the multiplication is done numerically after digitizing. That takes a lot of engineering and isn't always possible.

lowpass filter design whose bandwidth is about $\sqrt{f_1 f_2} - f_1$, and resonating all its elements at $f_0 = \sqrt{f_1 f_2}$.† The resulting frequency transformation maps the DC response of the lowpass to f_0 and makes the skirts symmetrical when plotted against $\log(f)$. Similarly for a highpass, you swap inductors and capacitors, keeping $|X(f_c)|$ the same, which maps f into $f_c{}^2/f$. Unfortunately, these nonlinear frequency stretches mess up the linear phase characteristic completely, so you can't build linear phase highpass and bandpass filters this way. Packages exist for designing filters of this sort, but we often wind up writing programs to do it.

Another subtlety is that in order for a bandpass or highpass filter to have good pulse fidelity, the trend line of the linear phase has to pass through the origin, or at least be offset by $2n\pi$ radians, and this need not be the case. Roughly speaking, if it doesn't, the carrier will be phase shifted with respect to the envelope, just as in group velocity dispersion in optics. This isn't too serious for narrowband filters, but for wider ones (say, >15% fractional bandwidth) you should check for it, because the position of the carrier peak within the envelope matters. Remember that if the trend line goes through (0 Hz, $\pi/2$ rad), you've built a bandpass Hilbert transformer, which may not be what you want.

13.8.9 How to Choose a Filter

Choosing a filter for your application starts with knowing what you need it for. A system that needs good pulse response while minimizing noise bandwidth will require a filter with good group delay characteristics. If the filter is inside a feedback loop (e.g., a fast AGC loop), the total delay is important as well as its flatness. In cases like that, a Bessel, Gaussian, or equiripple group delay filter is indicated. Most of the time, though, a little extra delay is not an excessive price for better skirt selectivity and narrower noise bandwidth for a given passband flatness, so that a group delay equalized Chebyshev or Cauer filter will do a better job.

The group delay characteristic is normally of interest only in the narrowest filter in the signal path. Elsewhere, we are concerned with rejecting images and spurious signals, and minimizing the noise bandwidth, while using the cheapest filter that will do the job. Cheapness has different meanings depending on what the intended use of our system is, and on how many are to be made. For one-offs or experimental systems, the optimal setup is a drawer full of Mini Circuits filters for chopping off harmonics, plus a few hand-wound coils and variable capacitors for the slightly more difficult cases, and perhaps a custom filter with a very flat passband and good skirts for the bandwidth setting filter.

Bandpass or bandstop filters with very wide bands (an octave or more) are difficult to build, so consider cascading a lowpass and a highpass instead. Narrow bandpass and bandstop filters have much steeper skirts than highpass or lowpass filters with the same complexity and cutoff frequency—the characteristic scale of the decay is δf rather than f_0. This is because a bandpass filter at $f_0 \pm \delta f$ is really a lowpass filter of bandwidth $2\delta f$ (see Section 15.8.6) whose elements have each been resonated at f_0 by an appropriate series or parallel element paired with each one; thus the skirts tend to roll off as inverse powers of $(f - f_0)/(\delta f)$ until the finite Q of the inductors limits them.

†Since the passband is mirrored about f_0, why isn't the equivalent lowpass bandwidth $(f_2 - f_1)/2$?

If you need to make your own *LC* filters, have a look at Chapters 15 and 19 for how to design and adjust them. There's a whole lot about practical filter prototyping in the ARRL handbook.

13.8.10 Matched Filtering and Pulses

Consider a narrowband signal in constant-amplitude additive white noise (e.g., Johnson noise but not shot noise), running into a linear phase filter of adjustable bandwidth. For discussion's sake, the signal is a train of narrow tone bursts 10 carrier periods long, spaced by 1000 carrier periods. A very narrow bandwidth will essentially sample the signal PSD at the carrier frequency, and the SNR will be just (signal PSD)/(noise PSD), evaluated at the carrier frequency. Seemingly, this is the best case, because the signal PSD drops off from there, and the noise PSD is constant. It would seem that the only reason to open up the filter would be to trade off SNR versus measurement speed, or possibly to overcome additive noise in later stages.

This analysis is exactly correct for a linear, time-invariant system with no additive noise after the filter, and so is relevant to unmodulated carriers. With a pulsed waveform, though, there is no information between the pulses, so we can turn off the receiver then and eliminate most of the noise. We'll put in this time-varying behavior by hand as usual, by paying attention to the instantaneous SNR within the pulse, instead of the average SNR.

Figure 13.17a shows a series of pulses in additive noise, with a SNR near 0 dB. Applying a narrow filter as in Figure 13.17b gets rid of most of the noise but also smears the pulses out completely, leaving only the DC, and that at a very low level because of the small duty cycle. As the filter gets wider as in Figure 13.17c, the total SNR starts out constant and gradually degrades, because at some point each additional hertz lets less and less signal get in, but the same amount of noise. The key observation, though, is that not only does the average power increase linearly with bandwidth (because of the nearly constant PSD close to the peak), but the pulses get shorter by a factor of $1/B$ too. Thus the instantaneous signal power in the pulses goes as B^2 for small bandwidths, and the instantaneous SNR within the pulse goes as B.

We keep winning by increasing B until the filter no longer limits the pulse width, and then even the peak-power/noise ratio starts to decline again, as we'd expect. Thus from the instantaneous SNR point of view, the best filter for this application is a linear phase bandpass, centered on the carrier, adjusted so as to stretch the pulses by a bit but not too much—say $\sqrt{2}$ times. The overall SNR improvement is maintained by thresholding (i.e., ignoring signals weaker than a threshold), by turning off the receiver between pulses, or ideally (as in Figure 13.17d), multiplying by the input pulse shape.†

Adding group delay dispersion to the filter will change it. This can be a good thing in special cases, for example, grating pairs used as pulse compressors in femtosecond lasers, but usually it hurts you by smearing out the pulses and adding artifacts without improving the noise.

As we can see, the choice of the narrowest filter in the signal path is the thorniest and really requires you to know what you're trying to optimize. If you like calculus of variations, you can easily prove that the maximum instantaneous SNR for a statistically stationary, deterministic signal in the presence of additive white noise is achieved with a

†These are both nonlinear operations, which is how we can get better total SNR than a linear analysis suggests.

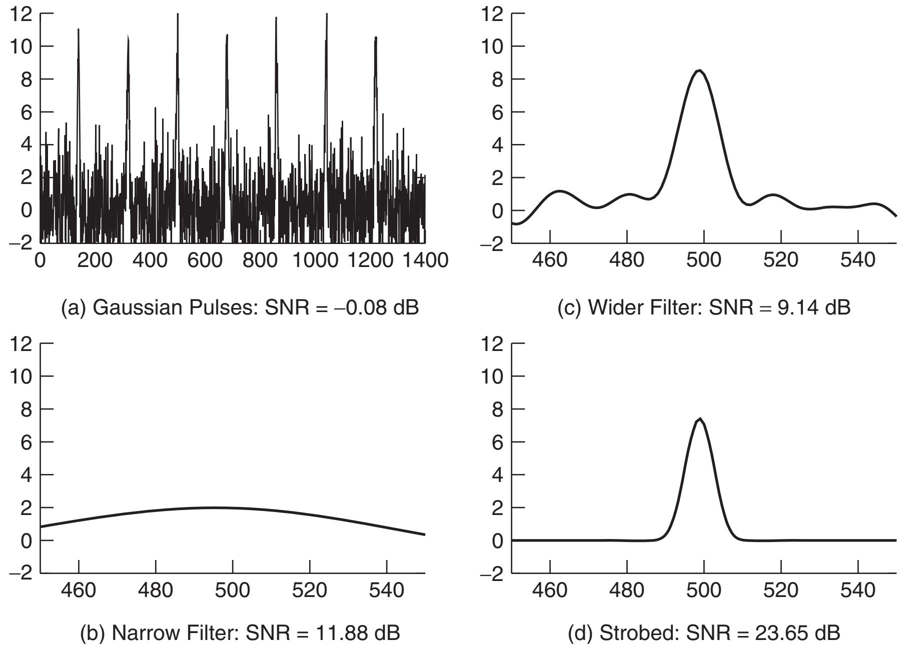

Figure 13.17. Pulse detection in the time domain. This is a fairly typical case, with a duty cycle on the big side (3%). The improvement is more dramatic with narrower pulses. Strobing is done by multiplying by the input pulse shape (Gaussian).

matched filter, that is, one whose transfer function is the complex conjugate of the signal spectrum (the complex conjugation undoes any group delay smear, see Example 1.12). If the noise isn't white, the optimal strategy is to whiten it (i.e., apply a filter that makes the noise PSD constant), apply the matched filter, and then unwhiten again to recover the narrow pulse shape.

If you are a bit unsure about what the best possible filter is for your application (and who isn't, at least to start with), choose a filter whose passband is fairly flat, and whose 3 dB bandwidth coincides with the 6 dB bandwidth of your signal waveform. Get one that has good group delay performance, no more than $\pm 2^\circ$ deviation from linear phase over the passband and transition band. This filter will be reasonably close almost always, and it errs in the right direction, namely, by being slightly too wide. Once you have your system working with this filter, you can fiddle with optimizing it to wring out the last decibel or so here, and find a cheaper filter in the process. The optimum is pretty flat, so it is unnecessary to bleed over errors of 10–15%. Note that we have implicitly assumed here that the desired signal is slowly varying, so that we don't need to consider sampled-data effects. Depending on the situation, these become important for signal bandwidths above 1–15% of the rep rate f_{rr}; see Chapter 17 for how to handle this case.

Aside: Matched Filters and Signal Averaging. In case you're worrying about signal averaging in this situation, a repetitive signal can indeed be improved by signal averaging after detection in a wide bandwidth, but in that case the signal PSD is zero except near harmonics of $1/P$, which are the only frequencies that survive averaging, so the signal averager is functioning as a matched filter.

Example 13.8: SNR Improvement with Pulsed Measurements. You don't do a pulsed measurement by shuttering a CW beam, but by crowding the available photons together in a high peak power pulse. If the average optical power is being kept constant as the duty cycle is reduced, the peak optical power and the peak photocurrent both go as d^{-1}, and peak electrical power as d^{-2}. If the measurement bandwidth needs to be at least $f_{\max}$, the pulse repetition rate f_p must be greater than $2f_{\max}$, so that the signal and noise are multiplied by the gating function $s(t)$,

$$s(t) = \text{rect}\left[\frac{f_p t}{d}\right] * \text{Ш}(f_p t),\; 0 < d \le 1, \tag{13.58}$$

where Ш is the sampling function (Cyrillic capital *sha*), a train of unit strength delta functions spaced at unit intervals. Note that this function becomes identically 1 as d goes to 1, so that we can connect the results for small d to the continuous case easily.†

The increased electrical power will be spread out over a wider frequency band, however. The signal has been multiplied by a pulse train, which smears its spectrum out through convolution with the transform of the pulse train,

$$S(f) = \frac{d}{f_p^2}\text{sinc}\left[\frac{fd}{f_p}\right]\text{Ш}\left[\frac{f}{f_p}\right], \tag{13.59}$$

so that the spectrum of the sampled signal is turned into a lot of little copies of itself, centered at all harmonics of f_p (including 0), whose amplitudes are governed by the sinc function envelope. To reconstruct our signal, we could just lowpass filter the pulse train, to reject all but the lobe at 0, which is identical with our original signal, no more and no less; all the extra electrical power has gone into the other lobes, and the noise bandwidth is just what it was before, so we will have gained exactly nothing by the exercise.

On the other hand, if we widen the filter bandwidth to about $f_{\max}/d$, we'll gain the d^{-2} advantage in signal power, at the price of a d^{-1} increase in the noise, even if we do nothing else. We can do even better, by time gating the noise as well, by using the same pulse waveform to gate the received signal. This of course does nothing to the signal but reduces the noise power by a factor‡ of d, so that our electrical SNR improvement is of order d^{-2}. Note that all this depends on our being able to pack the same number of photons into a shorter time. If the peak optical power doesn't grow as we shrink the duty cycle, we gain nothing—we've made a perfectly ordinary sampled-data system.

13.8.11 Pulsed Measurements and Shot Noise

A pulsed measurement concentrates all the signal power in a small time slice, where it can dominate the Johnson noise in the difficult photocurrent range from 1 pA to 1 μA.

†What happens to the spectrum as $d \to 1$?
‡Assuming the noise is white.

That's the only noise benefit, though, and you trade a lot away to get it.[†] Pulsed lasers are dramatically noisier than CW ones, have poorer spatial behavior, usually exhibit serious timing jitter, and are generally hard to use. Photodetector nonlinearity depends strongly on the peak photocurrent, so it gets far worse at short duty cycles. There is lots of spurious crud flying around at f_{rr} and its harmonics, from power supplies, Q switch drivers, and pulse gating electronics. Shrinking the duty cycle makes all these things worse rather than better.

Furthermore, shrinking the duty cycle doesn't help shot noise at all. There's no shot noise between pulses, of course, but if you're going to collect 10^{14} photoelectrons in one second, you're going to have 10^7 electrons rms noise in that time, whether the duty cycle is 1 or 10^{-6}. Thus as we improve the rest of the measurement, the need for pulses tends to go away; unless we're really stuck for photons, or there is some absolutely unavoidable noise source, for example, the shot noise of the self-luminosity of the sample,[‡] or you need the tall narrow pulses for another reason, for example, nonlinear optics or stroboscopy, you're a lot better off concentrating on getting to the shot noise than on fighting the annoyances of pulsed measurements.

13.9 SIGNAL DETECTION

13.9.1 Phase-Sensitive Detectors

From a signal processing theory point of view, a phase-sensitive detector is a mixer stage whose IF is at zero hertz. There are four common kinds of phase detectors that accept analog inputs: switching ones built from CMOS analog gates, track/hold amps, Gilbert cell multipliers, and diode bridge mixers. With a sinusoidal input, the amplitude and phase changes in the RF signal are converted to a baseband signal,

$$V_{\mathrm{IF}} \approx A_{\mathrm{RF}} \cos\theta. \tag{13.60}$$

Because phase detectors preserve phase information, filtering after the phase detector has nearly the same effect as doing it before; the only difference is the sideband folding, where USB and LSB signals produce the same output, so you can't separate them with a single mixer. This is in decided contrast to an ordinary rectifying detector, which will merrily intermodulate all noise components together, so that even distant noise will wind up landing on top of your desired signal. Even the sideband folding problem can be eliminated by using an SSB, making baseband filtering precisely equivalent to IF filtering.[§] This is a very useful fact, because you can make a 10 Hz wide filter at baseband with a 150 kΩ resistor and a 0.1 μF capacitor, which is a lot easier than doing the same trick at 10.7 MHz.

[†]It isn't the same as an ordinary AC measurement, because the signal power gets spewed out over a huge number of harmonics of the rep rate. We don't even get the benefit of getting out of the $1/f$ noise, because only the DC term is immune to timing jitter and other nonideal behavior.

[‡]Even if the sample is strongly self-luminous, a correctly designed CW measurement will usually win; the ISICL sensor of Chapter 1 needs only 50 photons to detect a particle in a plasma that's too bright to look at.

[§]Of course, if your sideband folding happens at optical frequencies (i.e., at the optical detector), you're dead. This is the origin of the 3 dB noise advantage of homodyne (IF at DC) measurements over heterodyne ones in optics. It's only the low frequency noise that tends to make homodyne measurements worse.

13.9.2 AM Detectors

There are two basic kinds of AM detectors: rectifying and synchronous. The classical AM detector is a 1N21 UHF diode wired up as a half-wave rectifier with an RC filter after it (nowadays you'd use a Schottky or back diode)—what could be simpler?. Unfortunately, this leads to horrible nonlinearity due to the soft forward turn-on behavior of the diode. The linear operating range of a diode detector is thus poor, about 20 dB, although there are circuit hacks that can be applied, as we'll see in Chapter 15.

Synchronous AM detectors are just phase detectors operated at $\theta = 0$ or π. Because of their excellent linearity, these make the best AM detectors, but of course there's the problem of where to get the LO signal. If you know its frequency and phase a priori (because you're generating it elsewhere), everything is easy—use the transmit signal as the LO, or put a very narrow filter on the input signal and use that. If you're not that lucky, you can fake it by putting the raw input signal through a comparator or limiter to make the LO—a pseudosynchronous detector. You have to worry about phase shifts and AM–PM conversion. A poor man's version of this that works only at low frequency is the (badly misnamed) perfect rectifier, which we'll encounter in Chapter 14. A common malady of these pseudosynchronous detectors is that they cannot distinguish signal from noise; the additive noise makes the LO phase noisier and noisier as the signal level drops, so that the low-SNR behavior of pseudosynchronous detectors is similar to that of rectifying ones.

A more complicated approach that yields excellent performance is to use an I/Q mixer to construct the envelope of the analytic signal directly, using the trick from Chapter 15. AM detectors are tougher to get right than you'd think. Design yours carefully.

13.9.3 PLL Detectors

Another approach, which will separate the signal from the noise a great deal better, is to phase lock to the signal (see below), and use the PLL output as the LO. This approach is probably the best if you anticipate working at low SNR a lot and want to be able to narrow the noise bandwidth after detection, as in a lock-in amplifier. This is not something for nothing; the PLL can be narrowband, so the SNR will be much better in the PLL's bandwidth than it is in the IF bandwidth. A corollary is that you can't usefully make the measurement bandwidth narrower than the PLL bandwidth. A minor drawback is that you need two phase detectors driven in quadrature, since one will be held at null by the operation of the PLL, and you need an in-phase LO for the AM detector. The walking ring counter approach is a good match here, since your LO is usually low enough that garden variety logic ICs can work at $4f_{\mathrm{LO}}$. In a real instrument you'll need the second PD for lock detection anyway (see Section 15.6.4).

13.9.4 FM/PM Detectors

FM and PM are indistinguishable if f_m is constant, but they're very different for wide modulation frequency bandwidths. They are generated differently, and you detect them differently too. For small modulation indexes $m \lesssim 0.1$ rad, you can just shove the signal into a phase detector operating near null, and use the IF output as a direct measure of the phase deviation (of course, you have to know the average phase in advance). Although the null is usually pretty stable with signal amplitude, the scale factor will change linearly, so that there is a danger of severe AM–PM conversion away from null. This problem is typically fixed by putting a limiter ahead of the phase detector.

Where to get the LO signal for an FM detector is a problem just as in the AM case. The usual way is to use an *LC* phase shift network whose phase varies with frequency (see Chapter 14). For wideband applications, the delay-line discriminator of Section 15.5.12 is useful. This is nothing but a mixer, whose LO is just the (limited) RF signal passed through a delay line made of a piece of coax cable cut to length. Its output voltage is theoretically proportional to $\cos(2\pi f\tau)$, where τ is the cable delay, but because of the nonsinusoidal characteristic of the mixer, it will actually be flattened a little on top and have a slight DC offset.

Phase demodulators are always limited to $|\Delta\phi| < \pi$, and those that operate on sine waves are limited by the nonlinearity of their output characteristic, $V(\phi) \approx V_0\cos(\phi)$. The linear range can be extended to $\pm\pi/2$ by using square waves fed into a CMOS exclusive-OR gate, such as a 74HC86.[†] After averaging the output with an *RC* filter, the result is a triangular function of the relative phase, so almost the full π range is available (the finite rise time of the edges of the output waveform make it very nonlinear if you go too close to 0 or π). Edge-sensitive phase detectors can accommodate a range of nearly 2π, also limited by rise time effects. Linearity requires flat-topped signal waveforms. They need not be square, though the steeper the edge, the more of the phase range we can use without running into edge effects. Thus we can get the same linear slope with sine waves by overdriving the inputs of a Gilbert cell multiplier, because overdriving a differential pair produces a roughly trapezoidal signal waveform with a very flat top (assuming that the collector load resistors are small enough that the transistors cut off before they saturate).

For ranges beyond 2π, it is possible to use frequency dividers in both the signal and comparison arms, so that the phase range can be as large as desired, at the expense of time delay[‡] and reduced sensitivity. A more subtle problem with dividers is the phase ambiguity, as we saw in Section 13.5.2. When the power is turned on, a divide-by-N counter can wake up in any one of N states, so the phase detector output can be at any of several places in its range, independent of the real phase relationship between the input signals. Resetting the counters externally can bring the phase detector output to a known state, but this will reflect the true phase shift modulo 2π, which does not fix the phase ambiguity. This problem remains difficult unless it is feasible to bring the two true phases into a known relationship before resetting the counters.

Besides limiters, another approach to eliminating the AM–PM conversion of a mixer is to force it to run at null, by putting it in a feedback loop with a phase shifter, and using the phase shifter control voltage as the output. You can do this in analog, with an op amp doing the feedback, or digitally by successive approximation.[§] (This is a lot quieter than fringe surfing. Why?)

[†]An exclusive-OR gate performs a cross-correlation between its inputs, if the two logic states are assigned numerical values of ±1 and the result is time averaged.

[‡]Time delay in a phase detector is somewhat subtle and may be of crucial importance in the frequency compensation of phase-locked loops. A digital phase detector's output cannot change until the next clock edge, because there's no phase information in a featureless logic level. On the other hand, sine waves exhibit phase information all the time, so an analog phase detector's output can begin changing immediately. The trade-off is a bit subtler than this because of the delays in any filters required to make the nice sine wave.

[§]P. C. D. Hobbs, High performance amplitude and phase digitizers at 60 MHz. *Rev. Sci. Instrum*. **58**(8), 1518–1522 (August 1987).

13.9.5 Phase-Locked Loops

A more common type of loop involving phase detectors is the phase-locked loop (PLL). PLLs are discussed in more detail in Section 15.6, but it's worth mentioning their use in FM and PM detection. A PLL has a voltage-controlled oscillator, a phase detector, an op amp, and some sort of filter connected in a closed loop. The feedback loop forces the VCO to follow the exact phase (and *a fortiori* frequency) of the input, for modulation frequencies well within its bandwidth. Near the unity gain bandwidth of the loop, the tracking degrades, and well outside the bandwidth, it hardly tracks at all.

You build an FM demodulator from a PLL by making the loop fast and using the VCO control voltage as the output. Since the VCO voltage has a 1:1 mapping to the frequency, this is a complete solution apart from perhaps needing linearizing. Even the linearization problem is not too severe, since it's easy to measure the VCO characteristics with a counter and voltmeter—it's a 1D calibration problem. PLL FM detection is a lowpass filtering operation and is slightly nonlinear if the signal has strong phase modulation near the edges of the loop bandwidth. An important side benefit is that you usually don't need a limiter with a PLL detector, so that the noise degradation at low SNR is avoided. (See Section 14.7.16.)

A PLL PM demodulator is made by making the loop as narrow as reasonably practicable, to reduce the noise impressed onto the oscillator by the loop, and using the phase detector output. This is inherently a highpass filtering operation, although you can sometimes use time gating of the control signal to approach DC-coupled operation. The VCO itself should be intrinsically quiet, so a voltage-controlled crystal oscillator (VCXO) is often used here.

13.9.6 *I* and *Q* Detection

In an interferometer with significantly unequal arms, (13.14) shows that we can get I and Q by a sleazy trick with the unintuitive name of *modulation-generated carrier*. If we pick m to be about 2.6, we get equal amounts of fundamental and second harmonic, and together they contain about 85% of the total AC power. Using a couple of plain AM detectors, we get I and Q almost free, because of the factors of j in the terms of (13.14). This trick is very convenient with a diode laser source, which can be frequency modulated by changing its bias current. It is especially popular in fiber optic sensors.

13.9.7 Pulse Detection

Many sorts of measurements require accurate determination of the height of a pulse. Examples are sizing of aerosol particles, pulsed laser spectroscopy, and time-domain reflectometry. It is tempting to use an analog peak detector circuit, but this is almost always the wrong thing to do, unless the pulses are very long (at least several microseconds), because peak detectors are inaccurate due to overshoot and sensitivity to the exact pulse shape. Unless you have a special situation where your pulses have all different shapes and widths, and what you need is really the maximum height, use a gated integrator instead. There's more in Chapter 14.

The other approach is to sample digitally, and do the peak detection in software. Provided the data rate is not too high, this is a reasonable approach, but sampling a waveform at 14 bits, say, and a 10 MHz sampling rate is difficult and expensive compared with a gated integrator of 30 ns time constant, with a fast comparator and a few logic gates to do the gating and a slow digitizer to follow. This is especially true if the event you're waiting for is infrequent, so that the 1000 MIPS DSP wastes most of its time idle.

13.10 REDUCING INTERFERENCE AND NOISE

Once your measurement has been moved away from 0 Hz, and the optical and environmental background, drift, and interference have been reduced adequately, the principal remaining way of rejecting noise and interference is by filtering. Sometimes the measurement is at a fixed frequency, in which case everything is easy: just pick a bandpass filter with the right bandwidth and adequate skirt selectivity, and you're almost done. This can be either an actual *LCR* or active filter, or else a bandwidth narrowing instrument such as a lock-in amplifier or signal averager.

Measurements where the carrier frequency uncertainty is large compared to the measurement bandwidth are a bit more difficult. For that, you'll be forced to choose between picking one frequency at a time and throwing the rest away, as in an AM radio, or using more hardware and detecting many frequency bins simultaneously, with a filter bank (analog or digital). It is also possible to increase the signal level while leaving the noise alone, without needing more photons, via pulse techniques such as time gating.

13.10.1 Lock-In Amplifiers

Lock-in amplifiers are described in Section 14.7.17. You probably already know that they're combination narrowband tunable filters and phase-sensitive detectors, designed to have a reasonably large dynamic range. Like all narrowband filters, lock-ins are slow, slow, slow. The speed problem inevitably accompanies the SNR improvement. A system with bandwidth B takes $10/B$ seconds to settle to 12 bits, even in the best case; if the number of data points is proportional to $1/B$ as well, as is the case in continuous-scan measurements, for example, the measurement time increases as $1/B^2$. A measurement that's 20 dB too noisy can take a graduate student's lifetime to complete if the approach to SNR improvement is just to dial down the lock-in bandwidth (this is even more of a problem if you don't have any graduate students). Such a slow measurement is highly vulnerable to drifts and external disturbances, e.g. bumping the table by accident and spoiling an hour's work, or someone opening the lab door and causing a temperature transient. Remember that that 20 dB SNR improvement will cost you a factor of 100 in speed for a fixed number of data points, and a factor of 10,000 for a continuous sweep. See also Section 14.7.4.

13.10.2 Filter Banks

Lock-ins and superhets work very well provided you know your signal frequency in advance. Sometimes this is not the case, as in Doppler shift measurements. What then? One approach is the use of filter banks. The RF signal is split into N separate transmission lines, each with its own sharp bandpass filter, chosen to be a good match for the expected signal bandwidth. The filters are overlapped somewhat, so that signals don't get missed if they happen to land in the hole between two adjacent filters. This gets expensive if you need more than 4 or 5 bands.

At sufficiently low frequency, we can use a fast Fourier transform to turn N samples into N frequency bins; this is an inexpensive way to get a whole lot of channels. The overlap between bins is controlled by windowing the data (multiplying it by a function that drops off toward the ends of the data run). The filter shapes are the DFT of the window function, so they're all identical, which is very useful. There's more on this in Section 17.5.

13.10.3 Time-Gated Detection

We've talked about improving the SNR by reducing the duty cycle d while keeping the average optical power constant. In that case, we gained SNR as the reciprocal of the duty cycle by simply widening the filter bandwidth, and by d^{-2} by time-gating the receiver as well. If we are using statistics-based methods such as thresholding, the situation is a bit more subtle.

A measurement with a 1% receiver duty cycle will have a false alarm rate 100 times smaller than a continuous measurement under identical conditions. This can really make a difference in photon counting experiments, which are limited by the dark count rate. Where the noise probability density is really exponentially decreasing,† this may make only a decibel or two of difference in the threshold, but if there are lots of outliers in the amplitude statistics, the gain can be much greater. If the peak optical power goes up by d^{-1}, the threshold can be raised considerably too.

Often the undesired pulses do not arrive in a Poisson process. In an optical time-domain reflectometer (OTDR), which sends pulses down a fiber and looks for reflections from breaks and discontinuities, there are big pulses at the beginning and the end of the fiber, which we'd like to reject well enough that they don't peg some amplifier and so mask real signals nearby. Time gating is good for that too.

The time gating need not be on–off. Some measurements have a signal strength that varies after a transmit pulse in a reproducible way. Time of flight ranging and optical tomography are good examples. It would be useful to change the system gain with time, so that huge signals from nearby objects can be accommodated along with faint ones from distant objects that arrive later. Variable PMT, MCP, or APD bias voltages can do this.

13.10.4 Signal Averaging

Pulsed measurements have another advantage, which is not always realized: most of the signal energy gets moved away from DC, up to harmonics of the sampling rate. Since most systems have $1/f$ noise and DC drifts at some level, this is an unequivocal benefit unless the $1/f$ noise of the pulses themselves dominates, as it frequently will.

Even CW measurements can benefit from this, through fast scanning plus signal averaging. Say you're making a scanned measurement, such as current-tunable diode laser spectroscopy. Instead of scanning slowly and accumulating a spectrum in one scan of a few seconds' duration, scan at 1 kHz and take N (say, 2000) spectra. Because the signal will be periodic, it will now occur in 0.5 Hz wide bands around each of the harmonics of 1 kHz, and, crucially, stuff happening much below 1 kHz can only produce a baseline shift. If your spectrum needs a frequency resolution of 1000 points, use a filter of 500 kHz to 1 MHz bandwidth, because you'll need at least the first 500 harmonics.

†For example, Gaussian or Poissonian.

When these traces are added, the signal will add coherently, and the noise only in power, leading to an SNR improvement of N times. This is equivalent to the slow scan approach for white noise, but with extra low frequency noise, it's a lot more stable and accurate. Do watch out for the temporal response of your filter—try scanning more slowly at first, and look for changes in the shape of the sharp features as you increase the scan rate, or use a triangular scan rather than a sawtooth, and force the retrace to lie on top of the forward scan on the scope. There's a more sophisticated look at this in Section 17.11.5.

13.10.5 Frequency Tracking

There are some signals, such as space-probe telemetry, where the carrier frequency of a narrowband signal varies slowly over many times the signal bandwidth. The usual method is a phase-locked loop with some lock-acquisition aid such as a slow frequency sweep, as in Section 15.6.5.

13.10.6 Modulation-Mixing Measurements

Section 10.9.4 gave some examples of modulation mixing. This trick is so useful that it's worth looking at it in some detail. The goal is to put the desired signal all by itself at some frequency that is easily filtered out from the various drive signals in the system—sort of an electronic version of a dark-field or fringe-based measurement (see Section 10.7.2). The usual method is to modulate the sample at one frequency and the beam at another frequency, then detect the difference frequency, using (13.11).

For instance, you can improve the sensitivity of a bolometer by biasing it with AC instead of DC and chopping the incoming light; signals at $f_{\text{chop}} \pm f_{\text{bias}}$ have to come from the mixing of the two, and you can put them at some distance from the effects of air currents and $1/f$ noise. The desired mixing is detected using a lock-in amplifier or tuned IF filter and detector (as in a superheterodyne radio, see Example 13.6). All the discussion of Section 13.7.2 on how to choose the frequencies to avoid spurious signals applies here as well.

Other examples are the Doppler-free spectrometer of Example 1.1, where the signal was at the sum and difference of the chopping frequencies, and many kinds of pump/probe experiments such as photothermal spectroscopy, where the pump and probe beams are modulated at different frequencies.

13.11 DATA ACQUISITION AND CONTROL

13.11.1 Quantization

A perfect N-bit digitizer will produce a numerical output indicating which of 2^N equal width bins the instantaneous signal voltage fell into at a time exactly coincident with the edge of its clock signal, and ignore what it does the rest of the time. This is obviously a nonlinear operation, and time variant too. How does it fit into our linear, time-invariant network picture? Well, by now you should foresee the answer: it doesn't, but it does nearly, and the errors can be fixed up by hand. The main things to keep straight are Nyquist's theorem and the noise due to quantization. Nyquist's theorem is covered in more detail in Chapter 17, but what it says is that if our signal contains no frequency components above f, then provided we sample periodically in time at a rate of more than $2f$, we can reconstruct the input signal perfectly from the samples at a later time, by interpolating in a certain way. This gives us confidence that sampling the data does not itself degrade it.

Quantization effects are usually dealt with by treating them as additive noise. Provided our signal swing is at least several ADUs,[†] the main effect of quantization is to contribute additive white noise of RMS voltage $1/\sqrt{12}$ ADU.[‡] This is of course not necessarily true if the signal has large components of the clock signal in it, or if your ADC autoranges. Together these give us a basis for specifying the digitizer subsystem. Real A/D converters are not of course perfect.[§] Unless you plan to do extensive numerical simulations of how different A/D errors contribute to your measurement inaccuracy, mentally reduce the number of bits of a converter by looking at the differential nonlinearity specification. The DNL of a perfect converter is 0, and of a good one is 0.5 ADU or less. If the DNL is 2 ADUs for a 12 bit converter, then what you've got is a 10 bit converter with some nice extra bits that may sometimes mean something, and won't make life worse. Integral nonlinearity and offsets are usually less serious, as something else in your system will probably be nonlinear as well, so calibration is likely to be needed anyway.

It's easy to see that the quantization error can't really be independent of the input signal, but (as the Widrow reference shows) assuming the signal is at least several ADUs in size, quantization noise is spread out pretty evenly through the fundamental interval of $[0, \nu/2)$. Since the total noise power is fixed, the noise PSD goes down as the bandwidth goes up, making quantization noise relatively less important in wideband systems for a fixed resolution.

Aside: When Is the RMS Error What Matters?. So far, we've blithely assumed that the RMS error is what matters. That's generally true when we're taking many measurements where the quantization error isn't likely to be correlated from sample to sample. On the other hand, if a maker of $3\frac{1}{2}$ digit DVMs claimed that they were accurate to 1 part in $2000\sqrt{12}$, that would clearly be specsmanship of the deepest dye; we're not making repetitive measurements of uncorrelated quantities there, and the relevant number is 0.5 ADU, or maybe 1 ADU, even if the digitizer is perfect. Thought is required.

13.11.2 Choosing a Sampling Strategy

Choosing how you'll filter and sample your data is a great deal like choosing an IF, so draw a frequency plan just as we did in Section 13.7.2. If at all possible, pick your sampling clock rate so that strong spurious signals (e.g., harmonics of 50 or 60 Hz from AC-powered equipment, or harmonics of your PWM motor controller clock) don't alias down into your frequency band of interest, and filter as close to the ADC as possible. If you can't avoid these spurs, try to put them where you can filter them out digitally, and pay very close attention to your grounding and shielding (see Chapter 16).

[†] B. Widrow, A study of rough amplitude quantization by means of Nyquist sampling theory. *IRE Trans. Circuit Theory* **3**, 266—276 (1956). A. B. Sripad and D. L. Snyder, A necessary and sufficient condition for quantization errors to be uniform and white. *IEEE Trans. Acoust. Speech Signal Process* **ASSP-25**, 442–448 (1977).

[‡] This noise has a peak-to-peak value of 1 ADU and is therefore not Gaussian.

[§] ADCs and DACs are discussed in some detail in Section 14.8, so look there for unfamiliar terms.

Or perhaps you need to sample a transient accurately, without worrying too much about its exact frequency content. In that case, the sampling theorem is more of a sanity check than a useful design criterion—you'll need to sample significantly faster than you might gather from Nyquist (see Section 17.4.4).

13.11.3 Designing with ADCs

Deciding on the resolution of your ADC is not necessarily trivial, as the following two examples demonstrate.

Example 13.9: CCD Digitizer. For instance, consider a CCD digitizer, to be operated at fairly low speed, say, 150k samples/s. Full scale on a commodity camcorder CCD is about 5×10^4 electrons, and the readout noise can be as small as 20 electrons. We saw above that an ADC contributes RMS noise voltage of $1/\sqrt{12}$ ADU, so in order for the readout noise to dominate the ADC contribution, we'd need

$$N_{\text{bits}} \geq \log_2 \left(\frac{5 \times 10^4}{20\sqrt{12}} \right) \tag{13.61}$$

or 10 bits, at least. Remember that it is not enough that the digitizer noise be smaller than the rest of the noise; allowing one stage to contribute as much noise as the rest of the instrument combined is a foolish cost/benefit choice unless that one is by far the most expensive. To reduce ADC noise power to 0.25 of the shot noise, so that the total SNR is degraded no more than 1 dB by the digitizer, we'll need an 11 bit ADC, which is unavailable. A 12 bit ADC will reduce the noise degradation to 0.26 dB, which is usually good enough. If you need to do this at high speed, that ADC will run into money. (If your photons are very expensive, it may pay to go to even higher ADC resolutions—a quarter of a decibel here and there adds up.)

Example 13.10: Soda Pop Bottles. Now let's say that you're building an inspection system for cola bottling plants. You have lots of light, so there's no reason to slow the sampling rate down, and you anticipate being near half-scale most of the time. Now the photon noise is not 20 electrons, but the square root of half the bucket capacity, or about 160, a factor of 8 higher than in the previous example. That means that an 8 bit digitizer is probably enough for that application. Video rate 8 bit ADCs are cheap enough to give away in cereal boxes, so you've saved some money and probably haven't lost any performance.

13.11.4 Choosing the Resolution

Rules of thumb for choosing a digitizer resolution are often simplistic and inadequate, and this one is no exception, because it assumes that the digitizer's performance is perfect. Ordinary problems like differential and integral nonlinearity are dealt with by mentally reducing the number of bits, but that's just the beginning. Real digitizers have noise and also exhibit serious time variations caused by the signal, so that they contribute spurious signals just as mixers do. If you are using a fast digitizer, and especially if your signal contains frequency components higher than 5% of the device's maximum rated sampling frequency, make sure you find out about its performance with large AC signals.

Many ADCs come with integral T/Hs; these are called sampling ADCs in the catalogs. The integrated T/Hs are often very good and well matched to the ADC behavior. Unless you have a pressing need to use a separate T/H, go for a sampling ADC. It'll just make your life easier. (A good example is putting the multiplexer between the sampling and digitizing steps to get simultaneous sampling of several channels.) Do keep an eye on the aperture uncertainty specification—it's critical but often overlooked. There's lots on that and other ADC specifications and problems in Chapter 14.

The bottom line of all this is that getting good performance out of your digitizer is nontrivial. To get the best measurement for your money, it is very important to make an end-to-end noise and error budget, so that if you can afford to give up 6 dB in sensitivity, you start out by spending most of it reducing the cost of the most expensive part (e.g., going to a lower power laser or reducing your collection NA), and then see how much is left for the rest of the system.

13.11.5 Keep Zero On-Scale

While we're on the subject of ADC resolution, we should note that even a 20 bit ADC has no resolution whatever beyond its voltage limits. Thus it is very important to ensure that all valid signals produce an on-scale reading. Normally people don't screw up the upper limit, but in these days of single-supply analog circuitry, an amazing number will put the zero-signal input to the ADC at 0 V, and not think anything of it. This is a disaster. First of all, you have no way of distinguishing between a zero signal condition and a fault (e.g., loss of the analog power supply). Also, everything has an offset voltage, and nobody can really pull down to 0, so there will be a region near zero signal where you can't get a reliable digital output.

Make sure your signal has a small, accurately known DC offset, enough to keep it on-scale over the worst case offset, supply, and nonlinearity variations (and then add a bit more for luck). Use 1% of full scale as a starting point.

Example 13.11: Error Budget for an Extinction Turbidity Meter. Turbidity meters are used in pollution monitoring, for example, to measure the nutrient concentration of water by looking for turbidity caused by bacteria and algae growth. A simple turbidity meter might use an inexpensive 600 nm LED driven with a 10 kHz square wave as a light source, a small solar cell as a detector, and a signal processor consisting of a transimpedance amp and active bandpass filter leading up to a synchronous AM detector to convert the AC signal into a voltage. The resulting voltage is to be digitized.

We anticipate needing high sensitivity only for small turbidities, since heavily contaminated sludge is not on the menu.

This arrangement should be relatively immune to etalon fringes, due to the very low temporal coherence of the LED source. Collimating the light and forcing it to pass through an aperture on the source side and a black tube baffle on the output side of the sample cell should eliminate false readings due to multiply scattered light. There should be no problems with pickup and $1/f$ noise, since it's working at a high audio frequency. Its time response will be set by the filter bandwidths but could easily be 1 ms. Let's say it is to be 10 ms, 10–90% . An RC lowpass filter has a rise time $t_r = 0.35/f_{-3\text{dB}}$, so the bandwidth should be about 35 Hz.

We'll use a CMOS analog gate, driven by the same square wave, as a phase-sensitive detector, followed by a 35 Hz RC lowpass filter. A single-pole RC lowpass has a noise bandwidth $BW_{noise} = (\pi/2)f_c$, or 55 Hz. If the peak photocurrent detected at zero turbidity is 100 μA, then the average is $i_{DC} = 50\mu$A, and the shot noise is 4 pA, a SNR of 141 dB in 1 Hz, or 124 dB in 55 Hz, so that we need at least a 19 bit ADC to preserve the noise performance. With care, the shot noise is easily reachable in a system like this, so the choice appears to be using a 20 bit ADC (slow ones exist) or allowing ourselves to be ADC limited. In the ADC-limited case, we should vary the drive current a bit between samples, because unless the signal varies by at least a few ADUs, quantization noise may not be additive. Since the sensitivity level required is not all that high, being ADC limited may not be so bad here.

There is another way of looking at this problem, though. We'll need to normalize the detected signal. The simplest way to do this is to send part of the light through twice, perhaps by mirror-coating part of the flow cell; a better way, which will help cancel effects due to gradual fouling of the windows, is to make the flow cell wedge-shaped and send two beams through regions of different width. The two photocurrents can be divided, perhaps by using the once-dimmed beam to produce the ADC reference current for digitizing the twice-dimmed one,[†] yielding a ratiometric measurement in which zero turbidity produces a (nearly) full scale output. Alternatively, they can be subtracted, with the beam strengths adjusted to null the resulting signal at zero turbidity; a zero-based signal results. This zero-based signal can be amplified and then digitized at much lower resolution, because the ADC resolution limits the maximum signal, rather than the sensitivity. (Remember to keep zero and maximum signal slightly inside the ends of the ADC range.) Very turbid water could be dealt with by having two ranges or by digitizing the logarithm instead (see Section 15.5.9), but multiple scattering would start to dominate, so another approach would probably be required anyway.

[†]Why does this work?