

14

Electronics*

Everything should be made as simple as possible, but not simpler.

Albert Einstein

14.1 Introduction

Electronics are a key component of all modern detector systems. Although experiments and their associated electronics can take very different forms, the same basic principles of the electronic readout and optimisation of signal-to-noise ratio apply to all. This chapter gives an introduction to electronic noise, signal processing and digital electronics. Because of space limitations, this can only be a brief overview. A more detailed discussion of electronics with emphasis on semiconductor detectors is given elsewhere [1]. Tutorials on detectors, signal processing and electronics are also available on the worldwide web [2].

The purpose of front-end electronics and signal processing systems is to

- (i) acquire an electrical signal from the sensor. Typically, this is a short current pulse.
- (ii) tailor the time response of the system to optimise
 - (a) the minimum detectable signal (detect hit/no hit),
 - (b) energy measurement,
 - (c) event rate,
 - (d) time of arrival (timing measurement),

* This chapter was contributed by Helmuth Spieler, Lawrence Berkeley National Laboratory, Berkeley, California, USA.

- (e) insensitivity to sensor pulse shape,
 - (f) or some combination of the above.
- (iii) digitise the signal and store for subsequent analysis.

Generally, these properties cannot be optimised simultaneously, so compromises are necessary. In addition to these primary functions of an electronic readout system, other considerations can be equally or even more important, for example radiation resistance, low power (portable systems, large detector arrays, satellite systems), robustness, and – last, but not least – cost.

14.2 Example systems

Figure 14.1 illustrates the components and functions of a radiation detector system. The sensor converts the energy deposited by a charged particle (or photon) to an electrical signal. This can be achieved in a variety of ways. In direct detection – semiconductor detectors, wire chambers, or other types of ionisation chambers – energy is deposited in an absorber and converted into charge pairs, whose number is proportional to the absorbed energy. The signal charge can be quite small, in semiconductor sensors about 50 aC ($5 \cdot 10^{-17}$ C) for 1 keV X rays and 4 fC ($4 \cdot 10^{-15}$ C) in a typical high-energy tracking detector, so the sensor signal must be amplified. The magnitude of the sensor signal is subject to statistical fluctuations, and electronic noise further ‘smears’ the signal. These fluctuations will be discussed below, but at this point we note that the sensor and preamplifier must be designed carefully to minimise electronic noise. A critical parameter is the total capacitance in parallel with the input, i.e. the sensor capacitance and input capacitance of the amplifier. The *signal-to-noise ratio* increases with decreasing capacitance. The contribution of electronic noise also relies critically on the next stage, the pulse shaper, which determines the bandwidth of the system and hence the




Fig. 14.1. Basic detector functions: radiation is absorbed in the sensor and converted into an electrical signal. This low-level signal is integrated in a preamplifier, fed to a pulse shaper, and then digitised for subsequent storage and analysis.




Fig. 14.2. In a scintillation detector absorbed energy is converted into visible light. The scintillation photons are commonly detected by a photomultiplier, which can provide sufficient gain to directly drive a threshold discriminator.

overall electronic noise contribution. The shaper also limits the duration of the pulse, which sets the maximum signal rate that can be accommodated. The shaper feeds an *analogue-to-digital converter* (ADC), which converts the magnitude of the analogue signal into a bit pattern suitable for subsequent digital storage and processing.

A scintillation detector (Fig. 14.2) utilises indirect detection, where the absorbed energy is first converted into visible light. The number of scintillation photons is proportional to the absorbed energy. The scintillation light is detected by a photomultiplier (PMT), consisting of a photocathode and an electron multiplier. Photons absorbed in the photocathode release electrons, whose number is proportional to the number of incident scintillation photons. At this point energy absorbed in the scintillator has been converted into an electrical signal whose charge is proportional to the energy. Increased in magnitude by the electron multiplier, the signal at the PMT output is a current pulse. Integrated over time this pulse contains the signal charge, which is proportional to the absorbed energy. Figure 14.2 shows the PMT output pulse fed directly to a threshold discriminator, which fires when the signal exceeds a predetermined threshold, as in a counting or timing measurement. The electron multiplier can provide sufficient gain, so no preamplifier is necessary. This is a typical arrangement used with fast plastic scintillators. In an energy measurement, for example using a NaI(Tl) scintillator, the signal would feed a pulse shaper and ADC, as shown in Fig. 14.1.

If the pulse shape does not change with signal charge, the peak amplitude – the pulse height – is a measure of the signal charge, so this measurement is called *pulse-height analysis*. The pulse shaper can serve multiple functions, which are discussed below. One is to tailor the pulse

shape to the ADC. Since the ADC requires a finite time to acquire the signal, the input pulse may not be too short and it should have a gradually rounded peak. In scintillation detector systems the shaper is frequently an integrator and implemented as the first stage of the ADC, so it is invisible to the casual observer. Then the system appears very simple, as the PMT output is plugged directly into a charge-sensing ADC.

A detector array combines the sensor and the analogue signal-processing circuitry together with a readout system. Figure 14.3 shows the circuit blocks in a representative *readout integrated circuit* (IC). Individual sensor electrodes connect to parallel channels of analogue signal-processing circuitry. Data are stored in an analogue pipeline[‡] pending a readout command. Variable write and read (R/W) pointers are used to allow simultaneous read and write. The signal in the time slot of interest is digitised, compared with a digital threshold, and read out. Circuitry is included to generate test pulses that are injected into the input to simulate a detector signal. This is a very useful feature in setting up the system and is also a key function in chip testing prior to assembly. Analogue control levels are set by *digital-to-analogue converters* (DACs). Multiple ICs




Fig. 14.3. Circuit blocks in a representative readout IC. The analogue processing chain is shown at the top. Control is passed from chip to chip by token passing.[†]

[‡] A form of parallel processing enabling streams of instructions or data to be executed or handled concurrently.

[†] In the technique of token passing only that part of an electronic circuitry is allowed to communicate that has a 'token'. When the information of that part is read out, the token is passed on to the next element which can then communicate.




Fig. 14.4. Multiple ICs are ganged to read out a strip detector. The rightmost chip IC1 is the master. A command on the control bus initiates the readout. When IC1 has written all of its data, it passes the token to IC2. When IC2 has finished, it passes the token to IC3, which in turn returns the token to the master IC1.

are connected to a common control and data output bus, as shown in Fig. 14.4. Each IC is assigned a unique address, which is used in issuing control commands for setup and *in situ* testing. Sequential readout is controlled by token passing. IC1 is the master, whose readout is initiated by a command (trigger) on the control bus. When it has finished writing data, it passes the token to IC2, which in turn passes the token to IC3. When the last chip has completed its readout, the token is returned to the master IC, which is then ready for the next cycle. The readout bit stream begins with a header that uniquely identifies a new frame. Data from individual ICs are labelled with a chip identifier and channel identifiers. Many variations on this scheme are possible. As shown, the readout is event oriented, i.e. all hits occurring within an externally set exposure time (e.g. time slice in the analogue buffer in Fig. 14.3) are read out together. For a concise discussion of data acquisition systems see [3].

In colliding beam experiments only a small fraction of beam crossings yields interesting events. The time required to assess whether an event is potentially interesting is typically of order microseconds, so hits from multiple beam crossings must be stored on-chip, identified by beam crossing or time stamp. Upon receipt of a trigger only the interesting data are digitised and read out ('sparsification').[§] This allows use of a digitiser that is slower than the collision rate. It is also possible to read out analogue signals and digitise them externally. Then the output stream is a

[§] Sparsification is a tool where only those elements of an electronic system that have relevant information are read out, thus providing a technique for speeding up the readout.




Fig. 14.5. Closeup of ICs mounted on a hybrid utilising a flexible polyimide substrate [4]. The three large rectangular objects are the readout chips, each with 128 channels. Within the chips the structure of the different circuit blocks is clearly visible, e.g. the 128 parallel analogue processing chains at the upper end. The 128 inputs at the upper edge are wire bonded to a pitch adapter to make the transition from the approximately $50\ \mu\text{m}$ pitch of the readout to the $80\ \mu\text{m}$ pitch of the silicon strip detector. The power, data and control lines are wire bonded at the lower edge. Bypass capacitors (the small rectangular objects with shiny contacts at the top and bottom) positioned between the readout chips connect to bond pads on the chip edges to reduce the series inductance to the on-chip circuitry. The ground plane is patterned as a diamond grid to reduce material. (Photograph courtesy of A. Ciocio.)

sequence of digital headers and analogue pulses. An alternative scheme only records the presence of a hit. The output of a threshold comparator signifies the presence of a signal and is recorded in a digital pipeline that retains the crossing number.

Figure 14.5 shows a closeup of ICs mounted on a hybrid circuit using a flexible polyimide substrate [4]. The wire bonds connecting the IC to the hybrid are clearly visible. Channels on the IC are laid out on a pitch of about $50\ \mu\text{m}$ and pitch adapters fan out to match the $80\ \mu\text{m}$ pitch of the silicon strip detector. The space between chips accommodates bypass capacitors and connections for control busses carrying signals from chip to chip.

14.3 Detection limits

The *minimum detectable signal* and the precision of the amplitude measurement are limited by fluctuations. The signal formed in the sensor fluctuates, even for a fixed energy absorption. In addition, electronic

noise introduces baseline fluctuations, which are superimposed on the signal and alter the peak amplitude. Figure 14.6a shows a typical noise waveform. Both the amplitude and time distributions are random. When superimposed on a signal, the noise alters both the amplitude and time dependence, as shown in Fig. 14.6b. As can be seen, the noise level determines the minimum signal whose presence can be discerned.

In an optimised system, the time scale of the fluctuations is comparable to that of the signal, so the peak amplitude fluctuates randomly above and below the average value. This is illustrated in Fig. 14.7, which shows the same signal viewed at four different times. The fluctuations in peak




Fig. 14.6. Waveforms of random noise (a) and signal + noise (b), where the peak signal is equal to the rms noise level ($S/N = 1$). The noiseless signal is shown for comparison.




Fig. 14.7. Signal plus noise for four different pulse examples, shown for a signal-to-noise ratio of about 20. The noiseless signal is superimposed for comparison.

amplitude are obvious, but the effect of noise on timing measurements can also be seen. If the timing signal is derived from a threshold discriminator, where the output fires when the signal crosses a fixed threshold, amplitude fluctuations in the leading edge translate into time shifts. If one derives the time of arrival from a centroid analysis, the timing signal also shifts (compare the top and bottom right figures). From this one sees that signal-to-noise ratio is important for all measurements – sensing the presence of a signal or the measurement of energy, timing or position.

14.4 Acquiring the sensor signal

The sensor signal is usually a short current pulse $i_s(t)$. Typical durations vary widely, from 100 ps for thin Si sensors to several μs for inorganic scintillators. However, the physical quantity of interest is often the deposited energy, so one has to integrate over the current pulse,

$$E \propto Q_s = \int i_s(t) dt . \quad (14.1)$$

This integration can be performed at any stage of a linear system, so one can

- (i) integrate on the sensor capacitance,
- (ii) use an integrating preamplifier ('charge-sensitive' amplifier),
- (iii) amplify the current pulse and use an integrating ADC ('charge-sensing' ADC),
- (iv) rapidly sample and digitise the current pulse and integrate numerically.

In large systems the first three options tend to be most efficient.

14.4.1 Signal integration

Figure 14.8 illustrates signal formation in an ionisation chamber connected to an amplifier with a very high input resistance. The ionisation chamber volume could be filled with gas, liquid or a solid, as in a silicon sensor. As mobile charge carriers move towards their respective electrodes, they change the induced charge on the sensor electrodes, which form a capacitor C_d . If the amplifier has a very small input resistance R_i , the time constant $\tau = R_i(C_d + C_i)$ for discharging the sensor is small, and the amplifier will sense the signal current (C_i is the dynamic input capacitance). However, if the input time constant is large compared to the duration of the current




Fig. 14.8. Charge collection and signal integration in an ionisation chamber.

pulse, the current pulse will be integrated on the capacitance and the resulting voltage at the amplifier input is

$$V_i = \frac{Q_s}{C_d + C_i} \tag{14.2}$$

The magnitude of the signal is dependent on the sensor capacitance. In a system with varying sensor capacitances, a Si tracker with varying strip lengths, for example, or a partially depleted semiconductor sensor, where the capacitance varies with the applied bias voltage, one would have to deal with additional calibrations. However, the charge-sensitive amplifiers, widely used for detector readout, would overcome this problem by providing an output signal height practically independent of the input capacity.

Figure 14.9 shows the principle of a *feedback amplifier* that performs integration. It consists of an inverting amplifier with voltage gain $-A$ and a feedback capacitor C_f connected from the output to the input. To simplify the calculation, let the amplifier have an infinite input impedance, so no current flows into the amplifier input. If an input signal produces a voltage v_i at the amplifier input, the voltage at the amplifier output is $-Av_i$. Thus, the voltage difference across the feedback capacitor is $v_f = (A+1)v_i$ and the charge deposited on C_f is $Q_f = C_f v_f = C_f(A+1)v_i$. Since no current can flow into the amplifier, all of the signal current must charge




Fig. 14.9. Principle of a *charge-sensitive amplifier*.

up the feedback capacitance, so $Q_f = Q_i$. The amplifier input appears as a 'dynamic' input capacitance

$$C_i = \frac{Q_i}{v_i} = C_f(A + 1) . \quad (14.3)$$

The voltage output per unit input charge is

$$A_Q = \frac{dv_o}{dQ_i} = \frac{Av_i}{C_i v_i} = \frac{A}{C_i} = \frac{A}{A + 1} \cdot \frac{1}{C_f} \approx \frac{1}{C_f} \quad (A \gg 1) , \quad (14.4)$$

so the charge gain is determined by a well-controlled component, the feedback capacitor.

The signal charge Q_s will be distributed between the sensor capacitance C_d and the dynamic input capacitance C_i . The ratio of measured charge to signal charge is

$$\frac{Q_i}{Q_s} = \frac{Q_i}{Q_d + Q_i} = \frac{C_i}{C_d + C_i} = \frac{1}{1 + C_d/C_i} , \quad (14.5)$$

so the dynamic input capacitance must be large compared to the sensor capacitance.

Another very useful byproduct of the integrating amplifier is the ease of charge calibration. By adding a test capacitor as shown in Fig. 14.10, a voltage step injects a well-defined charge into the input node. If the dynamic input capacitance C_i is much larger than the test capacitance C_T , the voltage step at the test input will be applied nearly completely across the test capacitance C_T , thus injecting a charge $C_T \Delta V$ into the input.

The preceding discussion assumed that the amplifiers are infinitely fast, so they respond instantaneously to the applied signal. In reality amplifiers

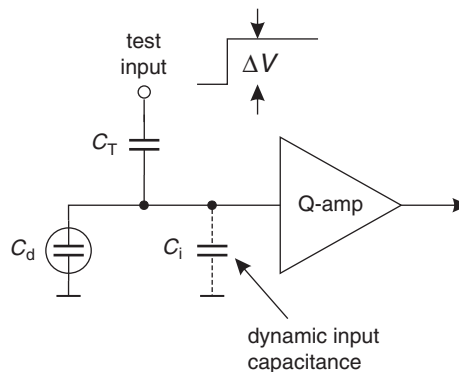


Fig. 14.10. Adding a test input to a charge-sensitive amplifier provides a simple means of absolute charge calibration.

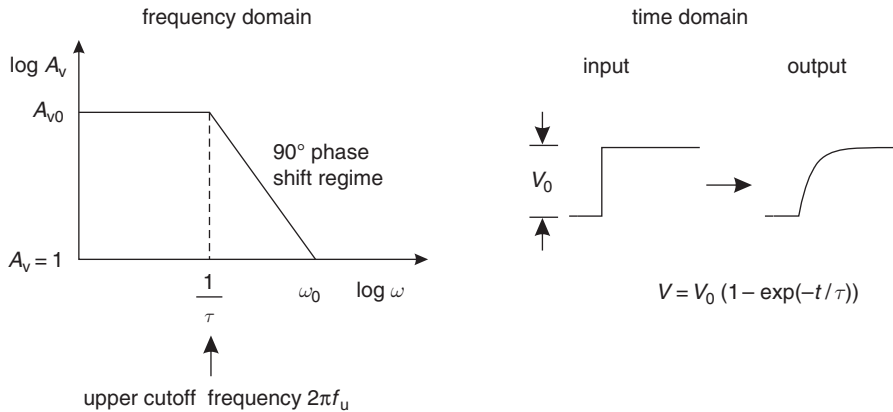


Fig. 14.11. The time constant of an amplifier τ affects both the frequency and the time response. The amplifier's cutoff frequency is $\omega_u = 1/\tau = 2\pi f_u$. Both the time and frequency response are fully equivalent representations.

have a limited bandwidth, which translates into a time response. If a voltage step is applied to the input of the amplifier, the output does not respond instantaneously, as internal capacitances must first charge up. This is shown in Fig. 14.11. In a simple amplifier the time response is determined by a single time constant τ , corresponding to a cutoff (corner) frequency $\omega_u = 1/\tau = 2\pi f_u$. In the frequency domain the gain of a simple single-stage amplifier is constant up to the cutoff frequency f_u and then decreases inversely proportional to frequency with an additional phase shift of 90° . In this regime the product of gain and bandwidth is constant, so extrapolation to unity gain yields the gain–bandwidth product $\omega_0 = A_{v0} \cdot \omega_u$. In practice, amplifiers utilise multiple stages, all of which contribute to the frequency response. However, for use as a feedback amplifier, only one time constant should dominate, so the other stages must have much higher cutoff frequencies. Then the overall amplifier response is as shown in Fig. 14.11, except that at high frequencies additional corner frequencies appear.

The frequency-dependent gain and phase affect the input impedance of a charge-sensitive amplifier. At low frequencies – where the gain is constant – the input appears capacitive, as shown in Eq. (14.3). At high frequencies where the additional 90° phase shift applies, the combination of amplifier phase shift and the 90° phase difference between voltage and current in the feedback capacitor leads to a resistive input impedance

$$Z_i = \frac{1}{\omega_0 C_f} \equiv R_i . \quad (14.6)$$

Thus, at low frequencies $f \ll f_u$ the input of a charge-sensitive amplifier appears capacitive, whereas at high frequencies $f \gg f_u$ it appears resistive.

Suitable amplifiers invariably have corner frequencies well below the frequencies of interest for radiation detectors, so the input impedance is resistive. This allows a simple calculation of the time response. The sensor capacitance is discharged by the resistive input impedance of the feedback amplifier with the time constant

$$\tau_i = R_i C_d = \frac{1}{\omega_0 C_f} \cdot C_d . \quad (14.7)$$

From this we see that the rise time of the charge-sensitive amplifier increases with sensor capacitance. The amplifier response can be slower than the duration of the current pulse from the sensor, as charge is initially stored on the detector capacitance, but the amplifier should respond faster than the peaking time of the subsequent pulse shaper. The feedback capacitance should be much smaller than the sensor capacitance. If $C_f = C_d/100$, the amplifier's gain-bandwidth product must be $100/\tau_i$, so for a rise time constant of 10 ns the gain-bandwidth product must be $\omega = 10^{10} \text{ s}^{-1} \hat{=} 1.6 \text{ GHz}$. The same result can be obtained using conventional operational amplifier feedback theory.

The mechanism of reducing the input impedance through shunt feedback leads to the concept of the *virtual ground*. If the gain is infinite, the input impedance is zero. Although very high gains (of order 10^5 – 10^6) are achievable in the kHz range, at the frequencies relevant for detector signals the gain is much smaller. The input impedance of typical charge-sensitive amplifiers in strip-detector systems is of order k Ω . Fast amplifiers designed to optimise power dissipation achieve input impedances of 100–500 Ω [5]. None of these qualifies as a 'virtual ground', so this concept should be applied with caution.

Apart from determining the signal rise time, the input impedance is critical in position-sensitive detectors. Figure 14.12 illustrates a silicon-strip sensor read out by a bank of amplifiers. Each strip electrode has a capacitance C_b to the backplane and a fringing capacitance C_{ss} to the neighbouring strips. If the amplifier has an infinite input impedance, charge induced on one strip will capacitively couple to the neighbours and the signal will be distributed over many strips (determined by C_{ss}/C_b). If, on the other hand, the input impedance of the amplifier is low compared to the interstrip impedance, practically all of the charge will flow into the amplifier, as current seeks the path of least impedance, and the neighbours will show only a small signal.




Fig. 14.12. To preserve the position resolution of strip detectors the readout amplifiers must have a low input impedance to prevent spreading of signal charge to the neighbouring electrodes.

14.5 Signal processing

As noted in the introduction, one of the purposes of signal processing is to improve the signal-to-noise ratio by tailoring the spectral distributions of the signal and the electronic noise. However, for many detectors electronic noise does not determine the resolution. This is especially true for counters using photomultipliers for signal readout and amplification. For example, in a NaI(Tl) scintillation detector measuring 511 keV gamma rays, say in a positron-emission tomography system, 25 000 scintillation photons are produced. Because of reflective losses, about 15 000 reach the photocathode. This translates to about 3000 electrons reaching the first dynode. The gain of the electron multiplier will yield about $3 \cdot 10^9$ electrons at the anode. The statistical spread of the signal is determined by the smallest number of electrons in the chain, i.e. the 3000 electrons reaching the first dynode, so the resolution is $\sigma_E/E = 1/\sqrt{3000} \approx 2\%$, which at the anode corresponds to about $3 \cdot 10^9 \cdot 2\% \approx 6 \cdot 10^7$ electrons. This is much larger than the electronic noise in any reasonably designed system. This situation is illustrated in Fig. 14.13a. In this case, signal acquisition and count-rate capability may be the prime objectives of the pulse processing system. Figure 14.13b shows the situation for high-resolution sensors with small signals. Examples are semiconductor detectors, photodiodes or ionisation chambers. In this case, low noise is critical. Baseline fluctuations




Fig. 14.13. Signal and baseline fluctuations add in quadrature. For large signal variance (a) as in scintillation detectors or proportional chambers the baseline noise is usually negligible, whereas for small signal variance as in semiconductor detectors or liquid-Ar ionisation chambers, baseline noise is critical.

can have many origins, external interference, artifacts due to imperfect electronics, etc., but the fundamental limit is electronic noise.

14.6 Electronic noise

Consider a current flowing through a sample bounded by two electrodes, i.e. n electrons moving with velocity v . The induced current depends on the spacing l between the electrodes (following ‘Ramo’s theorem’ [1, 6]), so

$$i = \frac{nev}{l} . \quad (14.8)$$

The fluctuation of this current is given by the total differential

$$\langle di \rangle^2 = \left(\frac{ne}{l} \langle dv \rangle \right)^2 + \left(\frac{ev}{l} \langle dn \rangle \right)^2 , \quad (14.9)$$

where the two terms add in quadrature, as they are statistically uncorrelated. From this one sees that two mechanisms contribute to the total noise, velocity and number fluctuations.

Velocity fluctuations originate from thermal motion. Superimposed on the average drift velocity are random velocity fluctuations due to thermal excitations. This *thermal noise* is described by the long-wavelength limit of Planck’s blackbody spectrum, where the spectral density, i.e. the power per unit bandwidth, is constant (*white noise*).

Number fluctuations occur in many circumstances. One source is carrier flow that is limited by emission over a potential barrier. Examples are thermionic emission or current flow in a semiconductor diode. The probability of a carrier crossing the barrier is independent of any other carrier being emitted, so the individual emissions are random and not correlated. This is called *shot noise*, which also has a ‘white’ spectrum. Another source of number fluctuations is carrier trapping. Imperfections in a crystal lattice or impurities in gases can trap charge carriers and release them after a characteristic lifetime. This leads to a frequency-dependent spectrum $dP_n/df = 1/f^\alpha$, where α is typically in the range of 0.5–2. Simple derivations of the spectral noise densities are given in [1].

The amplitude distribution of the noise is Gaussian, so superimposing a constant amplitude signal on a noisy baseline will yield a Gaussian amplitude distribution whose width equals the noise level (Fig. 14.14). Injecting a pulser signal and measuring the width of the amplitude distribution yields the noise level.

14.6.1 Thermal (Johnson) noise

The most common example of noise due to velocity fluctuations is the noise of resistors. The spectral noise power density versus frequency is

$$\frac{dP_n}{df} = 4kT \quad , \quad (14.10)$$




Fig. 14.14. Repetitive measurements of the signal charge yield a Gaussian distribution whose standard deviation equals the rms noise level Q_n . Often the width is expressed as the full width at half maximum (FWHM), which is 2.35 times the standard deviation.

where k is the Boltzmann constant and T the absolute temperature. Since the power in a resistance R can be expressed through either voltage or current,

$$P = \frac{V^2}{R} = I^2 R \quad , \quad (14.11)$$

the spectral voltage and current noise densities are

$$\frac{dV_n^2}{df} \equiv e_n^2 = 4kTR \quad \text{and} \quad \frac{dI_n^2}{df} \equiv i_n^2 = \frac{4kT}{R} \quad . \quad (14.12)$$

The total noise is obtained by integrating over the relevant frequency range of the system, the bandwidth, so the total noise voltage at the output of an amplifier with a frequency-dependent gain $A(f)$ is

$$v_{on}^2 = \int_0^\infty e_n^2 A^2(f) df \quad . \quad (14.13)$$

Since the spectral noise components are non-correlated, one must integrate over the noise power, i.e. the voltage squared. The total noise increases with bandwidth. Since small bandwidth corresponds to large rise times, increasing the speed of a pulse measurement system will increase the noise.

14.6.2 Shot noise

The spectral density of shot noise is proportional to the average current I ,

$$i_n^2 = 2eI \quad , \quad (14.14)$$

where e is the electron charge. Note that the criterion for shot noise is that carriers are injected independently of one another, as in thermionic or semiconductor diodes. Current flowing through an ohmic conductor does not carry shot noise, since the fields set up by any local fluctuation in charge density can easily draw in additional carriers to equalise the disturbance.

14.7 Signal-to-noise ratio versus sensor capacitance

The basic noise sources manifest themselves as either voltage or current fluctuations. However, the desired signal is a charge, so to allow a comparison we must express the signal as a voltage or current. This was illustrated for an ionisation chamber in Fig. 14.8. As was noted, when the input time constant $R_i(C_d + C_i)$ is large compared to the duration of the

sensor current pulse, the signal charge is integrated on the input capacitance, yielding the signal voltage $V_s = Q_s/(C_d + C_i)$. Assume that the amplifier has an input noise voltage V_n . Then the signal-to-noise ratio is

$$\frac{V_s}{V_n} = \frac{Q_s}{V_n(C_d + C_i)} \quad (14.15)$$

This is a very important result – the signal-to-noise ratio for a given signal charge is inversely proportional to the total capacitance at the input node. Note that zero input capacitance does not yield an infinite signal-to-noise ratio. As shown in [1], this relationship only holds when the input time constant is about ten times greater than the sensor current pulse width. The dependence of signal-to-noise ratio on capacitance is a general feature that is independent of amplifier type. Since feedback cannot improve signal-to-noise ratio, Eq. (14.15) holds for charge-sensitive amplifiers, although in that configuration the charge signal is constant, but the noise increases with total input capacitance (see [1]). In the noise analysis the feedback capacitance adds to the total input capacitance (the passive capacitance, not the dynamic input capacitance), so C_f should be kept small.

14.8 Pulse shaping

Pulse shaping has two conflicting objectives. The first is to limit the bandwidth to match the measurement time. Too large a bandwidth will increase the noise without increasing the signal. Typically, the pulse shaper transforms a narrow sensor pulse into a broader pulse with a gradually rounded maximum at the peaking time. This is illustrated in Fig. 14.15. The signal amplitude is measured at the peaking time T_P .

The second objective is to constrain the pulse width so that successive signal pulses can be measured without overlap (pile-up), as illustrated in Fig. 14.16. Reducing the pulse duration increases the allowable signal rate, but at the expense of electronic noise.




Fig. 14.15. In energy measurements a pulse processor typically transforms a short sensor current pulse to a broader pulse with a peaking time T_P .




Fig. 14.16. Amplitude pile-up occurs when two pulses overlap (left). Reducing the shaping time allows the first pulse to return to the baseline before the second pulse arrives.

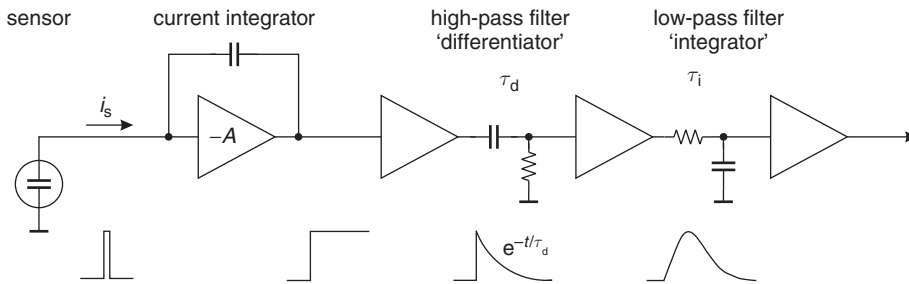


Fig. 14.17. Components of a pulse shaping system. The signal current from the sensor is integrated to form a step impulse with a long decay. A subsequent high-pass filter ('differentiator') limits the pulse width and the low-pass filter ('integrator') increases the rise time to form a pulse with a smooth cusp.

In designing the shaper it is necessary to balance these conflicting goals. Usually, many different considerations lead to a 'non-textbook' compromise; 'optimum shaping' depends on the application.

A simple shaper is shown in Fig. 14.17. A high-pass filter sets the duration of the pulse by introducing a decay time constant τ_d . Next a low-pass filter with a time constant τ_i increases the rise time to limit the noise bandwidth. The high-pass filter is often referred to as a 'differentiator', since for short pulses it forms the derivative. Correspondingly, the low-pass filter is called an 'integrator'. Since the high-pass filter is implemented with a CR section and the low-pass with an RC , this shaper is referred to as a $CR-RC$ shaper. Although pulse shapers are often more sophisticated and complicated, the $CR-RC$ shaper contains the essential features of all pulse shapers, a lower frequency bound and an upper frequency bound.

After peaking the output of a simple $CR-RC$ shaper returns to baseline rather slowly. The pulse can be made more symmetrical, allowing higher signal rates for the same peaking time. Very sophisticated circuits have




Fig. 14.18. Pulse shape versus number of integrators in a $CR-nRC$ shaper. The time constants are scaled with the number of integrators to maintain the peaking time.

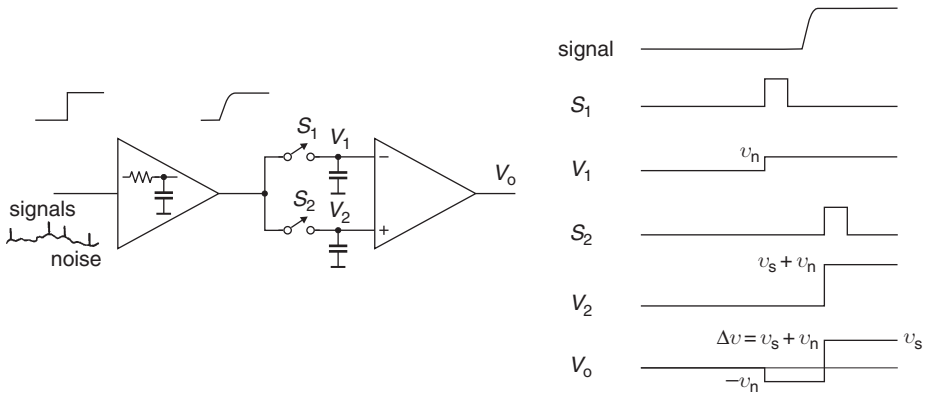


Fig. 14.19. Principle of a shaper using correlated double sampling. v_s and v_n are the signal and noise voltages.

been developed towards this goal, but a conceptually simple way is to use multiple integrators, as illustrated in Fig. 14.18. The integration and differentiation time constants are scaled to maintain the peaking time. Note that the peaking time is a key design parameter, as it dominates the noise bandwidth and must also accommodate the sensor response time.

Another type of shaper is the correlated double sampler, illustrated in Fig. 14.19. This type of shaper is widely used in monolithically integrated circuits, as many CMOS processes (see Sect. 14.11.1) provide only capacitors and switches, but no resistors. This is an example of a time-variant filter. The $CR-nRC$ filter described in Fig. 14.18 acts continuously on the

signal, whereas the correlated double sample changes filter parameters versus time. Input signals are superimposed on a slowly fluctuating baseline. To remove the baseline fluctuations the baseline is sampled prior to the signal. Next, the signal plus baseline is sampled and the previous baseline sample subtracted to obtain the signal. The prefilter is critical to limit the noise bandwidth of the system. Filtering after the sampler is useless, as noise fluctuations on time scales shorter than the sample time will not be removed. Here the sequence of filtering is critical, unlike a time-invariant linear filter, where the sequence of filter functions can be interchanged.

14.9 Noise analysis of a detector and front-end amplifier

To determine how the pulse shaper affects the signal-to-noise ratio, consider the detector front end in Fig. 14.20. The detector is represented by the capacitance C_d , a relevant model for many radiation sensors. Sensor bias voltage is applied through the resistor R_b . The bypass capacitor C_b shunts any external interference coming through the bias supply line to ground. For high-frequency signals this capacitor appears as a low impedance, so for sensor signals the ‘far end’ of the bias resistor is connected to ground. The coupling capacitor C_c blocks the sensor bias voltage from the amplifier input, which is why a capacitor serving this rôle is also called a ‘blocking capacitor’. The series resistor R_s represents any resistance present in the connection from the sensor to the amplifier input. This includes the resistance of the sensor electrodes, the resistance of the connecting wires or traces, any resistance used to protect the amplifier against large voltage transients (‘input protection’), and parasitic resistances in the input transistor.

The following implicitly includes a constraint on the bias resistance, whose rôle is often misunderstood. It is often thought that the signal current generated in the sensor flows through R_b and the resulting voltage drop is measured. If the time constant $R_b C_d$ is small compared to the




Fig. 14.20. A detector front-end circuit and its equivalent circuit for noise calculations.

peaking time of the shaper T_P , the sensor will have discharged through R_b and much of the signal will be lost. Thus, we have the condition $R_b C_d \gg T_P$, or $R_b \gg T_P / C_d$. The bias resistor must be sufficiently large to block the flow of signal charge, so that all of the signal is available for the amplifier.

To analyse this circuit a voltage amplifier will be assumed, so all noise contributions will be calculated as a noise voltage appearing at the amplifier input. Steps in the analysis include: (1) determine the frequency distribution of all noise voltages presented to the amplifier input from all individual noise sources, (2) integrate over the frequency response of the shaper (for simplicity a CR - RC shaper) and determine the total noise voltage at the shaper output, and (3) determine the output signal for a known input signal charge. The equivalent noise charge (ENC) is the signal charge for which $S/N = 1$.

The *equivalent circuit* for the noise analysis (second panel of Fig. 14.20) includes both current and voltage noise sources. The ‘shot noise’ i_{nd} of the sensor leakage current is represented by a current noise generator in parallel with the sensor capacitance. As noted above, resistors can be modelled either as a voltage or current generator. Generally, resistors shunting the input act as noise current sources and resistors in series with the input act as noise voltage sources (which is why some in the detector community refer to current and voltage noise as *parallel* and *series noise*). Since the bias resistor effectively shunts the input, as the capacitor C_b passes current fluctuations to ground, it acts as a current generator i_{nb} and its noise current has the same effect as the shot-noise current from the detector. The shunt resistor can also be modelled as a noise voltage source, yielding the result that it acts as a current source. Choosing the appropriate model merely simplifies the calculation. Any other shunt resistances can be incorporated in the same way. Conversely, the series resistor R_s acts as a voltage generator. The electronic noise of the amplifier is described fully by a combination of voltage and current sources at its input, shown as e_{na} and i_{na} .

Thus, the noise sources are

$$\begin{aligned} \text{sensor bias current: } & i_{nd}^2 = 2eI_d \text{ ,} \\ \text{shunt resistance: } & i_{nb}^2 = \frac{4kT}{R_b} \text{ ,} \\ \text{series resistance: } & e_{ns}^2 = 4kTR_s \text{ ,} \\ \text{amplifier: } & e_{na}, i_{na} \text{ ,} \end{aligned}$$

where e is the electron charge, I_d the sensor bias current, k the Boltzmann constant, and T the temperature. Typical amplifier noise parameters e_{na}

and i_{na} are of order nV/\sqrt{Hz} and fA/\sqrt{Hz} (FETs)[¶] to pA/\sqrt{Hz} (bipolar transistors). Amplifiers tend to exhibit a ‘white’ noise spectrum at high frequencies (greater than order kHz), but at low frequencies show excess noise components with the spectral density

$$e_{nf}^2 = \frac{A_f}{f}, \tag{14.16}$$

where the noise coefficient A_f is device specific and of order 10^{-10} – $10^{-12} V^2$.

The noise voltage generators are in series and simply add in quadrature. White noise distributions remain white. However, a portion of the noise currents flows through the detector capacitance, resulting in a frequency-dependent noise voltage $i_n/(\omega C_d)$, so the originally white spectrum of the sensor shot noise and the bias resistor now acquires a $1/f$ dependence. The frequency distribution of all noise sources is further altered by the combined frequency response of the amplifier chain $A(f)$. Integrating over the cumulative noise spectrum at the amplifier output and comparing to the output voltage for a known input signal yields the signal-to-noise ratio. In this example the shaper is a simple CR – RC shaper, where for a given differentiation time constant the signal-to-noise ratio is maximised when the integration time constant equals the differentiation time constant, $\tau_i = \tau_d \equiv \tau$. Then the output pulse assumes its maximum amplitude at the time $T_P = \tau$.

Although the basic noise sources are currents or voltages, since radiation detectors are typically used to measure charge, the system’s noise level is conveniently expressed as an equivalent noise charge Q_n . As noted previously, this is equal to the detector signal that yields a signal-to-noise ratio of 1. The equivalent noise charge is commonly expressed in Coulombs, the corresponding number of electrons, or the equivalent deposited energy (eV). For the above circuit the equivalent noise charge is

$$Q_n^2 = \left(\frac{e^2}{8}\right) \left[\left(2eI_d + \frac{4kT}{R_b} + i_{na}^2\right) \cdot \tau + (4kTR_s + e_{na}^2) \cdot \frac{C_d^2}{\tau} + 4A_f C_d^2 \right]. \tag{14.17}$$

The prefactor $e^2/8 = \exp(2)/8 = 0.924$ normalises the noise to the signal gain. The first term combines all noise current sources and increases with shaping time. The second term combines all noise voltage sources and decreases with shaping time, but increases with sensor capacitance. The third term is the contribution of amplifier $1/f$ noise and, as a voltage source, also increases with sensor capacitance. The $1/f$ term is

[¶] FET – field effect transistor.

independent of shaping time, since for a $1/f$ spectrum the total noise depends on the ratio of upper to lower cutoff frequency, which depends only on shaper topology, but not on the shaping time.

The equivalent noise charge can be expressed in a more general form that applies to all types of pulse shapers:

$$Q_n^2 = i_n^2 F_i T_S + e_n^2 F_v \frac{C^2}{T_S} + F_{vf} A_f C^2, \quad (14.18)$$

where F_i , F_v and F_{vf} depend on the shape of the pulse determined by the shaper and T_S is a characteristic time, for example, the peaking time of a CR - nRC shaped pulse or the prefilter time constant in a correlated double sampler. C is the total parallel capacitance at the input, including the amplifier input capacitance. The shape factors F_i and F_v are easily calculated,

$$F_i = \frac{1}{2T_S} \int_{-\infty}^{\infty} [W(t)]^2 dt, \quad F_v = \frac{T_S}{2} \int_{-\infty}^{\infty} \left[\frac{dW(t)}{dt} \right]^2 dt. \quad (14.19)$$

For time-invariant pulse shaping $W(t)$ is simply the system's impulse response (the output signal seen on an oscilloscope) with the peak output signal normalised to unity. For a time-variant shaper the same equations apply, but $W(t)$ is determined differently. See [7–10] for more details.

A CR - RC shaper with equal time constants $\tau_i = \tau_d$ has $F_i = F_v = 0.9$ and $F_{vf} = 4$, independent of the shaping time constant, so for the circuit in Fig. 14.17 Eq. (14.18) becomes

$$Q_n^2 = \left(2q_e I_d + \frac{4kT}{R_b} + i_{na}^2 \right) F_i T_S + (4kT R_s + e_{na}^2) F_v \frac{C^2}{T_S} + F_{vf} A_f C^2. \quad (14.20)$$

Pulse shapers can be designed to reduce the effect of current noise, to mitigate radiation damage, for example. Increasing pulse symmetry tends to decrease F_i and increase F_v , e.g. to $F_i = 0.45$ and $F_v = 1.0$ for a shaper with one CR differentiator and four cascaded RC integrators.

Figure 14.21 shows how equivalent noise charge is affected by shaping time. At short shaping times the voltage noise dominates, whereas at long shaping times the current noise takes over. Minimum noise is obtained where the current and voltage contributions are equal. The noise minimum is flattened by the presence of $1/f$ noise. Also shown is that increasing the detector capacitance will increase the voltage noise contribution and shift the noise minimum to longer shaping times, albeit with an increase in minimum noise.




Fig. 14.21. Equivalent noise charge versus shaping time. At small shaping times (large bandwidth) the equivalent noise charge is dominated by voltage noise, whereas at long shaping times (large integration times) the current noise contributions dominate. The total noise assumes a minimum where the current and voltage contributions are equal. The ‘1/f’ noise contribution is independent of shaping time and flattens the noise minimum. Increasing the voltage or current noise contribution shifts the noise minimum. Increased voltage noise is shown as an example.

For quick estimates one can use the following equation, which assumes a field effect transistor (FET) amplifier (negligible i_{na}) and a simple CR – RC shaper with peaking time τ . The noise is expressed in units of the electron charge e and C is the total parallel capacitance at the input, including C_d , all stray capacitances and the amplifier’s input capacitance

$$Q_n^2 = 12 \left[\frac{e^2}{\text{nA ns}} \right] I_d \tau + 6 \cdot 10^5 \left[\frac{e^2 \text{ k}\Omega}{\text{ns}} \right] \frac{\tau}{R_b} + 3.6 \cdot 10^4 \left[\frac{e^2 \text{ ns}}{(\text{pF})^2 (\text{nV})^2 / \text{Hz}} \right] e_n^2 \frac{C^2}{\tau} \quad (14.21)$$

The noise charge is improved by reducing the detector capacitance and leakage current, judiciously selecting all resistances in the input circuit, and choosing the optimum shaping time constant. The noise parameters of a well-designed amplifier depend primarily on the input device. Fast, high-gain transistors are generally best.

In field effect transistors, both junction field effect transistors (JFETs) and metal oxide semiconductor field effect transistors (MOSFETs), the

noise current contribution is very small, so reducing the detector leakage current and increasing the bias resistance will allow long shaping times with correspondingly lower noise. The equivalent input noise voltage is $e_n^2 \approx 4kT/g_m$, where g_m is the transconductance,^{||} which increases with operating current. For a given current, the transconductance increases when the channel length is reduced, so reductions in feature size with new process technologies are beneficial. At a given channel length, minimum noise is obtained when a device is operated at maximum transconductance. If lower noise is required, the width of the device can be increased (equivalent to connecting multiple devices in parallel). This increases the transconductance (and required current) with a corresponding decrease in noise voltage, but also increases the input capacitance. At some point the reduction in noise voltage is outweighed by the increase in total input capacitance. The optimum is obtained when the FET's input capacitance equals the external capacitance (sensor + stray capacitance). Note that this capacitive matching criterion only applies when the input-current noise contribution of the amplifying device is negligible.

Capacitive matching comes at the expense of power dissipation. Since the minimum is shallow, one can operate at significantly lower currents with just a minor increase in noise. In large detector arrays power dissipation is critical, so FETs are hardly ever operated at their minimum noise. Instead, one seeks an acceptable compromise between noise and power dissipation (see [1] for a detailed discussion). Similarly, the choice of input devices is frequently driven by available fabrication processes. High-density integrated circuits tend to include only MOSFETs, so this determines the input device, even where a bipolar transistor would provide better performance.

In bipolar transistors the shot noise associated with the base current I_B is significant, $i_{nB}^2 = 2eI_B$. Since $I_B = I_C/\beta_{DC}$, where I_C is the collector current and β_{DC} the direct current gain, this contribution increases with device current. On the other hand, the equivalent input noise voltage

$$e_n^2 = \frac{2(kT)^2}{eI_C} \quad (14.22)$$

decreases with collector current, so the noise assumes a minimum at a specific collector current,

$$Q_{n,\min}^2 = 4kT \frac{C}{\sqrt{\beta_{DC}}} \sqrt{F_i F_v} \quad \text{at} \quad I_C = \frac{kT}{e} C \sqrt{\beta_{DC}} \sqrt{\frac{F_v}{F_i}} \frac{1}{T_S} \quad (14.23)$$

^{||} Transconductance is the ratio of the current change at the output and the corresponding voltage change at the input, $g_m = \Delta I_{out}/\Delta V_{in}$.

For a CR – RC shaper and $\beta_{DC} = 100$,

$$Q_{n,\min} \approx 250 \left[\frac{e}{\sqrt{\text{pF}}} \right] \cdot \sqrt{C} \quad \text{at} \quad I_C = 260 \left[\frac{\mu\text{A ns}}{\text{pF}} \right] \cdot \frac{C}{T_S}. \quad (14.24)$$

The minimum obtainable noise is independent of shaping time (unlike FETs), but only at the optimum collector current I_C , which does depend on shaping time.

In bipolar transistors the input capacitance is usually much smaller than the sensor capacitance (of order 1 pF for $e_n \approx 1 \text{ nV}/\sqrt{\text{Hz}}$) and substantially smaller than in FETs with comparable noise. Since the transistor input capacitance enters into the total input capacitance, this is an advantage. Note that capacitive matching does not apply to bipolar transistors, because their noise current contribution is significant. Due to the base current noise bipolar transistors are best at short shaping times, where they also require lower power than FETs for a given noise level.

When the input noise current is negligible, the noise increases linearly with sensor capacitance. The noise slope

$$\frac{dQ_n}{dC_d} \approx 2e_n \cdot \sqrt{\frac{F_v}{T}} \quad (14.25)$$

depends both on the preamplifier (e_n) and the shaper (F_v , T). The zero intercept can be used to determine the amplifier input capacitance plus any additional capacitance at the input node.

Practical noise levels range from $< 1e$ for charge-coupled devices (CCDs) at long shaping times to $\approx 10^4 e$ in high-capacitance liquid-Ar calorimeters. Silicon strip detectors typically operate at $\approx 10^3$ electrons, whereas pixel detectors with fast readout provide noise of 100–200 electrons. Transistor noise is discussed in more detail in [1].

14.10 Timing measurements

Pulse-height measurements discussed up to now emphasise measurement of signal charge. *Timing measurements* seek to optimise the determination of the time of occurrence. Although, as in amplitude measurements, signal-to-noise ratio is important, the determining parameter is not signal-to-noise, but slope-to-noise ratio. This is illustrated in Fig. 14.22, which shows the leading edge of a pulse fed into a threshold discriminator (comparator), a *leading-edge trigger*. The instantaneous signal level is modulated by noise, where the variations are indicated by the




Fig. 14.22. Fluctuations in signal amplitude crossing a threshold translate into timing fluctuations (a). With realistic pulses the slope changes with amplitude, so minimum timing jitter occurs with the trigger level at the maximum slope.

shaded band. Because of these fluctuations, the time of threshold crossing fluctuates. By simple geometrical projection, the timing variance or jitter is

$$\sigma_t = \frac{\sigma_n}{(dS/dt)_{S_T}} \approx t_r \frac{\sigma_n}{S} , \tag{14.26}$$

where σ_n is the rms noise and the derivative of the signal dS/dt is evaluated at the trigger level S_T . To increase dS/dt without incurring excessive noise, the amplifier bandwidth should match the rise time of the detector signal. The 10%–90% rise time of an amplifier with bandwidth f_u (see Fig. 14.11) is

$$t_r = 2.2 \tau = \frac{2.2}{2\pi f_u} = \frac{0.35}{f_u} . \tag{14.27}$$

For example, an oscilloscope with 350 MHz bandwidth has a 1 ns rise time. When amplifiers are cascaded, which is invariably necessary, the individual rise times add in quadrature:

$$t_r \approx \sqrt{t_{r1}^2 + t_{r2}^2 + \dots + t_{rn}^2} . \tag{14.28}$$

Increasing signal-to-noise ratio improves time resolution, so minimising the total capacitance at the input is also important. At high signal-to-noise ratios the time jitter can be much smaller than the rise time.

The second contribution to time resolution is time walk, where the timing signal shifts with amplitude as shown in Fig. 14.23. This can be corrected by various means, either in hardware or software. For more detailed tutorials on timing measurements, see [1, 11].




Fig. 14.23. The time at which a signal crosses a fixed threshold depends on the signal amplitude, leading to ‘time walk’.

14.11 Digital electronics

Analogue signals utilise continuously variable properties of the pulse to impart information, such as the pulse amplitude or pulse shape. Digital signals have constant amplitude, but the presence of the signal at specific times is evaluated, i.e. whether the signal is in one of two states, ‘low’ or ‘high’. However, this still involves an analogue process, as the presence of a signal is determined by the signal level exceeding a threshold at the proper time.

14.11.1 Logic elements

Figure 14.24 illustrates several functions utilised in digital circuits (‘logic’ functions). An *AND gate* provides an output only when all inputs are




Fig. 14.24. Basic logic functions include gates (AND, OR, Exclusive OR) and flip-flops. The outputs of the AND and D flip-flop show how small shifts in relative timing between inputs can determine the output state.

high. An *OR* gives an output when any input is high. An *eXclusive OR* (*XOR*) responds when only one input is high. The same elements are commonly implemented with inverted outputs, then called NAND and NOR gates, for example. The D flip-flop is a bistable memory circuit that records the presence of a signal at the data input D when a signal transition occurs at the clock input CLK. This device is commonly called a *latch*. Inverted inputs and outputs are denoted by small circles or by superimposed bars, e.g. \bar{Q} is the inverted output of a flip-flop, as shown in Fig. 14.25.

Logic circuits are fundamentally amplifiers, so they also suffer from bandwidth limitations. The pulse train of the AND gate in Fig. 14.24 illustrates a common problem. The third pulse of input B is going low at the same time that input A is going high. Depending on the time overlap, this can yield a narrow output that may or may not be recognised by the following circuit. In an XOR this can occur when two pulses arrive nearly at the same time. The D flip-flop requires a minimum setup time for a level change at the D input to be recognised, so changes in the data level may not be recognised at the correct time. These marginal events may be extremely rare and perhaps go unnoticed. However, in complex systems the combination of ‘glitches’ can make the system ‘hang up’, necessitating a system reset. Data transmission protocols have been developed to detect such errors (parity checks, Hamming codes, etc.), so corrupted data can be rejected.

Some key aspects of logic systems can be understood by inspecting the circuit elements that are used to form logic functions. In an *n*-channel metal oxide semiconductor (NMOS) transistor a conductive channel is formed when the input electrode is biased positive with respect to the

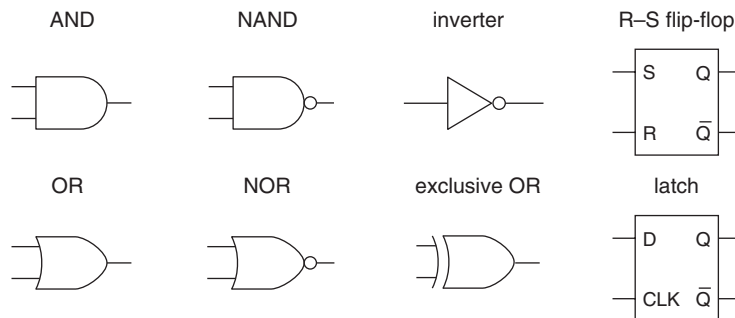


Fig. 14.25. Some common logic symbols. Inverted outputs are denoted by small circles or by a superimposed bar, as for the latch output \bar{Q} . Additional inputs can be added to gates as needed. An R-S flip-flop (R-S: Reset-Set) sets the Q output high in response to an S input. An R input resets the Q output to low.

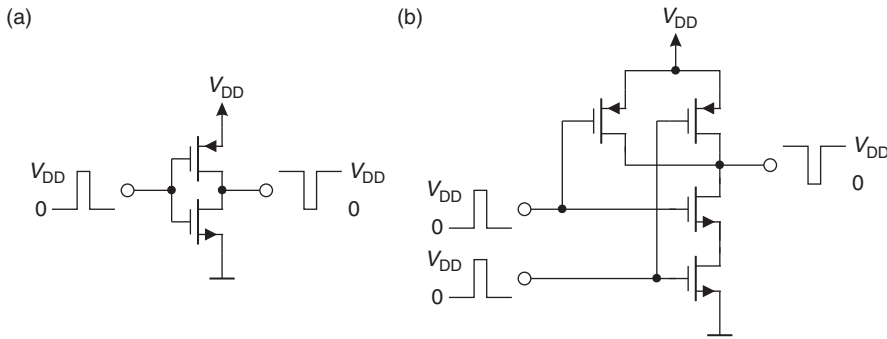


Fig. 14.26. A CMOS inverter (a) and NAND gate (b).

channel. The input, called the ‘gate’, is capacitively coupled to the output channel connected between the ‘drain’ and ‘source’ electrodes. A *p*-channel (PMOS) transistor is the complementary device, where a conductive channel is formed when the gate is biased negative with respect to the source.

Complementary MOS (CMOS) logic utilises both NMOS and PMOS transistors as shown in Fig. 14.26. In the inverter the lower (NMOS) transistor is turned off when the input is low, but the upper (PMOS) transistor is turned on, so the output is connected to V_{DD} , taking the output high. Since the current path from V_{DD} to ground is blocked by either the NMOS or PMOS device being off, the power dissipation is zero in both the high and low states. Current only flows during the level transition when both devices are on as the input level is at approximately $V_{DD}/2$. As a result, the power dissipation of CMOS logic is significantly lower than in NMOS or PMOS circuits, which draw current in either one or the other logic state. However, this reduction in power is obtained only in logic circuitry. CMOS analogue amplifiers are not fundamentally more power efficient than NMOS or PMOS circuits, although CMOS allows more efficient circuit topologies.

14.11.2 Propagation delays and power dissipation

Logic elements always operate in conjunction with other circuits, as illustrated in Fig. 14.27. The wiring resistance together with the total load capacitance increases the rise time of the logic pulse and as a result delays the time when the transition crosses the logic threshold. The energy dissipated in the wiring resistance R is

$$E = \int i^2(t)R dt . \quad (14.29)$$

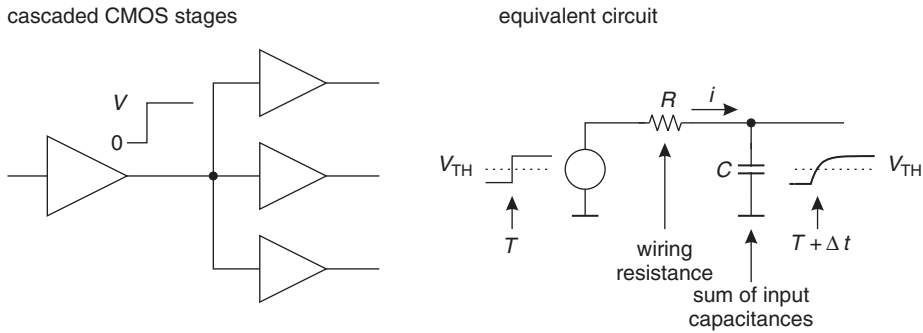


Fig. 14.27. The wiring resistance together with the distributed load capacitance delays the signal.

The current flow during one transition is

$$i(t) = \frac{V}{R} \exp\left(-\frac{t}{RC}\right), \quad (14.30)$$

so the dissipated energy per transition (either positive or negative)

$$E = \frac{V^2}{R} \int_0^{\infty} \exp\left(-\frac{2t}{RC}\right) dt = \frac{1}{2} CV^2. \quad (14.31)$$

When pulses occur at a frequency f , the power dissipated in both the positive and negative transitions is

$$P = fCV^2. \quad (14.32)$$

Thus, the power dissipation increases with clock frequency and the square of the logic swing.

Fast logic is time-critical. It relies on logic operations from multiple paths coming together at the right time. Valid results depend on maintaining minimum allowable overlaps and set-up times as illustrated in Fig. 14.24. Each logic circuit has a finite propagation delay, which depends on circuit loading, i.e. how many loads the circuit has to drive. In addition, as illustrated in Fig. 14.27 the wiring resistance and capacitive loads introduce delay. This depends on the number of circuits connected to a wire or trace, the length of the trace and the dielectric constant of the substrate material. Relying on control of circuit and wiring delays to maintain timing requires great care, as it depends on circuit variations and temperature. In principle all of this can be simulated, but in complex systems there are too many combinations to test every one. A more robust solution is to use synchronous systems, where the timing of all transitions




Fig. 14.28. Complex logic circuits are commonly implemented using logic arrays that as an integrated block provide the desired outputs in response to specific input combinations.

is determined by a master clock. Generally, this does not provide the utmost speed and requires some additional circuitry, but increases reliability. Nevertheless, clever designers frequently utilise asynchronous logic. Sometimes it succeeds... and sometimes it does not.

14.11.3 Logic arrays

Commodity integrated circuits with basic logic blocks are readily available, e.g. with four NAND gates or two flip-flops in one package. These can be combined to form simple digital systems. However, complex logic systems are no longer designed using individual gates. Instead, logic functions are described in a high-level language (e.g. VHDL**), synthesised using design libraries, and implemented as custom ICs – *application-specific ICs* (ASICs) – or programmable logic arrays. In these implementations the digital circuitry no longer appears as an ensemble of inverters, gates and flip-flops, but as an integrated logic block that provides specific outputs in response to various input combinations. This is illustrated in Fig. 14.28. Field Programmable Gate or logic Arrays (FPGAs) are a common example. A representative FPGA has 512 pads usable for inputs and outputs, $\approx 10^6$ gates, and ≈ 100 K of memory. Modern design tools also account for propagation delays, wiring lengths, loads and temperature dependence. The design software also generates ‘test vectors’ that can be used to test finished parts. Properly implemented, complex digital designs can succeed on the first pass, whether as ASICs or as logic or gate arrays.

14.12 Analogue-to-digital conversion

For data storage and subsequent analysis the analogue signal at the shaper output must be digitised. Important parameters for *analogue-to-digital converters* (ADCs or A/Ds) used in detector systems are as follows:

** VHDL – VHSIC Hardware Description Language; VHSIC – Very High Speed Integrated Circuit.

- (i) Resolution: The ‘granularity’ of the digitised output.
- (ii) Differential non-linearity: How uniform are the digitisation increments?
- (iii) Integral non-linearity: Is the digital output proportional to the analogue input?
- (iv) Conversion time: How much time is required to convert an analogue signal to a digital output?
- (v) Count-rate performance: How quickly can a new conversion commence after completion of a prior one without introducing deleterious artifacts?
- (vi) Stability: Do the conversion parameters change with time?

Instrumentation ADCs used in industrial data acquisition and control systems share most of these requirements. However, detector systems place greater emphasis on differential non-linearity and count-rate performance. The latter is important, as detector signals often occur randomly, in contrast to systems where signals are sampled at regular intervals. As in amplifiers, if the DC gain is not precisely equal to the high-frequency gain, the baseline will shift. Furthermore, following each pulse it takes some time for the baseline to return to its quiescent level. For periodic signals of roughly equal amplitude these baseline deviations will be the same for each pulse, but for a random sequence of pulses with varying amplitudes, the instantaneous baseline level will be different for each pulse and broaden the measured signal.

Conceptually, the simplest technique is *flash conversion*, illustrated in Fig. 14.29. The signal is fed in parallel to a bank of threshold comparators. The individual threshold levels are set by a resistive divider. The comparator outputs are encoded such that the output of the highest-level comparator that fires yields the correct bit pattern. The threshold levels can be set to provide a linear conversion characteristic where each bit corresponds to the same analogue increment, or a non-linear characteristic to provide increments proportional to the absolute level, which provides constant relative resolution over the range, for example.

The big advantage of this scheme is speed; conversion proceeds in one step and conversion times < 10 ns are readily achievable. The drawbacks are component count and power consumption, as one comparator is required per conversion bin. For example, an 8-bit converter requires 256 comparators. The conversion is always monotonic and differential non-linearity is determined by the matching of the resistors in the threshold divider. Only relative matching is required, so this topology is a good match for monolithic integrated circuits. *Flash ADCs* are available with

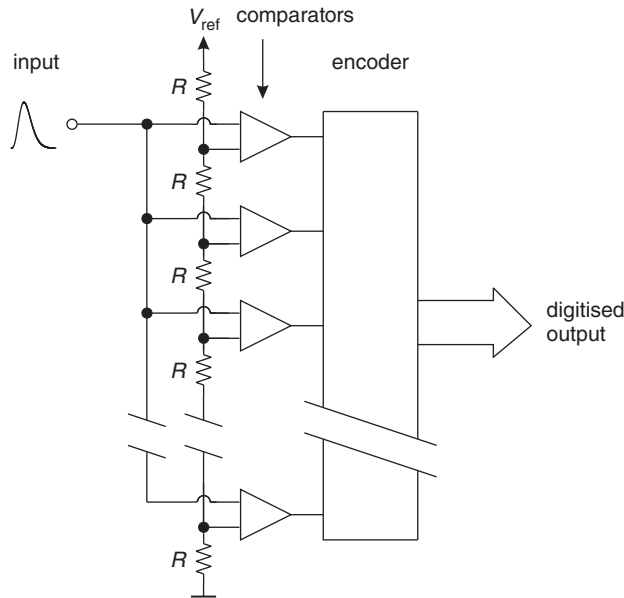


Fig. 14.29. Block diagram of a flash ADC.

conversion rates > 500 MS/s (megasamples per second) at 8-bit resolution and a power dissipation of about 5 W.

The most commonly used technique is the successive-approximation ADC, shown in Fig. 14.30. The input pulse is sent to a pulse stretcher, which follows the signal until it reaches its cusp and then holds the peak value. The stretcher output feeds a comparator, whose reference is provided by a digital-to-analogue converter (DAC). The DAC is cycled beginning with the most significant bits. The corresponding bit is set when the comparator fires, i.e. the DAC output becomes less than the pulse height. Then the DAC cycles through the less significant bits, always setting the corresponding bit when the comparator fires. Thus, n -bit resolution requires n steps and yields 2^n bins. This technique makes efficient use of circuitry and is fairly fast. High-resolution devices (16–20 bits) with conversion times of order μ s are readily available. Currently a 16-bit ADC with a conversion time of 1 μ s (1 MS/s) requires about 100 mW.

A common limitation is differential non-linearity (DNL), since the resistors that set the DAC levels must be extremely accurate. For $\text{DNL} < 1\%$ the resistor determining the 2^{12} -level in a 13-bit ADC must be accurate to $< 2.4 \cdot 10^{-6}$. As a consequence, differential non-linearity in high-resolution successive-approximation converters is typically 10%–20% and often exceeds the 0.5 LSB (least significant bit) required to ensure monotonic response.

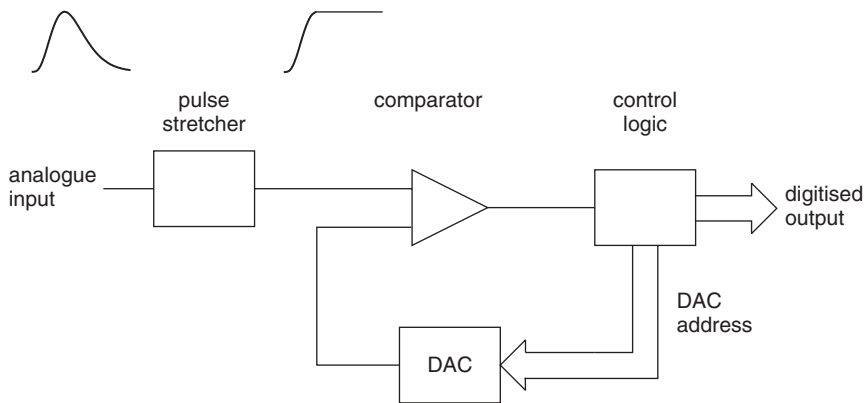


Fig. 14.30. Principle of a successive-approximation ADC. The DAC is controlled to sequentially add levels proportional to $2^n, 2^{n-1}, \dots, 2^0$. The corresponding bit is set if the comparator output is high (DAC output $<$ pulse height).




Fig. 14.31. Principle of a Wilkinson ADC. After the peak amplitude has been acquired, the output of the peak detector initiates the conversion process. The memory capacitor is discharged by a constant current while counting the clock pulses. When the capacitor is discharged to the baseline level V_{BL} , the comparator output goes low and the conversion is complete.

The *Wilkinson ADC* [12] has traditionally been the mainstay of precision pulse digitisation. The principle is shown in Fig. 14.31. The peak signal amplitude is acquired by a combined peak detector/pulse stretcher and transferred to a memory capacitor. The output of the peak detector initiates the conversion process:

- (i) The memory capacitor is disconnected from the stretcher,
- (ii) a current source is switched on to linearly discharge the capacitor with current I_R , and simultaneously
- (iii) a counter is enabled to determine the number of clock pulses until the voltage on the capacitor reaches the baseline level V_{BL} .

The time required to discharge the capacitor is a linear function of pulse height, so the counter content provides the digitised pulse height. The clock pulses are provided by a crystal oscillator, so the time between pulses is extremely uniform and this circuit inherently provides excellent differential linearity. The drawback is the relatively long conversion time T_C , which is proportional to the pulse height, $T_C = n \cdot T_{clk}$, where the channel number n corresponds to the pulse height. For example, a clock frequency of 100 MHz provides a clock period $T_{clk} = 10$ ns and a maximum conversion time $T_C = 82 \mu\text{s}$ for 13 bits ($n = 8192$). Clock frequencies of 100 MHz are typical, but > 400 MHz have been implemented with excellent performance ($\text{DNL} < 10^{-3}$). This scheme makes efficient use of circuitry and allows low power dissipation. Wilkinson ADCs have been implemented in 128-channel readout ICs for silicon strip detectors [13]. Each ADC added only $100 \mu\text{m}$ to the length of a channel and a power of $300 \mu\text{W}$ per readout channel.

14.13 Time-to-digital converters (TDCs)

The combination of a clock generator with a counter is the simplest technique for *time-to-digital conversion*, as shown in Fig. 14.32. The clock pulses are counted between the start and stop signals, which yields a direct readout in real time. The limitation is the speed of the counter, which in current technology is limited to about 1 GHz, yielding a time resolution of 1 ns. Using the stop pulse to strobe the instantaneous counter status into a register provides multi-hit capability.

Analogue techniques are commonly used in high-resolution digitisers to provide resolution in the range of ps to ns. The principle is to convert a time interval into a voltage by charging a capacitor through a switchable current source. The start pulse turns on the current source and the stop pulse turns it off. The resulting voltage on the capacitor C is $V = Q/C = I_T(T_{\text{stop}} - T_{\text{start}})/C$, which is digitised by an ADC. A convenient implementation switches the current source to a smaller discharge current I_R and uses a Wilkinson ADC for digitisation, as illustrated in Fig. 14.33. This technique provides high resolution, but at the expense of dead time and multi-hit capability.

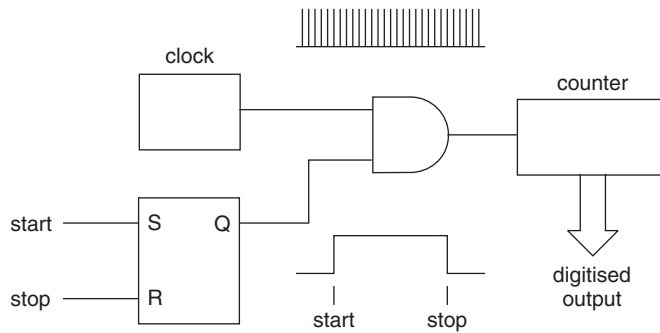


Fig. 14.32. The simplest form of time digitiser counts the number of clock pulses between the start and stop signals.




Fig. 14.33. Combining a time-to-amplitude converter with an ADC forms a time digitiser capable of ps resolution. The memory capacitor C is charged by the current I_T for the duration $T_{\text{start}} - T_{\text{stop}}$ and subsequently discharged by a Wilkinson ADC.

14.14 Signal transmission

Signals are transmitted from one unit to another through *transmission lines*, often coaxial cables or ribbon cables. When transmission lines are not terminated with their characteristic impedance, the signals are reflected. As a signal propagates along the cable, the ratio of instantaneous voltage to current equals the cable's characteristic impedance $Z_0 = \sqrt{L/C}$, where L and C are the inductance and capacitance per unit length. Typical impedances are $50\ \Omega$ or $75\ \Omega$ for coaxial cables and $\approx 100\ \Omega$ for ribbon cables. If at the receiving end the cable is connected to a resistance different from the cable impedance, a different ratio of voltage to current must be established. This occurs through a reflected signal. If the termination is less than the line impedance, the voltage must be smaller and the reflected voltage wave has the opposite sign. If the termination is greater than the line impedance, the voltage wave is reflected with the same polarity. Conversely, the current in the reflected wave is of like sign when the termination is less than the line impedance and of opposite sign when the termination is greater. Voltage pulse reflections are illustrated in Fig. 14.34. At the sending end the reflected pulse appears after twice the propagation delay of the cable. Since in the presence of a dielectric the velocity of propagation is $v = c/\sqrt{\epsilon}$, in typical coaxial and ribbon cables the delay is $5\ \text{ns/m}$.

Cable drivers often have a low output impedance, so the reflected pulse is reflected again towards the receiver, to be reflected again, etc. This is

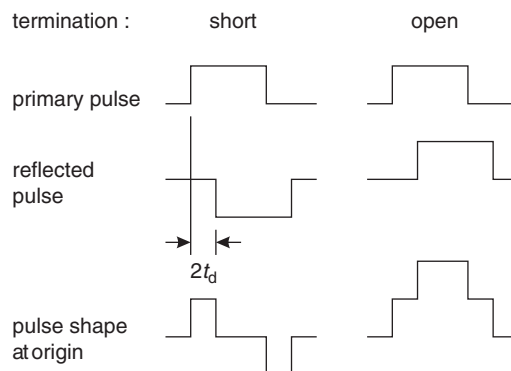


Fig. 14.34. Voltage pulse reflections on a transmission line terminated either with a short (left) or open circuit (right). Measured at the sending end, the reflection from a short at the receiving end appears as a pulse of opposite sign delayed by the round-trip delay of the cable. If the total delay is less than the pulse width, the signal appears as a bipolar pulse. Conversely, an open circuit at the receiving end causes a reflection of like polarity.




Fig. 14.35. (a) Signal observed in an amplifier when a low-impedance driver is connected to an amplifier through a 4 m long coaxial cable. The cable impedance is $50\ \Omega$ and the amplifier input appears as $1\ \text{k}\Omega$ in parallel with $30\ \text{pF}$. When the receiving end is properly terminated with $50\ \Omega$, the reflections disappear (b).




Fig. 14.36. Cables may be terminated at the receiving end (top, shunt termination) or sending end (bottom, series termination).

shown in Fig. 14.35, which shows the observed signal when the output of a low-impedance pulse driver is connected to a high-impedance amplifier input through a 4 m long $50\ \Omega$ coaxial cable. When feeding a counter, a single pulse will be registered multiple times, depending on the threshold level. When the amplifier input is terminated with $50\ \Omega$, the reflections disappear and only the original 10 ns wide pulse is seen.

There are two methods of terminating cables, which can be applied either individually or – in applications where pulse fidelity is critical – in combination. As illustrated in Fig. 14.36 the termination can be applied at the receiving or the sending end. Receiving-end termination absorbs the signal pulse when it arrives at the receiver. With sending-end termination the pulse is reflected at the receiver, but since the reflected pulse is absorbed at the sender, no additional pulses are visible at the receiver. At the sending end the original pulse is attenuated two-fold by the voltage divider formed by the series resistor and the cable impedance. However, at the receiver the pulse is reflected with the same polarity,

so the superposition of the original and the reflected pulses provides the original amplitude.

This example uses voltage amplifiers, which have low output and high input impedances. It is also possible to use current amplifiers, although this is less common. Then, the amplifier has a high output impedance and low input impedance, so shunt termination is applied at the sending end and series termination at the receiving end.

Terminations are never perfect, especially at high frequencies where stray capacitance becomes significant. For example, the reactance of 10 pF at 100 MHz is 160 Ω . Thus, critical applications often use both series and parallel termination, although this does incur a 50% reduction in pulse amplitude. In the μs regime, amplifier inputs are usually designed as high impedance, whereas timing amplifiers tend to be internally terminated, but one should always check if this is the case. As a rule of thumb, whenever the propagation delay of cables (or connections in general) exceeds a few per cent of the signal rise time, proper terminations are required.

14.15 Interference and pickup

The previous discussion analysed random noise sources inherent to the sensor and front-end electronics. In practical systems external noise often limits the obtainable detection threshold or energy resolution. As with random noise, external *pickup* introduces baseline fluctuations. There are many possible sources, radio and television stations, local radio frequency (RF) generators, system clocks, transients associated with trigger signals and data readout, etc. Furthermore, there are many ways through which these undesired signals can enter the system. Again, a comprehensive review exceeds the allotted space, so only a few key examples of pickup mechanisms will be shown. A more detailed discussion is given in [1, 2]. Ott [14] gives a more general treatment and texts by Johnson and Graham [15, 16] give useful details on signal transmission and design practices.

14.15.1 Pickup mechanisms

The most sensitive node in a detector system is the input. Figure 14.37 shows how very small spurious signals coupled to the sensor backplane can inject substantial charge. Any change in the bias voltage ΔV directly at the sensor backplane will inject a charge $\Delta Q = C_d \Delta V$. Assume a silicon strip sensor with 10 cm strip length. Then the capacitance C_d from the backplane to a single strip is about 1 pF. If the noise level is 1000 electrons ($1.6 \cdot 10^{-16}$ C), ΔV must be much smaller than $Q_n/C_d = 160 \mu\text{V}$. This can be introduced as noise from the bias supply (some voltage supplies are quite noisy; switching power supplies can be clean, but most are




Fig. 14.37. Noise on the detector bias line is coupled through the detector capacitance to the amplifier input.




Fig. 14.38. Shared current paths introduce common voltage drops to different circuits.

not) or noise on the ground plane can couple through the capacitor C . Naively, one might assume the ground plane to be 'clean', but it can carry significant interference for the following reason.

One of the most common mechanisms for cross-coupling is shared current paths, often referred to as 'ground loops'. However, this phenomenon is not limited to grounding. Consider two systems: The first is transmitting large currents from a source to a receiver. The second is similar, but is attempting a low-level measurement. Following the prevailing lore, both systems are connected to a massive ground bus, as shown in Fig. 14.38. Current seeks the path of least resistance, so the large current from source V_1 will also flow through the ground bus. Although the ground bus is massive, it does not have zero resistance, so the large current flowing through the ground system causes a voltage drop ΔV .

In system 2 (source V_2) both signal source and receiver are also connected to the ground system. Now the voltage drop ΔV from system 1 is in series with the signal path, so the receiver measures $V_2 + \Delta V$. The cross-coupling has nothing to do with grounding per se, but is due to

the common return path. However, the common ground caused the problem by establishing the shared path. This mechanism is not limited to large systems with external ground busses, but also occurs on the scale of printed circuit boards and micron-scale integrated circuits. At high frequencies the impedance is increased due to skin effect and inductance. Note that for high-frequency signals the connections can be made capacitively, so even if there is no DC path, the parasitic capacitance due to mounting structures or adjacent conductor planes can be sufficient to close the loop.

The traditional way of dealing with this problem is to reduce the impedance of the shared path, which leads to the *copper braid syndrome*. However, changes in the system will often change the current paths, so this 'fix' is not very reliable. Furthermore, in many detector systems – tracking detectors, for example – the additional material would be prohibitive. Instead, it is best to avoid the root cause.

14.15.2 Remedial techniques

Figure 14.39 shows a sensor connected to a multistage amplifier. Signals are transferred from stage to stage through definite current paths. It is critical to maintain the integrity of the signal paths, but this does not depend on grounding – indeed Fig. 14.39 does not show any ground connection at all. The most critical parts of this chain are the input, which is the most sensitive node, and the output driver, which tends to circulate the largest current. Circuit diagrams usually are not drawn like Fig. 14.39; the bottom common line is typically shown as ground. For example, in Fig. 14.37 the sensor signal current flows through capacitor C and reaches




Fig. 14.39. The signal is transferred from the sensor to the input stage and from stage to stage via local current loops.

the return node of the amplifier through 'ground'. Clearly, it is critical to control this path and keep deleterious currents from this area.

However superfluous *grounding* may be, one cannot let circuit elements simply float with respect to their environment. Capacitive coupling is always present and any capacitive coupling between two points of different potential will induce a signal. This is illustrated in Fig. 14.40, which represents individual detector modules mounted on a support/cooling structure. Interference can couple through the parasitic capacitance of the mount, so it is crucial to reduce this capacitance and control the potential

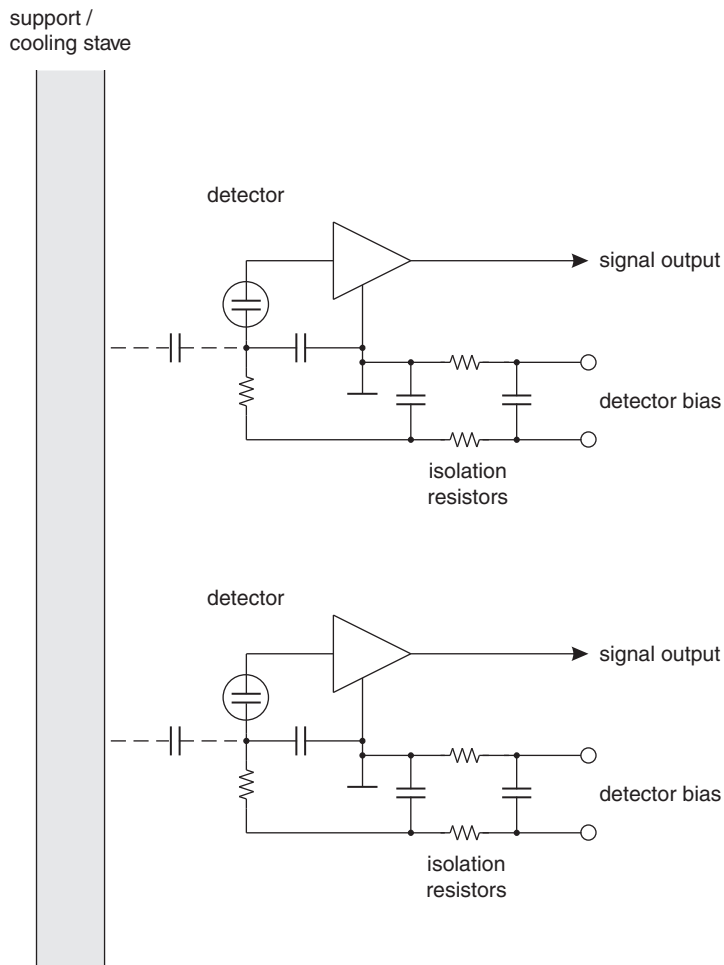


Fig. 14.40. Capacitive coupling between detectors or detector modules and their environment introduces interference when relative potentials and stray capacitance are not controlled.

of the support structure relative to the detector module. Attaining this goal in reality is a challenge, which is not always met successfully. Nevertheless, paying attention to signal paths and potential references early on is much easier than attempting to correct a poor design after it is done. *Troubleshooting* is exacerbated by the fact that current paths interact, so doing the ‘wrong’ thing sometimes brings improvement. Furthermore, only one mistake can ruin system performance, so if this has been designed into the system from the outset, one is left with compromises. Nevertheless, although this area is rife with myths, basic physics still applies.

14.16 Conclusion

Signal processing is a key part of modern detector systems. Proper design is especially important when signals are small and electronic noise determines detection thresholds or resolution. Optimisation of noise is well understood and predicted noise levels can be achieved in practical experiments within a few per cent of predicted values. However, systems must be designed very carefully to avoid extraneous pickup.

14.17 Problems

- 14.1** In a time-of-flight system the start detector has a time resolution of 100 ps and the stop detector has 50 ps resolution. What is the overall time resolution?
- 14.2** Consider a spectroscopy system whose resolution is determined by electronic noise.
- (a) The current noise contribution is 120 eV and the voltage noise contribution is 160 eV. What is the total noise?
 - (b) After cooling the detector, the current noise is 10 eV and the voltage noise remains unchanged at 160 eV. What is the total noise?
- 14.3** An X-ray spectroscopy system is to resolve the Tl $K_{\alpha 1}$ and $K_{\alpha 2}$ emissions from a ^{203}Hg source. The $K_{\alpha 1}$ and $K_{\alpha 2}$ energies are 72.87 keV and 70.83 keV, at about equal intensities.
- (a) Determine the energy resolution required to separate the two X-ray peaks.

- (b) The intrinsic energy resolution of the detector is $\sigma_{\text{det}} = 160 \text{ eV}$. What is the allowable electronic noise contribution?

14.4 A spectroscopy system has the front-end components shown in Fig. 14.20. The Si detector draws a reverse-bias current of 100 nA and has a capacitance of 100 pF . The bias resistor has $R_{\text{b}} = 10 \text{ M}\Omega$ and the total resistance of the connections between the detector and the preamplifier input is 10Ω . The preamplifier has an equivalent input noise voltage of $1 \text{ nV}/\sqrt{\text{Hz}}$ with negligible $1/f$ and current noise.

- (a) The system utilises a simple CR – RC pulse shaper with integration and differentiation time constants of $1 \mu\text{s}$. What is the electronic noise expressed in electrons and in eV? How large are the contributions of the individual noise sources?
- (b) Assume a CR – RC shaper with adjustable peaking time, where both the integration and differentiation time constants are adjusted simultaneously to be equal. What are the noise current and noise voltage contributions at $1 \mu\text{s}$ shaping time? Determine the time constant that yields minimum noise.
- (c) Using the CR – RC shaper at the optimum shaping time determined in (b), what is the minimum value of bias resistor that will degrade the overall noise by less than 1%?

14.5 The signal at the input of a voltage-sensitive amplifier is a 10 mV pulse with a rise time of 10 ns (10%–90%). The equivalent input noise of the amplifier is $10 \mu\text{V}$ rms. The amplifier feeds a simple threshold comparator.

- (a) Assume a comparator threshold of 5 mV . What is the timing jitter?
- (b) Still keeping the threshold at 5 mV , how much does the output of the comparator shift when the signal changes from 10 mV to 50 mV ? As an approximation assume a perfectly linear transition:

$$t(V_{\text{T}}) = \frac{V_{\text{T}}}{V_{\text{s}}} t_{\text{r}} + t_0 ,$$

where t_0 is the time at which the pulse attains 10% of its peak, i.e. 1 ns .

References

- [1] H. Spieler, *Semiconductor Detector Systems*, Oxford University Press, Oxford (2005) ISBN 0-19-852784-5
- [2] www-physics.lbl.gov/~spieler
- [3] J. Butler, Triggering and Data Acquisition General Considerations, *Instrumentation in Elementary Particle Physics, AIP Conf. Proc.* **674** (2003) 101–29
- [4] T. Kondo *et al.*, Construction and Performance of the ATLAS Silicon Microstrip Barrel Modules, *Nucl. Instr. Meth.* **A485** (2002) 27–42
- [5] I. Kipnis, H. Spieler & T. Collins, A Bipolar Analog Front-End Integrated Circuit for the SDC Silicon Tracker, *IEEE Trans. Nucl. Sci.* **NS-41(4)** (1994) 1095–103
- [6] S. Ramo, Currents Induced by Electron Motion, *Proc. IRE* **27** (1939) 584–5
- [7] F.S. Goulding, Pulse Shaping in Low-Noise Nuclear Amplifiers: A Physical Approach to Noise Analysis, *Nucl. Instr. Meth.* **100** (1972) 493–504
- [8] F.S. Goulding & D.A. Landis, Signal Processing for Semiconductor Detectors, *IEEE Trans. Nucl. Sci.* **NS-29(3)** (1982) 1125–41
- [9] V. Radeka, Trapezoidal Filtering of Signals from Large Germanium Detectors at High Rates, *Nucl. Instr. Meth.* **99** (1972) 525–39
- [10] V. Radeka, Signal, Noise and Resolution in Position-Sensitive Detectors, *IEEE Trans. Nucl. Sci.* **NS-21** (1974) 51–64
- [11] H. Spieler, Fast Timing Methods for Semiconductor Detectors, *IEEE Trans. Nucl. Sci.* **NS-29(3)** (1982) 1142–58
- [12] D.H. Wilkinson, A Stable Ninety-Nine Channel Pulse Amplitude Analyser for Slow Counting, *Proc. Cambridge Phil. Soc.* **46(3)** (1950) 508–18
- [13] M. Garcia-Sciveres *et al.*, The SVX3D Integrated Circuit for Dead-Timeless Silicon Strip Readout, *Nucl. Instr. Meth.* **A435** (1999) 58–64
- [14] H.W. Ott, *Noise Reduction Techniques in Electronic Systems*, 2nd edition, Wiley, New York (1988) ISBN 0-471-85068-3, TK7867.5.087
- [15] H. Johnson & M. Graham, *High-Speed Digital Design*, Prentice-Hall PTR, Upper Saddle River (1993) ISBN 0-13-395724-1, TK7868.D5J635
- [16] H. Johnson & M. Graham, *High-Speed Signal Propagation*, Prentice-Hall PTR, Upper Saddle River (2002) ISBN 0-13-084408-X, TK5103.15.J64