PART IV

RELATED CIRCUITS AND DEVICES (DATA-ACQUISITION-PERIPHERAL COMPONENTS)

Chapter Eighteen

Sample-Hold Circuits

A sample-hold amplifier (SHA) is a device that has a signal input, output, and a control input. As its name implies, a SHA has two steady-state operating modes. In the *sample* (or *track*) mode, the output tracks the input as faithfully as possible until the hold command is applied at the control input. In the *hold* mode, the output retains the last value of the input signal that it had at the time the Hold command was applied.

In data-acquisition systems, sample-hold amplifiers (SHA) are frequently required to "freeze" fast-moving signals prior to their processing by the system. Accurately holding the amplitude of the signal for an appropriate length of time is critical in measurement systems involving analog-digital conversion, peak detection, multiplexing, and other functions where precise timing is of the essence.

The control input is usually TTL- or ECL-logic-compatible; one logic state initiates and maintains the Sample or Track mode; the other does the same for the Hold condition. Most track-holds and sample-holds are identical in both function and circuit implementation. The only distinction between them is how they are used in the system. A *sample-and-hold* implies that the device samples the input for a short time and stays in the hold mode for the duration of the duty cycle. A *track-and-hold*, on the other hand, spends most of the time tracking the input and is switched into the hold mode for only brief intervals.

In data-acquisition systems operating at high update rates (greater than 1 MHz), the terms track-hold and sample-hold lose their distinction; for the purposes of this discussion, all such devices will be referred to as SHAs. The circuitry, characteristics, and applications of various types of SHAs will be discussed.

18.1 SAMPLE-HOLD OPERATION

Regardless of the circuit details or type of SHA in question, all such devices have four major components. The input amplifier, energy storage device (hold capacitor), output buffer, and switching circuits are common to all SHAs, as shown in the typical configuration of Figure 18.1.

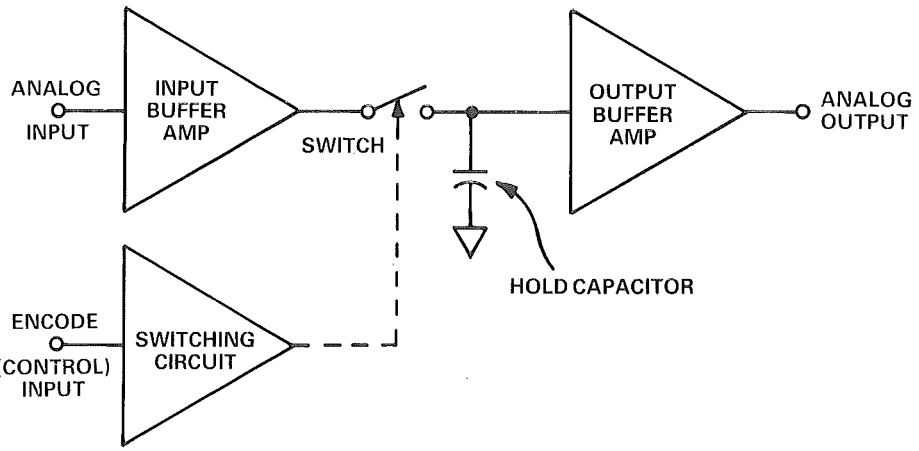

Figure 18.1. Basic sample-hold-amplifier structure.

The energy-storage device, the heart of the SHA, is almost always a capacitor. The input amplifier buffers the input by presenting a high impedance to the signal source and providing current gain to charge the hold capacitor. In the track mode, the hold capacitor usually determines the frequency response of the device; in the hold mode, the capacitor retains the voltage existing before it was disconnected from the input buffer. The output buffer offers a high impedance to the hold capacitor to keep the held voltage from discharging prematurely. The switching circuit and its driver form the mechanism by which the hold capacitor is alternately switched between Track and Hold.

18.2 SPECIFICATIONS

There are four groups of specifications that properly describe SHA operation. They are the *static* and *dynamic* characteristics that describe operation in and between the *track* and *hold* modes. Unique to SHAs are the dynamic specifications that describe the transitions from Track to Hold, and Hold to Track. An understanding of the terminology used to describe these devices is of key importance to the proper selection and use of SHAs. Figure 18.2 shows errors (exaggerated) during a complete cycle from Track to Hold and back.

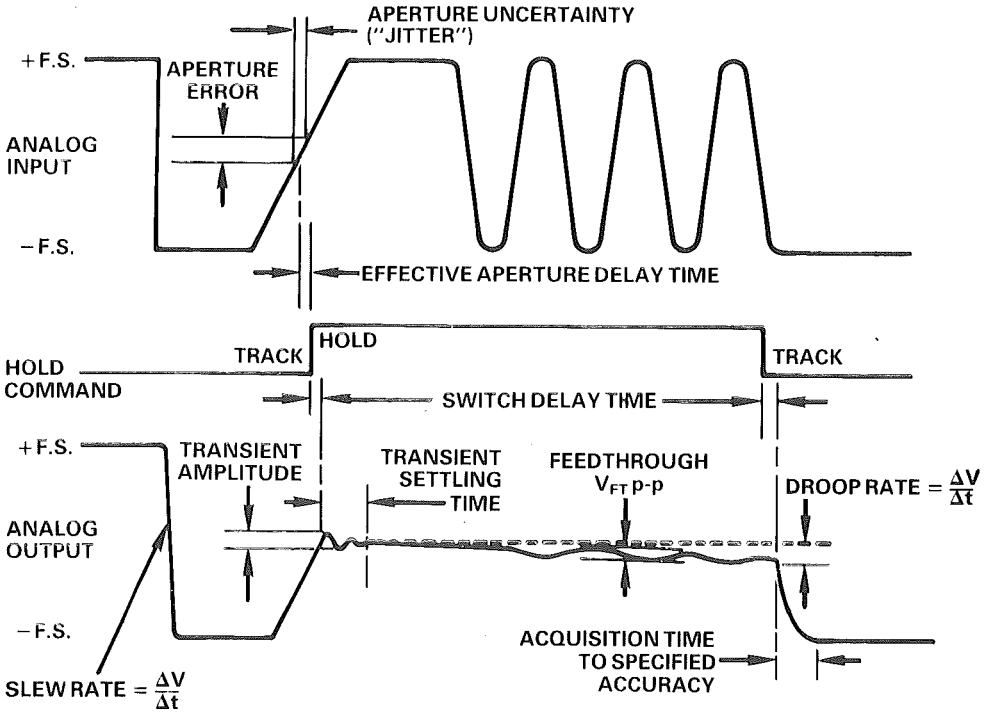

Figure 18.2. Sources of error (exaggerated) in a sample-hold.

18.2.1 TRACK MODE

While in the track—or sample—mode of operation, the SHA is simply a limited-bandwidth amplifier that may or may not provide gain. Operation in this mode is described by the same specifications that are used to characterize any analog amplifier. Indeed, many SHAs are little more than an operational amplifier with a capacitor and a switch. The principal specifications used to describe sample-mode operation are:

Offset: For zero input, the extent to which the output deviates from zero, over time and temperature.

Nonlinearity: The amount by which the plot of output vs. input deviates from a "best straight line." It is usually specified as a percentage of full-scale.

Gain: The multiplication factor describing the input to output "dc" transfer function.

Settling Time: The time required for the output to attain its final value within a specified fraction of full scale when a full-scale analog input step is applied, either as an input step or by a hold-to-sample transition (0 to ±FS, −FS to +FS, or +FS to −FS).

Bandwidth: Describes the frequency response in terms of output attenuation over frequency; it is usually characterized by the −3-dB value.

Slew(ing) Rate: The maximum rate of change of the output voltage when an analog step is applied, either as an input step or by a hold-to-sample transition.

18.2.2 TRACK-TO-HOLD TRANSITION

The specifications used to describe the transition from Track to Hold and from Hold to Track seem to be the most confusing to users of SHAs. These terms are unique to SHA devices and deserve special attention. Perhaps the most misunderstood and misused specifications are those that include the word "aperture" in their definition.

Aperture Time: The most essential dynamic property of a SHA is its ability to disconnect quickly the hold capacitor from the input buffer amplifier. The short—but non-zero—interval required for this action is called *aperture time*. The actual value of voltage that gets held at the end of this interval is a function of both the input signal and the errors introduced by the switching operation itself.

Sample-to-Hold Offset or Pedestal: There is a step error, which causes the held value to differ from the last value in *sample*; it is caused by charge dumped onto the hold capacitor via stray capacitance from the switch-control circuit. A design objective is to keep this error, resulting from a non-ideal switch, independent of the input signal level; the degree to which it in fact deviates from a constant over the input signal range can result in nonlinearities in the output with respect to the input signal. The constant offset portion of the error, called *charge transfer*, or *offset step*, can be compensated by coupling a signal of opposite phase onto the hold capacitor through an auxiliary switching circuit and compensation capacitor.

Thus, the *sample-to-hold-offset*, or *pedestal*, specification depends on the actual device configuration. For SHAs having fixed internal hold capacitance, it includes the residual uncorrected step and the *offset nonlinearity*. For integrated-circuit SHAs, requiring external capacitors, it is the residual step error after the *charge transfer* is accounted for and/or cancelled. In a device for which the capacitance can be chosen by the user, these effects can be reduced by increasing the capacitance in proportion, but at the cost of increased acquisition time.

Aperture Delay (Or, more descriptively, *Effective Aperture Delay Time)*: This specification is important because it helps the SHA user know when to strobe the device with respect to the input-signal timing. Figure 18.3 shows the sequence of what happens when the hold command is applied with an input signal of arbitrary slope (for clarity, the sample-to-hold offset error and switching transients are ignored). The value that finally gets held is a delayed version of the input signal, averaged over the aperture time of the switch. Effective Aperture Delay Time is defined as the interval between the leading edge of

the hold command and the instant when the *input* signal was equal to the held value. This is a more useful specification than aperture time alone, because it includes the effects of the analog and digital propagation delays, as well as the aperture time.

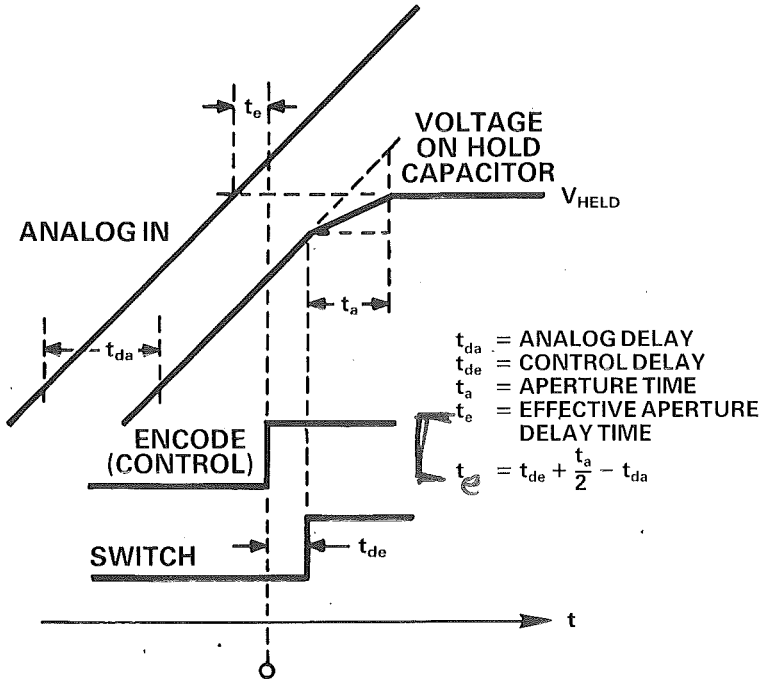

Figure 18.3. Internal sample-hold-amplifier timing.

The formula in Figure 18.3 is based on the assumption that the value of voltage on the hold capacitor is approximately equal to the average value of the signal applied to the switch over the interval during which the switch changes from a low to high impedance (aperture time).

Aperture Uncertainty (Jitter): Aperture uncertainty, or "jitter", is the result of noise which modulates the phase of the hold command. This jitter shows up as a sample-to-sample variation in the value of the analog signal which is being "frozen".

Aperture uncertainty manifests itself as an aperture error, as shown in Figure 18.2. The amplitude of the error is related to the rate-of-change of the analog input. For any given value of aperture uncertainty, aperture error will increase as the input dV/dt increases.

Switching Transient: As shown in Figure 18.2, most SHA specifications include the maximum amplitude and duration of the transient that appears at the output as a result of the sample-to-hold transition. A similar transient appears during the hold-to-sample transition but is usually not noticeable because of the dominant effect of acquisition time.

Switch Delay Time: The interval between the edges of the hold command and the beginning of change of state at the analog output. This delay, as shown in Figure 18.2, occurs at both the sample-to-hold and hold-to-sample transition.

18.2.3 HOLD MODE

During the hold mode there are errors due to imperfections in the switch, the output amplifier, and the hold capacitor.

Droop: A constant drift of the output voltage due to charge leakage from the hold capacitor through the switch, output buffer, circuit board or substrate—or within the capacitor itself ($dV/dt = I/C$). This error can be reduced by increasing the hold capacitance (at the cost of increased acquisition time) and/or reducing leakage currents by component choice, component placement, and shielding or guarding.

Feedthrough: The fraction of input signal that appears at the output in Hold, caused primarily by capacitance across the switch. Usually measured by applying a full-scale sinusoidal input at a fixed frequency (e.g., 20 V peak-to-peak at 10 kHz) and observing the output during Hold.

Dielectric Absorption: The tendency of charges within a capacitor to redistribute themselves over a period of time, resulting in "creep" to a new level when allowed to rest after large, fast changes. This effect, less than 0.01% for good polystyrene and teflon capacitors, can be as large as several percent for ceramic and mylar capacitors.

18.2.4 HOLD-TO-SAMPLE TRANSITION

Acquisition Time: The length of time during which the SHA must remain in the sample mode in order for the hold capacitor to acquire a full-scale step input; adequate acquisition time makes it possible for the subsequent hold-mode output to be within a specified error band of the final value.

Acquisition time is a key SHA dynamic specification. The maximum sample rate of any SHA is limited by the sum of the time intervals required for the sample and the hold modes. The interval spent in the hold mode (after transients have settled) is primarily determined by the system in which the SHA is used. The minimum time spent in the Sample mode, however, is established by the sample-hold's acquisition time for a given degree of accuracy.

When acquisition time is measured as the interval from the hold-to-sample transition to the instant when the output buffer has settled, the resulting value of acquisition time is generally pessimistic. As defined, acquisition time measures the time required to acquire the signal *at the hold capacitor, not necessarily at the output of the buffer* (unless they are the same). For devices in which the hold capacitor is buffered by a follower, rather than used as an integrator (see the next Section), a measurement at the output includes output buffer settling

time as well as switch delay time. In practice, the voltage has already settled at the hold capacitor, and a Hold command can be applied before the output buffer completely settles.

Figure 18.4 is a sketch of waveforms in a method of determining acquisition time that is independent of output-buffer effects. The input analog square wave is sampled and converted at twice the analog square-wave frequency with relatively narrow Sample pulses. The Figure indicates that, on the transition to Sample, the output starts to change, in order to follow the step input. Initially set wide enough for the signal to be accurately acquired, the pulsewidth, t_{aq}, is reduced until the output waveform, measured digitally, begins to collapse to some defined percentage of full scale. When that occurs, the width of the Sample pulse is equal to the acquisition time, because the hold capacitor can no longer acquire the signal accurately with further reduction in t_{aq}.

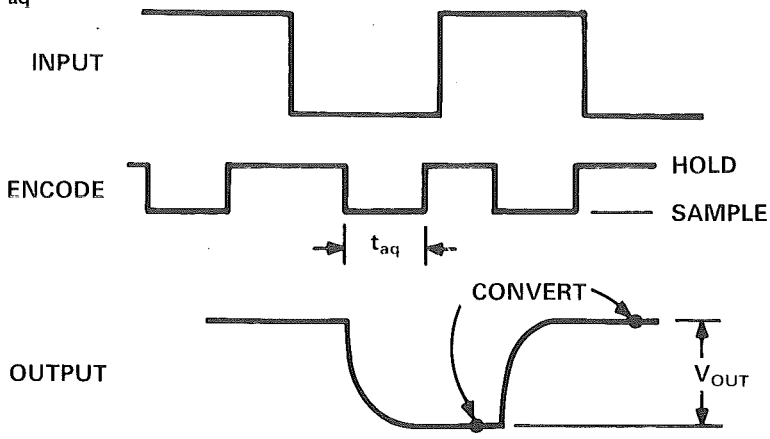

Figure 18.4. Acquisition-time measurement waveforms.

18.3 SAMPLE-HOLD CIRCUITS

Most SHA designs fall into one of two categories, open- or closed-loop circuits. Closed-loop SHAs exploit the accuracy, low drift, and gain flexibility available with operational amplifiers. Open-loop designs take advantage of the high-speed characteristics of unity-gain buffer amplifiers.

18.3.1 OPEN-LOOP CIRCUITS

Figure 18.5 shows the conceptually simplest SHA circuit. When the switch is closed, the capacitor charges exponentially to the input voltage, and the amplifier's output follows the capacitor's voltage. When the switch is opened, the charge remains on the capacitor. The capacitor's acquisition time depends on the series resistance and the current available to charge the capacitor. Once the charge is acquired with the appropriate accuracy, the switch can be opened, even though the amplifier has not yet settled, without affecting the final output value.

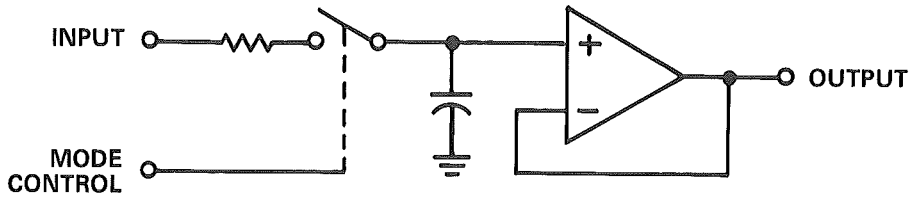

Figure 18.5. SHA structure – simple follower.

A disadvantage of this circuit is that the switched capacitor dynamically loads the input source, which may not have low-enough output impedance and sufficient current-drive capacity. The circuit of Figure 18.6 is similar but includes an input buffer amplifier to isolate the source. The Analog Devices HTS-0025 and HTS-0010 SHAs use this scheme to achieve acquisition times as low as 10 ns. These designs utilize a high speed diode switching bridge for sampling wide-bandwidth signals at update rates of up to 50 MHz. The high speeds are achieved by employing buffer amplifiers that do not use voltage feedback; the nonlinearities of these amplifiers limit their use to systems requiring resolutions of 12 bits or less.

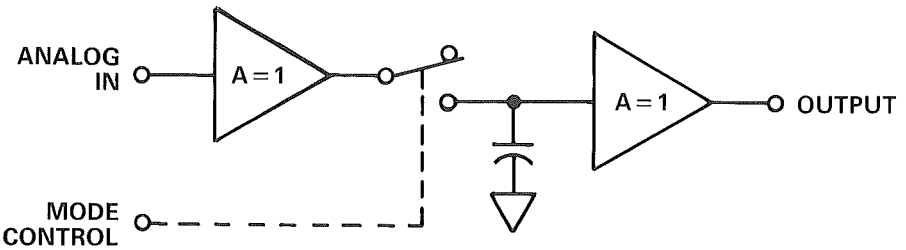

Figure 18.6. SHA structure – follower with input buffer.

The HTS-0010 is an example of a SHA that has the hold capacitor's active terminal brought out so that additional capacitance can be employed to improve the droop rate and sample-to-hold offset (at the expense of bandwidth and acquisition time, as noted earlier).

SHA specifications and terminology are generally applicable to both open-loop and closed-loop SHAs. Applications requiring update rates higher than 10 MHz and acquisition times less than 100 ns usually require open-loop circuits. At these speeds, subtle distinctions in terminology can mislead the system designer. Before selecting a high-speed SHA, the user should have a clear understanding of the intent and meaning of the manufacturer's usage of such specifications as aperture time and acquisition time.

18.3.2 CLOSED-LOOP CIRCUITS

The circuits of Figures 18.5 and 18.6 have the essential advantage of potentially fast acquisition and settling time because they are open-loop devices.

If low-frequency tracking accuracy is more important than speed, this can be accomplished by closing the loop around the storage capacitor and using high loop gain to enforce tracking accuracy.

Figure 18.7 shows a configuration in which the input follower of Figure 18.6 is replaced by a high-gain difference amplifier. Now, when the switch is closed, the output (which represents the charge on the capacitor) is forced to track the input, within the capability of input amplifier's gain, bandwidth, common-mode error, and current-driving capabilities. SHAs using this circuit configuration to achieve higher accuracy than that available in open loop circuits include the Analog Devices AD582, AD583, and ADSHC-85.

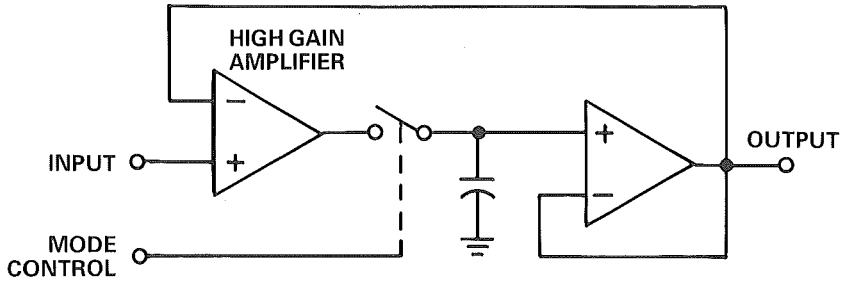

Figure 18.7. SHA structure – feedback circuit with output buffer follower.

In Figure 18.8 (which illustrates the architecture of the closed-loop AD585 and SHA1144), an integrator configuration is used to charge the capacitor, permitting the switch to operate at ground potential; this simplifies leakage problems.

Because both the output and the input affect the charge on the capacitor, the acquisition time and the settling time are identical in the circuits of both Figure 18.7 and Figure 18.8. If the circuit of Figure 18.7 is switched into Hold before the output has settled at its final value, the sample may be in error. In addition, since the loop is open during Hold, the input stage must reacquire the input when the circuit is returned to Sample, even if the input

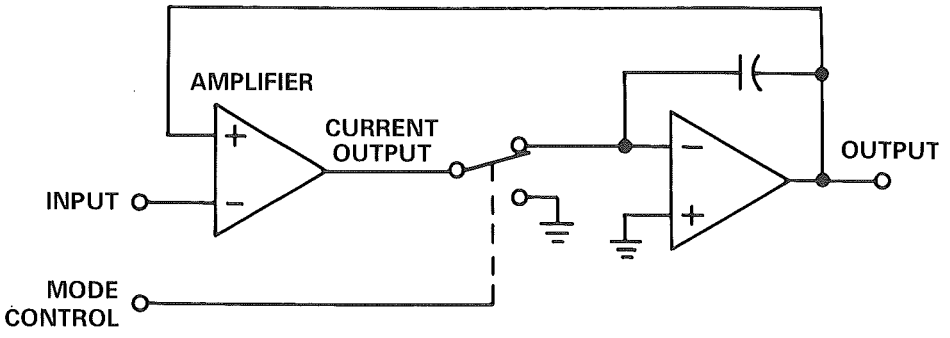

Figure 18.8. SHA structure – feedback circuit with output integrator.

is unchanged. As a rule, this will result in a spike if the input amplifier has high voltage gain.

If input impedance and acquisition time are not critical, and sufficient drive current is available from the source, the circuit of Figure 18.9 may be desirable. Here, only a single operational amplifier is required; and the input impedance and gain are a function of the choice of R_I and R_F. In both Sample and Hold, the input impedance to ground is R_I. In the Sample mode, the input circuit sees a virtual ground through R_I, and the hold capacitor is charged by the amplifier. In the Hold mode, the resistance node is switched to analog ground potential to disconnect the capacitor while minimizing signal feedthrough and maintaining a constant input impedance. The Analog Devices AD346, AD389, HTC-0300, and HTC-0500 all utilize this principle.

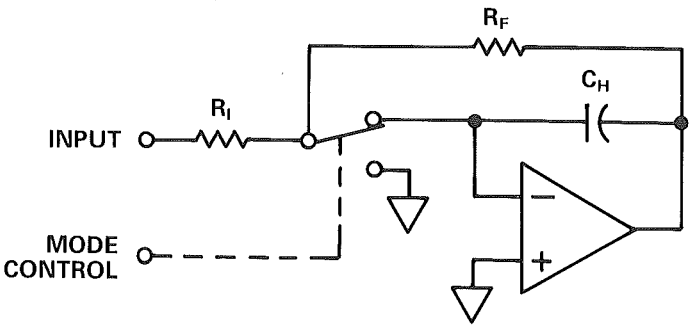

Figure 18.9. SHA structure – inverting integrator switched at summing point.

18.4 APPLICATIONS

Sample-holds are most widely used in data acquisition systems, typically as shown in Figure 18.10. The sample-hold maintains the input to the a/d con-

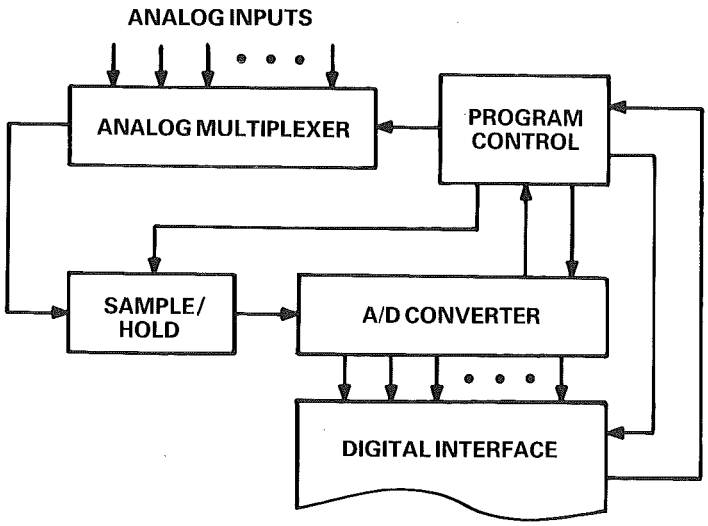

Figure 18.10. Typical data-acquisition system.

verter constant during the conversion interval; meanwhile, the multiplexer is seeking the next channel to be converted, either randomly or sequentially. As soon as conversion is completed, the sample-hold samples the newly established input, and the cycle is repeated.

This mode of operation is known as synchronous sampling; the sample-hold operates in synchronism with the other system elements. In another mode (viz., asynchronous), a large number of sample-holds may be used to acquire and store data at rates pertinent to each individual channel. They are then either interrogated by analog multiplexers, or the signals are individually converted asynchronously, and then multiplexed digitally.

In data distribution, 0.01% sample-holds may be less costly than large numbers of D/A converters having comparable accuracy. A typical data distribution system is shown in Figure 18.11. A fast, accurate D/A converter updates a number of sample-holds at speed and accuracy levels appropriate to the individual channels.

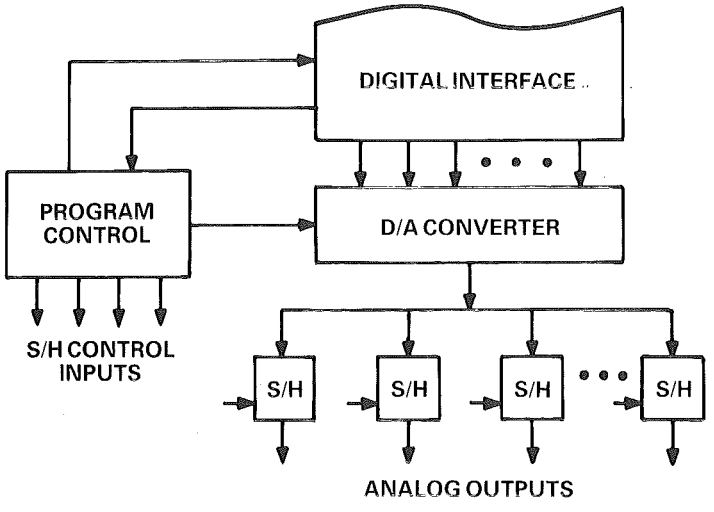

Figure 18.11. Data distribution system with analog storage.

Sample-holds are also used to "deglitch" d/a converters in systems that are sensitive to the D/A glitches that occur at transition points. Figure 18.12 shown the timing relationship used when deglitching DACs. Just prior to latching the new data into the d/a converter, the SHA is put into the hold mode, so that D/A glitches are isolated from the output. SHAs used as deglitchers must have very small sample-to-hold and hold-to-sample transients and pedestal errors, as well as fast acquisition times.

Undoubtedly, the principal usage of SHAs is ahead of analog-to-digital converters. "Aperture time," in an a/d converter system, differs from the SHA specification. In a/d conversion, *aperture time* refers to the period of time over which the analog input must remain stable in order for an accurate conversion

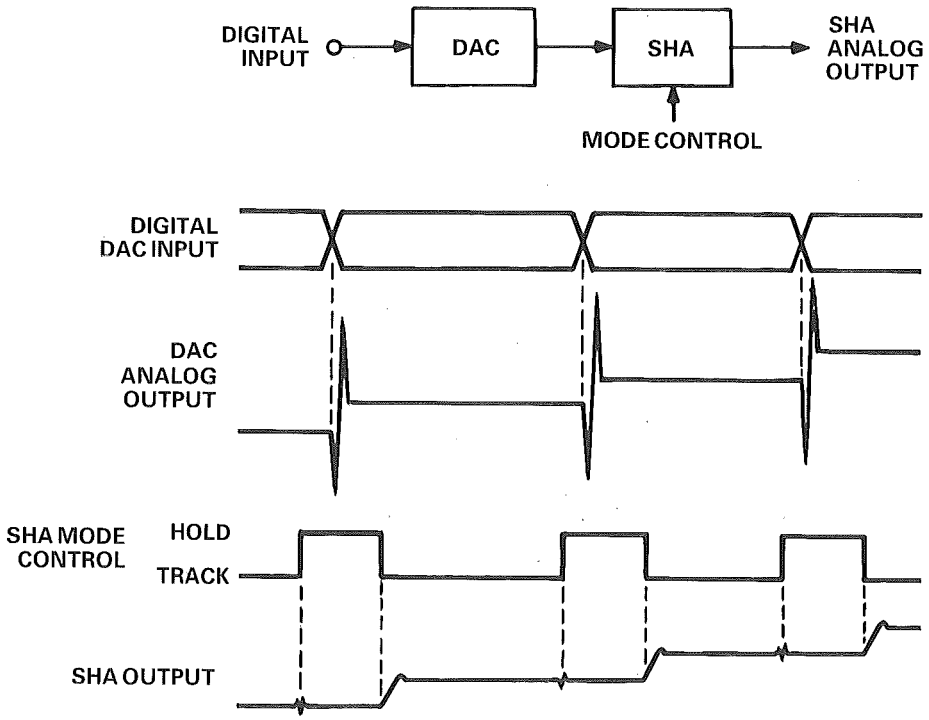

Figure 18.12. Deglitching the output of a d/a converter.

to occur. It is this interval which limits the analog bandwidth (maximum dV/dt) of the A/D converter.

The purpose of the SHA is, in effect, to reduce the aperture time of the a/d converter. By holding the rapidly changing analog input for the period required by the ADC, the analog bandwidth of the system can be increased substantially. In addition, accuracy and linearity—even of low-frequency signals—are enhanced. See Chapters 2 and 13 for discussions of this point.

Most successive-approximation and subranging-type a/d converters require a SHA for all but the lowest analog bandwidths. For this reason, some converters are offered with the SHA incorporated into the design. This makes available to the system designer an ADC-SHA combination in which the timing has already been optimized for the best performance. On the other hand, many state-of-the-art flash type a/d converters offer extremely low aperture times; the inclusion of a SHA with such devices can improve high-frequency performance but may degrade low-frequency performance because of droop associated with the necessarily small capacitors.

Another application of the SHA is as a peak detector. A typical example is shown in Figure 18.13.

When the input is greater than the SHA output, the comparator's positive output causes the SHA to track. When the input decreases and becomes less

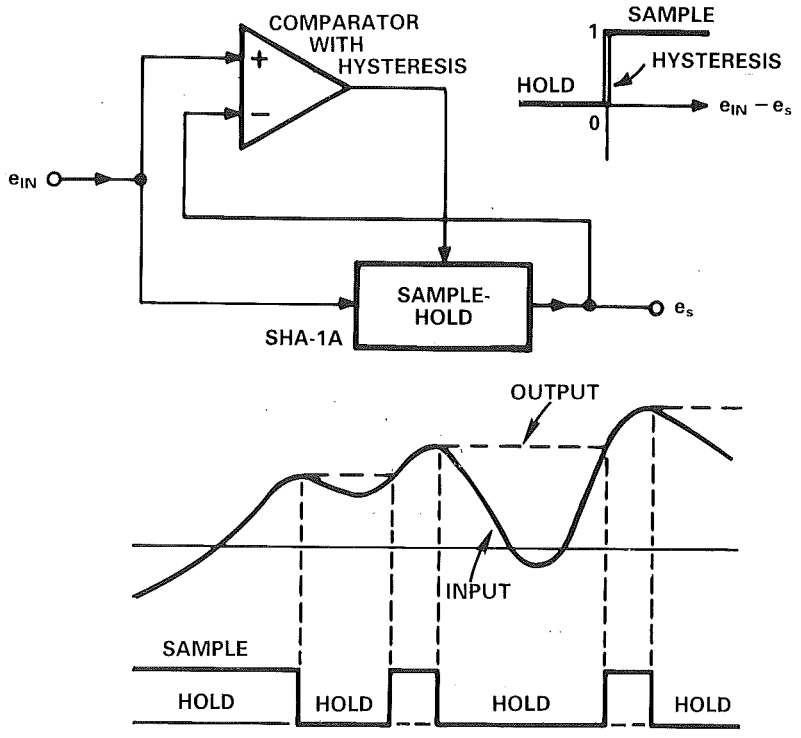

Figure 18.13. Peak follower, using a SHA and a comparator.

than the SHA output, the comparator's "0" output causes the SHA to hold until the input again becomes greater than the output. To reset, the control input is arbitrarily switched into Sample, and the lowest expected level is applied at the input. The sample-hold output (or the comparator input) is biased by a few millivolts of hysteresis to avoid ambiguity during step inputs and minimize false triggering by noise.

18.5 SELECTING AND USING SHAS

In a data-acquisition system, the point where the analog sample is taken—the sample-hold—is often the weakest link, because many system designers overlook the challenging nature of the requirements on both its static and dynamic performance. As an example of the variety of performance characteristics, Table 18.1 lists a number of SHAs available from Analog Devices at the end of 1984 and some of their critical specifications.

In those devices that require an external hold capacitor, the charge transfer is given in picocoulombs (pC). The actual sample-to-hold offset error is given by:

$$V_{error} = \frac{Q}{C} = \frac{\text{Charge Transfer (pC)}}{\text{Hold Capacitance (pF)}}$$

In choosing capacitors for devices with this option, the sample-to-hold (pedestal) error is only one of the specifications to consider. The settling time, droop rate, bandwidth, and acquisition time are all functions of the hold capacitance.

When comparing SHA specifications, be sure to take into account the conditions under which the measurements are made. For example, settling-time measurements made while driving a large capacitive load will rarely be as good as manufacturers' specifications, based on measurements with a simple resistive load. In short, system designs incorporating SHAs require consideration of the associated circuits and layout, in addition to understanding and appreciation of the device's specifications.

With all this for the user to consider, it is little wonder that many manufacturers offer ADCs and SHAs packaged (and specified) together. This is another dimension of choice for the system designer: the advantages of an ADC, complete with SHA; or the flexibility offered by the immense variety of individually available SHAs and ADCs.

	Acquisition Time	Droop Rate or Current	Aperture Uncertainty (ps, rms)	Pedestal Error (Voltage or Charge)	Offset vs. Temp.	Nonlinearity	Technology
HTS-0010	14 ns (0.1%)	100 μV/μs	5	5mV	125μV/°C	0.1%	Hybrid
HTS-0025	25 ns (0.1%)	200 μV/μs	20	5 mV	100 μV/°C	0.1%	Hybrid
HTC-0300A	150 ns (0.1%)	5 μV/μs	100	5 mV	100 μV/°C	0.01%	Hybrid
HTC-0500	700 ns (0.1%)	0.5 μV/μs	60	5 mV	100 μV/°C	0.01%	Hybrid
AD346	1 μs (0.01%)	0.1 mV/ms	400	10 mV	–	–	Hybrid
AD389	1.5 μs (0.01%)	0.1 μV/μs	400	2 mV	–	0.001%	Hybrid
AD585	3 μs (0.01%)	1 mV/ms	500	0.3 pC	–	–	Monolithic
AD583	4 μs (0.1%)	5 pA	5,000	10 pC	–	–	Monolithic
ADSHC-85	4.5 μs (0.01%)	0.2 mV/ms	500	1 mV	25 μV/°C	0.01%	Hybrid
SHA1144	6 μs (0.003%)	1μV/μs	500	1mV	30μV/°C	0.001%	Module
AD582	6 μs (0.1%)	100 pA	15,000	5 pC	–	0.01%	Monolithic

TABLE 18.1. Brief Summary of Sample-Hold Performance

Chapter Nineteen

Analog Switching and Multiplexing

The semiconductor analog switch has become an essential and ubiquitous component in the design of many electronic systems. As an integrated-circuit electronic building block, it has matured greatly since the early days when switches were constructed using discrete semiconductors.

The electromechanical relay continues to be widely used for signal switching. But the many advantages associated with semiconductor analog switches, particularly with regard to their high speed of operation, reliability, and ease of interfacing to microprocessors, have created significant new roles and markets, which are effectively barred to electromechanical devices.

Table 19.1 is a brief comparison of the advantages and disadvantages of relays and analog switches; it serves to indicate the suitability of the two categories to various application areas in analog signal handling.

19.1 POPULAR IC SWITCH PROCESSES

Several circuit and process techniques are used to fabricate analog switches. The most common of these are JFET, PMOS and CMOS. Following a very brief discussion of the three approaches, this chapter will concentrate on the characteristics and applications of *CMOS* switches, which are becoming widely used in multiplexing and general-purpose switching.

19.1.1 JUNCTION FIELD-EFFECT TRANSISTOR (JFET) TYPES.

JFET switches exhibit constant ON resistance with varying signal voltage. This makes JFET devices particularly suitable for low-distortion signal switching.

RELAYS

ADVANTAGES

- Low ON resistance
- High OFF resistance
- Galvanic isolation
- Switch characteristics are relatively temperature independent

DISADVANTAGES

- Limited number of operations
- R_{ON} and R_{OFF} deteriorate throughout the usable life of the device
- High power dissipation
- Large physical size/weight
- Slow (1-millisecond t_{ON})
- Prone to "bounce"
- High cost per channel
- Not compatible with standard logic (TTL, CMOS)
- Acoustically noisy

ANALOG SWITCHES

ADVANTAGES

- Fast (100-ns t_{ON})
- TTL/CMOS compatible
- No switch "bounce"
- No channel degradation during switch life
- Virtually unlimited number of switch operations
- Small size and weight
- Low power dissipation
- Low cost per channel
- Rugged construction
- Acoustically quiet

DISADVANTAGES

- Relatively High R_{ON}
- Potentially Lower R_{OFF} than relays
- Switch characteristics are temperature dependent
- Lack of galvanic isolation

Table 19.1 Electronic analog switches vs. relays

However, in order for an n-channel JFET switch to be fully turned off, it requires a gate voltage that is more negative than the most negative value of signal voltage to be switched, by at least V_p volts (where V_p is the pinchoff voltage of the FET). The negative supply of the JFET switch driver provides this negative voltage; this means that only signals that are more positive than $(-V_s + V_p)$, where V_s is the negative supply voltage, can be switched.

Figure 19.1 illustrates the range of signal voltages that can be handled by a typical JFET analog switch.

19.1.2 P-TYPE METAL-OXIDE-SEMICONDUCTORS (PMOS)

Analog switches employing PMOS technology are relatively simple to fabricate because of the low number of masking stages necessary during fabrication.

However, the PMOS switch channel exhibits a large variation in ON-resistance with changing signal voltage; this can cause unacceptable degradation of circuit performance.

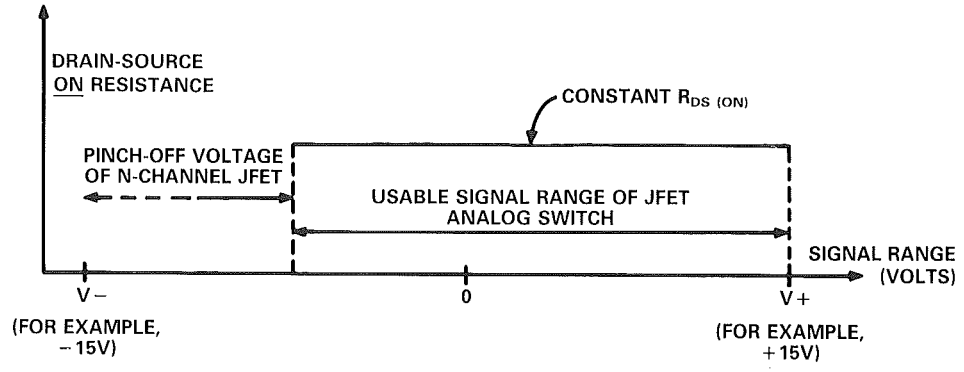

Figure 19.1. Voltage-range limitation of junction FET switches.

For a specified channel ON-resistance, PMOS switches occupy much larger chip areas than either NMOS or CMOS alternatives. Thus, junction capacitances, which are proportional to junction area, are relatively large, making PMOS switches less suitable for high frequency switching.

19.1.3 COMPLEMENTARY METAL-OXIDE-SEMICONDUCTORS

When a PMOS switch is paralleled with an NMOS switch, we arrive at another switch technology—CMOS. CMOS is by far the most popular process used for the fabrication of analog switches. CMOS switches have extremely low quiescent power dissipation, require little drive or supply current while switching, and are low in cost. Their ON resistance is low and varies slightly (typically 10%) with applied voltage. In the OFF condition, leakage is quite small, both across the gate and to the drive and supply circuits. Most types respond to TTL/DTL—as well as CMOS—logic levels.

Like most other processes, CMOS has undergone many and diverse developments over the past decade. Many CMOS process variations have been developed to fulfill the parametric needs of analog switches. These process variations can be conveniently divided into two main types: *junction-isolated CMOS* and *dielectrically isolated CMOS*.

All the various CMOS fabrication technologies make it possible to have low ON-resistance variation with changing signal voltage. Figure 19.2 shows the ON-resistance profiles of separate NMOS and PMOS devices and the effect on R_{ON} when paralleling the two complementary device types to form a CMOS channel.

Junction-Isolated CMOS

This relies on reverse-biased p-n junctions to form the electrical isolation between different devices on the same chip. While junction isolation provides effective isolation under normal conditions, fault conditions can cause junc-

tions to link together to form parasitic bipolar transistors within the CMOS structure; under certain conditions, they can provide the latched low-resistance path that is characteristic of a silicon controlled-rectifier (SCR). While the device behaves normally with specified signal and control voltages, reversed supplies or overvoltage may trigger the parasitic SCR, resulting in high and potentially destructive fault currents in an unprotected device. The phenomenon is known as "latchup."

The reverse-biased isolating junctions exhibit capacitance proportional to junction area; this capacitance tends to degrade high-frequency switch performance. The junction capacitance also induces ac coupling between the switch input/output terminals and the device power supplies, which are effectively at ac ground. This can result in increased insertion loss, reduced OFF-isolation, and susceptibility to high-frequency pickup from the supply leads.

Dielectrically Isolated CMOS

Each DI device is diffused in its own separate pocket of silicon surrounded by an isolating dielectric layer, usually silicon dioxide. The advantages of DI are threefold:

• It allows the IC designer to incorporate devices on-chip which would be difficult, if not impossible, to include using a junction-isolated process. Notable examples of this are current-limiting resistors, which are useful in protecting against input overvoltage.

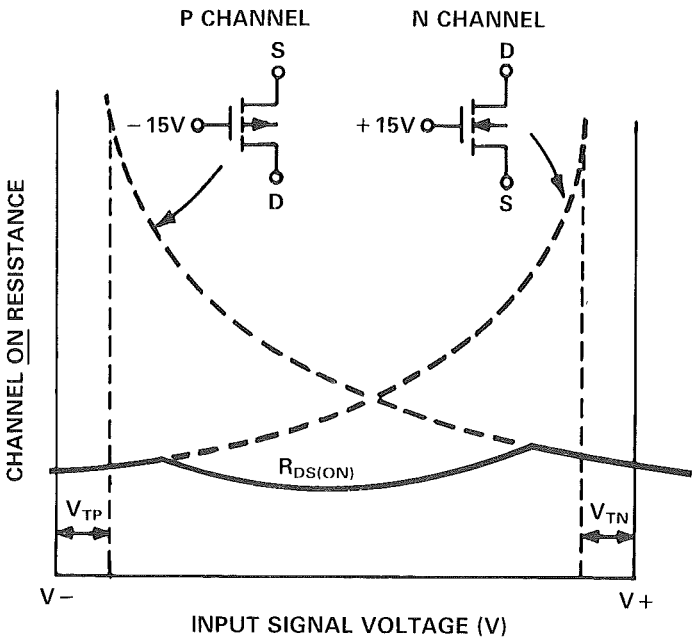

Figure 19.2. ON-resistance profiles of MOS switches.

• It eliminates the parasitic SCRs that can be formed within junction-isolated processes, thus completely avoiding device latchup.

• It reduces circuit parasitic capacitances because of the significant reduction of reverse-biased junction area.

19.1.4 SWITCH FORMS

Analog switches are available in three basic functional forms:

- Single-Pole-Single-Throw—SPST (Figure 19.3a)
- Single-Pole-Double-Throw—SPDT (Figure 19.3b)
- Double-Pole-Single-Throw—DPST (Figure 19.3c)

Note that the DPST function can be made up using two SPST switch channels with their digital control inputs connected together.

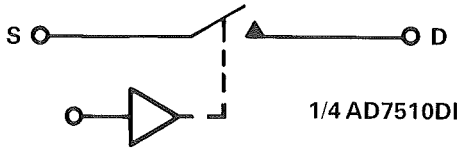

a. Single-pole, single-throw (SPST).

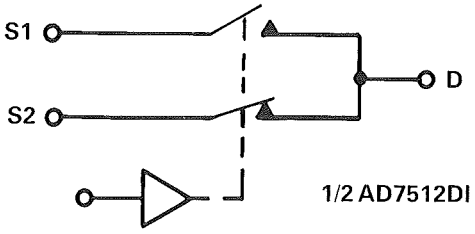

b. Single-pole, double-throw (SPDT).

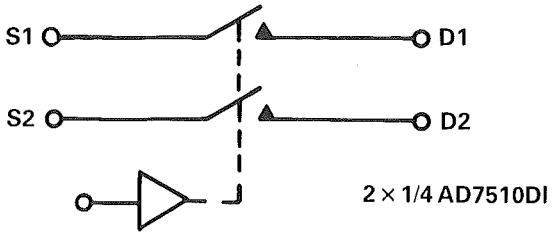

c. Double-pole, single-throw (DPST).

Figure 19.3. CMOS analog switch functions.

Figure 19.4 shows the basic schematic of a typical single-pole-single-throw switch channel. The switch driver includes level shifters, which translate the TTL or CMOS input logic levels into a pair of voltages—which swing between the supply rails—suitable for driving the gates of the two signal-carrying complementary MOSFETS.

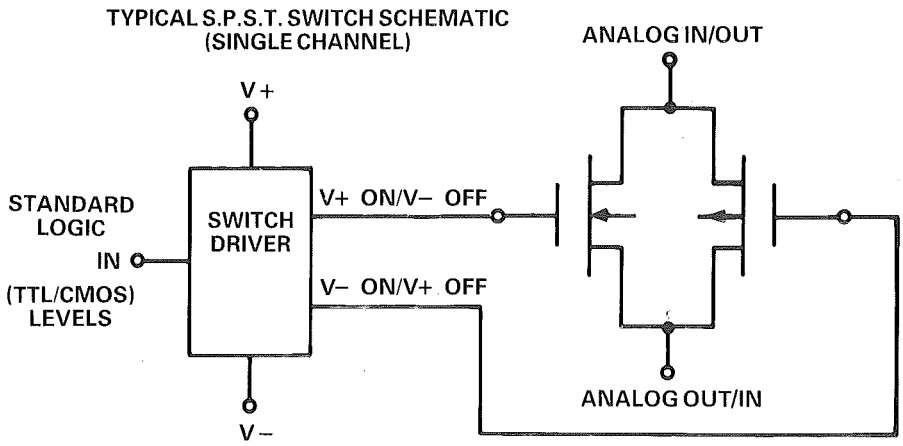

Figure 19.4. Single-pole, single throw CMOS analog switch.

The bodies (back gates) of the n-channel and p-channel FETS are usually connected to the negative and positive supply rails respectively. This ensures that the source-to-body and drain-to-body junctions remain reverse biased for all signal voltages within the supply-rail potentials.

19.2 SWITCH CHARACTERISTICS AND SPECIFICATIONS

There are many switch properties specified on switch and multiplexer data sheets. Key phenomena and specifications will be discussed in the sections that follow. A list of parameters, their abbreviations and definitions, appears in Tables 19.2 (switches) and 19.3 (multiplexers).

19.2.1 R_{ON} VARIATION WITH SIGNAL VOLTAGE

Although the resistance change is not large at higher supply voltages, the ON-resistance of CMOS switches does vary with changes in applied signal voltage. In Figure 19.5, the effect of this phenomenon in the AD7510 is illustrated for various supply voltages.

R_{ON} can introduce attenuation, and its variation with applied voltage can introduce distortion into the signal path. Figure 19.6 shows a typical circuit that exhibits this effect. The nominal 100-ohm resistance of R_{ON} introduces a 1% attenuation, which can be compensated for by a gain trim. However, the $\pm 10\%$ variations of R_{ON} with signal level will modulate the attenuation, introducing a nonlinearity of $\pm 0.1\%$. If the input and feedback resistance are

R_{DS}	Ohmic resistance between terminals D and S.
I_D, I_S	Current at terminals D or S. This is a leakage current when the switch is OFF.
I_{DS}	Current flowing through the closed switch.
$I_D - I_S$	Leakage current that flows from the closed switch into the body. (This leakage will show up as the difference between the current, I_D, going into the switch, and the outgoing current, I_S.
V_D, V_S	Analog voltage on terminal D or S.
C_S, C_D	Capacitance between terminal S or D and ground. (This capacitance is specified for the switch open and closed.)
C_{DS}	Capacitance between terminals D and S. (This will determine the switch OFF isolation as a function of frequency.)
C_{DD}, C_{SS}	Capacitance between terminals D or S of any two switches. (This will determine the cross-coupling between switches as a function of frequency.)
t_{ON}	Delay time between the 50% points of the digital input and switch ON condition.
t_{OFF}	Delay time between the 50% points of the digital input and switch OFF condition.
V_{INL}	Threshold voltage for the low state.
V_{INH}	Threshold voltage for the high state.
I_{INL}, I_{INH}	Input current of the digital input.
C_{IN}	Input capacitance to ground of the digital input.
V_{DD}	Most positive voltage supply.
V_{SS}	Most negative voltage supply.
I_{DD}	Positive supply current.
I_{SS}	Negative supply current.

Table 19.2 Switch Characteristics

sufficiently high, distortion effects can be neglected; for example, 100-kΩ resistors will reduce the nonlinearity to ±0.01%. However, excessively large resistance values can incur high noise levels. An acceptable trade-off between noise performance and distortion may not be possible using this circuit configuration.

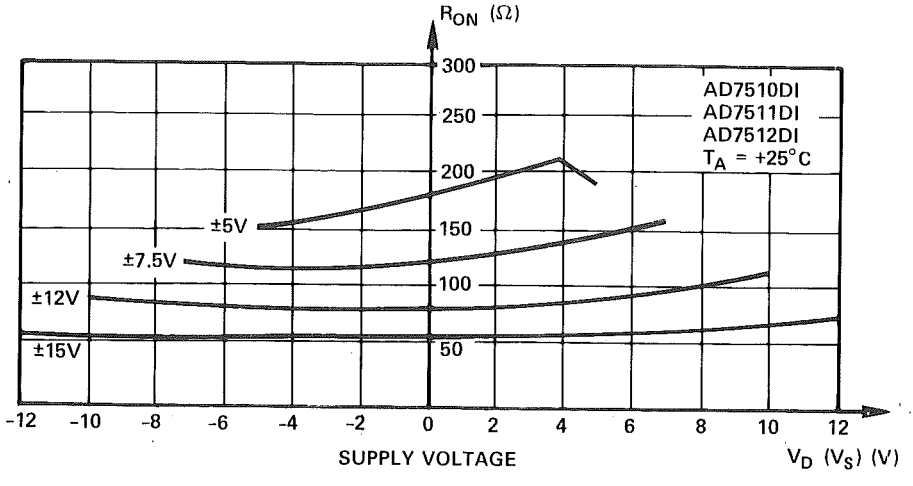

Figure 19.5. Switch ON resistance vs. V_D or V_S, as a function of supply voltage.

R_{ON}	Ohmic resistance between the output and an addressed input.
R_{ON} vs. Temperature	R_{ON} drift over the temperature range.
ΔR_{ON} between switches	Difference between the R_{ON} values of any two switches.
R_{ON} vs. Temperature between switches	Difference between the R_{ON} drifts of any two switches.
I_S	Current at any switch input, S1 through SN. This is a leakage current when the switch is open.
I_{OUT}	Current at the output. This is a leakage current when all switches are open.
$I_{OUT}-I_S$	Difference between the current flowing into terminal S and the current flowing out of terminal "out" when terminal S is addressed.
V_{INL}	Digital threshold voltage for the low state.
V_{INH}	Digital threshold voltage for the high state.
C_S	Capacitance between any open terminal, S, and ground.
C_{OUT}	Capacitance between the output terminal and ground with all switches open.
$C_{S\text{-}OUT}$	Capacitance between any open terminal, S, and the output terminal.
C_{SS}	Capacitance between any two "S" terminals.
$t_{transition}$	Delay time when switching from one address state to another.
t_{open}	"OFF" time of both switches when switching from one address state to another.
$t_{on(En)}$	Delay time between the 50% points of the enable input and the switch "ON" condition.
$t_{off(En)}$	Delay time between the 50% points of the enable input and the switch "OFF" condition.
V_{DD}	Most positive voltage supply.
V_{SS}	Most negative voltage supply.
I_{DD}	Positive supply current.
I_{SS}	Negative supply current.

Table 19.3 Multiplexer Characteristics

An alternative is to make the circuit insensitive to loading by eliminating or buffering the load resistance. In Figure 19.7, the switches are buffered by a non-inverting unity-gain stage. Distortion is substantially eliminated because the modulation of ON-resistance has no effect on the gain of the amplifier. Since the resistors have been eliminated, the noise effects associated with them in Figure 19.6, due to the noise they generate, as well as the voltage developed across them by the input current noise of the amplifier, are also eliminated.

Another way to improve the situation of Figure 19.6 is to connect the switches at the summing point of the unity-gain inverter, as shown in Figure 19.8. The signal voltage across the switch is kept small; thus, ON-resistance modulation effects due to signal-voltage variations are minimized, although the series resistance of the switches must still be accounted for. An interesting advantage of this circuit is that differing values of input resistance may be used to program different gains for the various input signals.

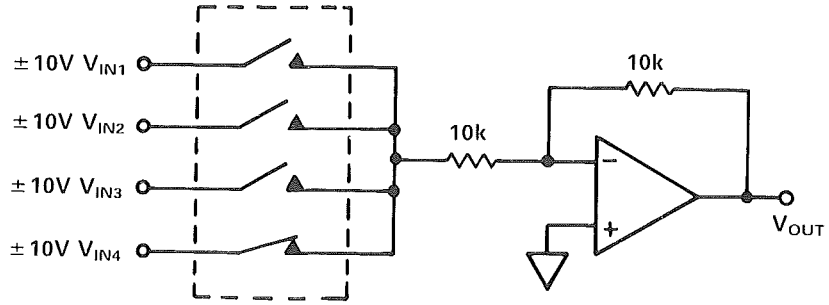

FOR $R_{ON} = 100\Omega \pm 10\%$

$$V_{OUT} = -V_{IN4}\left[\frac{10^4\Omega}{1.01 \times 10^4\Omega} \pm 0.1\%\right]$$

Figure 19.6. Quad switch selects an input for a unity-gain inverter.

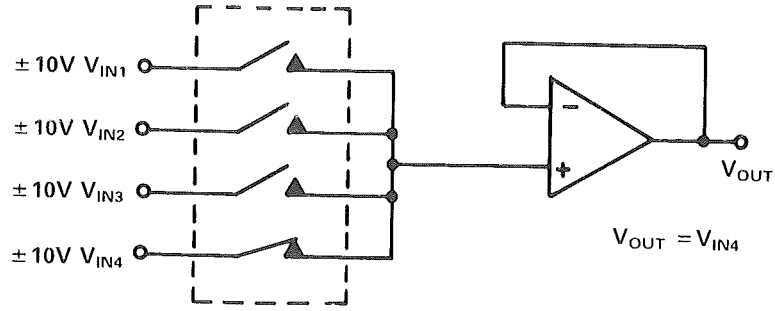

Figure 19.7. Switching the input of a follower-amplifier.

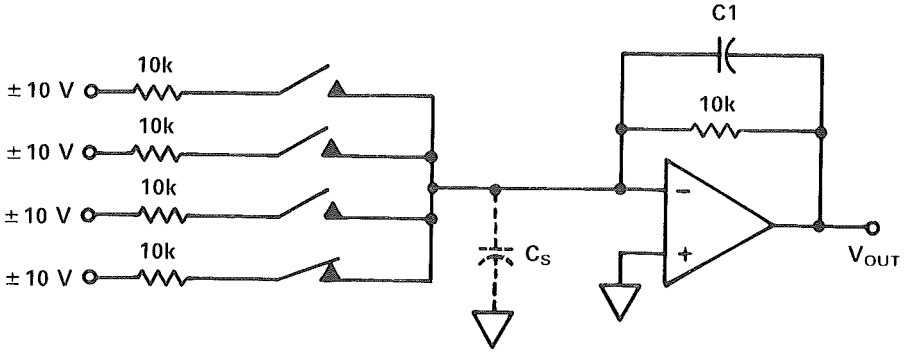

Figure 19.8. Switching at the summing point.

However, this configuration can introduce dynamic problems. First, the switch capacitance, C_S, which appears between the op-amp summing point and ground, can reduce phase margin. To compensate, capacitor C1, with approximately equal capacitance, is used to maintain stability—at the cost of reduced frequency range.

In addition, switching transients, coupled in by the capacitance between the switch drive and the switch points, will be coupled directly to the amplifier summing point, introducing output noise. This is not as much of a problem in Figures 19.6 and 19.7, because the coupled charge is dumped into the (usually low) output impedance of the connected signal source.

Since the switch resistances tend to be approximately equal, a permanently ON switch channel can be inserted in series with the feedback resistor, as shown in Figure 19.9, to reduce the gain error caused by the R_{ON} of the switches in series with the circuit resistors. In addition, since the switch resistances tend to track with temperature, gain drift due to temperature-induced resistor variation can be reduced. However this improvement in amplifier performance is achieved at the cost of one signal channel.

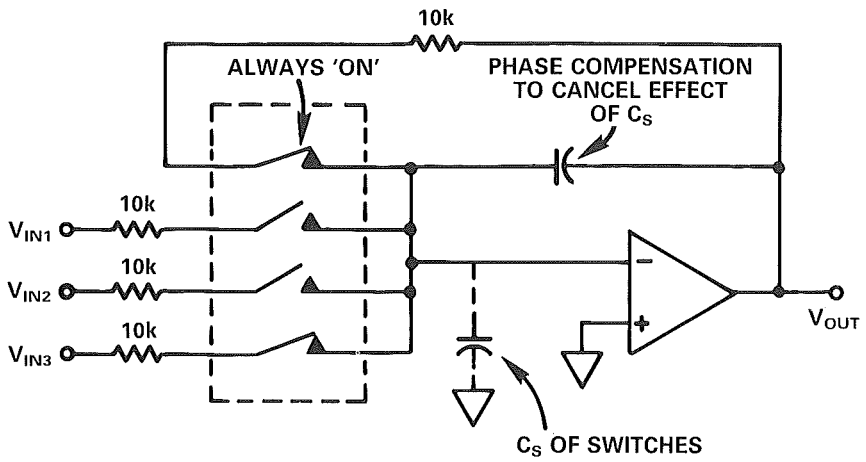

Figure 19.9. Compensating for switch resistance and capacitance.

19.2.2 SWITCH LEAKAGE

The dc OFF-isolation of an analog switch is determined by the leakage current flowing between the source and drain terminals.

However leakage current does not only flow between source and drain. Leakage currents measured at the source or drain connection flow from various parts of the circuit, as well as across the switch channel. The various contributions to switch leakage are difficult to quantify separately. An easier, and often more useful, parametric test is the total leakage at the source or drain. This model current can be used to calculate signal voltage errors due to the flow of leakage currents through channel resistance.

The problem of leakage gets worse at elevated temperatures. Leakage tends to double with every 10°C increase in temperature. ON-resistance also has a positive temperature coefficient.

The cumulative effect is that worst-case ON-resistance-leakage voltage errors occur at the maximum system operating temperature and are normally computed for that temperature.

An illustration of this is shown in Figure 19.10, where the total leakage current is shown as the sum of the ON leakage current for the ON channel, and 3 × the OFF leakage current for the three OFF channels. If $I_{D\ ON}$ is 9 nA max at +25°C and 600 nA max over temperature; and $I_{D\ OFF}$ is 3 nA max at +25°C and 200 nA max over temperature; then for $R_{FB} = 10\,k\Omega$,

$$V_{OUT} = 18 \times 10^{-9}\,A \times 10^{4}\,\Omega = 180\,\mu V, \text{at } +25°C$$

$$V_{OUT} = 1.2 \times 10^{-6}\,A \times 10^{4}\,\Omega = 12\,mV, -55°C \text{ to } +125°C$$

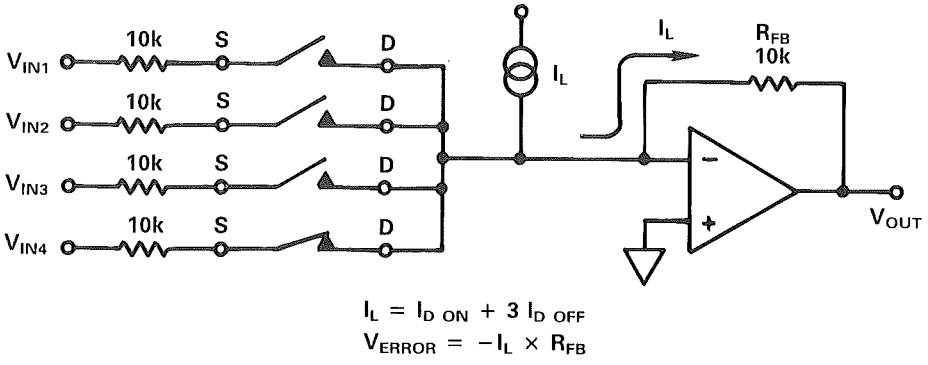

Figure 19.10. Switch leakage effects.

19.2.3 CHARGE INJECTION

As mentioned earlier, charge can be injected into the signal line due to capacitive coupling between the digital control and the source and drain pins.

The coupling is mainly associated with three separate capacitances:

Output switch n-channel gate-to-drain capacitance (C_{GDN}).
Output switch p-channel gate-to-drain capacitance (C_{GDP}).
Digital-control-pin to output-switch capacitance (C_{LIO}).

These capacitances are illustrated in Figure 19.11.

For a positive logic (logic High = Switch ON) SPST analog switch the respective voltage transitions present on these three capacitances when the switch turns OFF are (assuming ± 15-volt supplies):

+ 15 to − 15 volts. (n-channel gate capacitance)
− 15 to + 15 volts. (p-channel gate capacitance)
− 5 volts (TTL Logic) or − 15 volts (CMOS Logic)

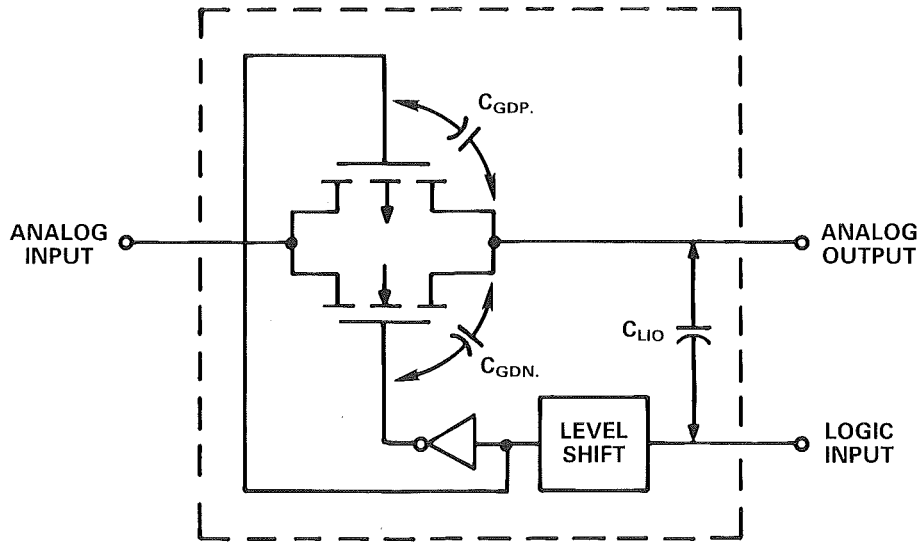

Figure 19.11. Parasitic capacitances.

The charge injections, due to each of these transitions acting on the relevant capacitance, sum algebraically to form a composite charge injection at the switch output.

The p-channel gate capacitance is usually about two to three times greater than the n-channel gate capacitance. The logic-input-to-analog-output capacitance depends largely on the device pin configuration. If these two pins are far enough apart (and the associated wiring maintains the separation), this capacitance can be neglected.

The total charge injection is given by the approximate sum:

$$Q_{INJ}\,(\text{p-channel}) + Q_{INJ}\,(\text{n-channel}) + Q_{INJ}\,(\text{logic input, Off transition})$$

$$Q_{INJ}\,(\text{total}) = C_{GDP}\,\Delta V_{GP} + C_{GDN}\,\Delta V_{GN} + C_{LIO}\,\Delta V_{L}$$

$$= +C_{GDP}\,30\text{V} - C_{GDN}\,30\text{V} - C_{LIO}\,5\text{V}$$

for $\pm$ 15 volt supplies and TTL-logic, with $V_S = 0$ volts, $R_S = 0$.

This is an oversimplified quantitative approach to computing the charge injection phenomenon. Other parameters which influence the value of charge injection are source voltage, source impedance, and effective load impedance. A typical charge-injection-vs.-source-voltage characteristic is shown in Figure 19.12.

A charge-injection measurement circuit is illustrated in Figure 19.13a, and typical test waveforms that one might see, using this circuit, are shown in Figure 19.13b.

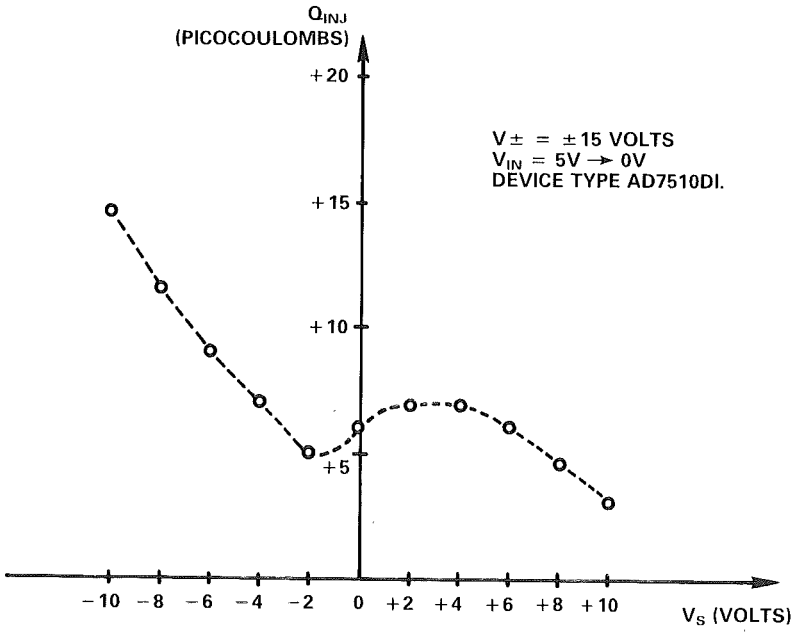

Figure 19.12. Measured drain charge injection as a function of source voltage.

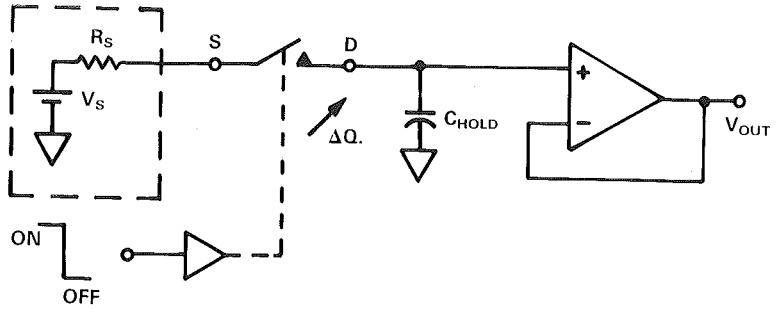

a. Charge-injection measurement circuit.

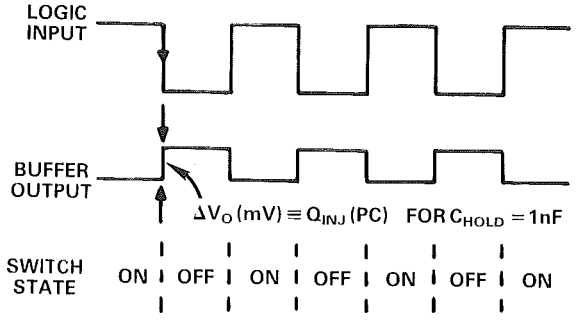

b. Waveforms.

Figure 19.13. Charge-injection measurement circuit and waveforms.

19.2.4 PUMPBACK

Figure 19.14 shows a typical analog switch configuration which can exhibit the phenomenon known as pumpback.

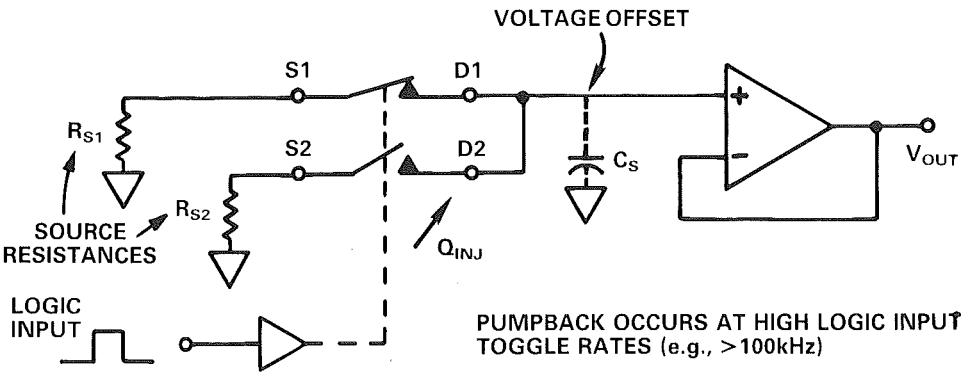

Figure 19.14. Circuit to demonstrate pumpback.

Charge, which is injected into the signal path via D1 when switch 1 turns OFF, normally discharges through the ON-resistance of switch 2, in series with the source impedance at channel 2 (R_{S2}). If the source impedance is low, the injected charge will dissipate quickly. However, if the source impedance is high, discharge will take longer and—if the frequency of changeover between channel 1 and channel 2 is high—the charge may not be sufficiently discharged by the time the next injection of charge, due to the OFF transition of switch 1, takes place. The result of this situation is a standing offset at the switch output. This phenomenon is most readily measured with both switch input voltages at zero. As the switch toggle frequency is increased, this "pumpback" offset voltage will increase.

Pumpback can thus impose an upper limit on the sampling rate of high-speed data-acquisition systems. Reduction of pumpback to increase maximum sampling rates can be effected by selecting low-charge-injection switches and by reducing source impedances.

19.2.5 SWITCHING SPEED

Analog switches and multiplexers are commonly used to sample several input signals for subsequent analog-to-digital conversion. The maximum sampling rate of the switches is determined by the propagation delay of the switch drivers and the time taken for the switch output to settle to within the required error band of the a/d converter.

Propagation delay is quantified by the turn-on and turn-off times of the switch. Turn-on time is the time elapsed from an arbitrary point on the logic

input transition (usually 50%) to an arbitrary point (usually 90% of the final value) on the analog output transition. Switching speed and settling time measurements are illustrated in Figure 19.15.

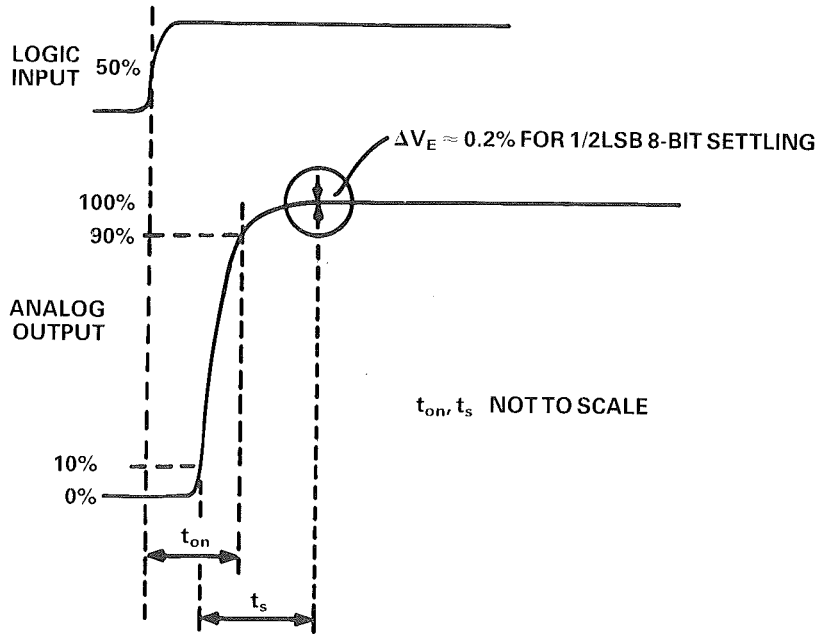

Figure 19.15. Switching speed and settling time.

19.2.6 SETTLING TIME

Settling time, from an arbitrary point on the analog output transition (e.g., 10%) to within a specified percentage of the final value, is a function of the signal source impedance, the switch ON-resistance and the capacitance at the switch output. Measurements of settling time require great care because of the high speed, the role of stray capacitance, and the transient response of the measurement circuitry.

19.2.7 OFF ISOLATION

Analog switches and multiplexers have a limited ability to isolate ac and dc input signals from the switch output. OFF Isolation is limited at dc by the leakage current across the switch. At higher frequencies, the switch capacitances reduce the OFF impedance of the switch. An OFF-Isolation specification must be accompanied by a statement of the frequency of measurement and the load impedance, as both affect the switch isolation. OFF isolation is defined as the ratio, V_{OUT}/V_{IN}, with the switch OFF, and is expressed in dB as $20 \log_{10}(V_{OUT}/V_{IN})$.

19.2.8 INSERTION LOSS

This measure of the ON-channel attenuation in a signal path is defined, at a specific frequency and a given load, as the ratio, V_{OUT}/V_{IN}, with the switch ON. It is expressed in dB as $20 \log_{10} (V_{OUT}/V_{IN})$.

At dc, the insertion loss is approximately $R_L /(R_L + R_{ON})$, where R_L is the load resistance and R_{ON} is the channel ON resistance.

19.2.9 CROSSTALK

This is the amount of spurious signal crossover from the signal input of an OFF-channel to the input of a nearby ON channel.

Crosstalk in analog switches and multiplexers is mainly due to capacitive coupling between channels and is defined as the ratio $V_{IN\ ON} /V_{IN\ OFF}$, where $V_{IN\ OFF}$ is the signal input to the OFF channel and $V_{IN\ ON}$ is the resulting signal at the ON-channel input.

It is important to note that the inherent parameters of the switching device set a *minimum* to the system's expected high-frequency isolation, crosstalk, and insertion loss at high frequencies. These parameters can be further greatly influenced by circuit-board layout.

19.3 MULTIPLEXING AND MULTIPLE SWITCHES

To achieve analog switching of multiple inputs to a common output (or vice versa), using a single IC device, the switching function is designed to make optimum use of each available package pin. The analog multiplexer IC invariably uses binary address-decoding to minimize the number of channel-select control pins.

Features included in the standard multiplexer are an Enable pin, which allows the user to turn off all channels, regardless of the chosen (decoded) channel, and break-before-make switch action, which helps protect the input sources by ensuring that no two channel inputs are connected to one another.

Figure 19.16 shows a typical multiplexed-input data-acquisition system. The 12-bit a/d converter must wait until the output of the multiplexer has settled to within ½ LSB of the decoded input's voltage range before starting a conversion.

The settling time of a multiplexer is directly dependent on the time constant established by the load capacitance and the sum of the input source impedance and the multiplexer ON-resistance.

The time from the application of the multiplexer decode pulse to the point where the a/d converter can start converting is the sum of the switching time (t_{ON}) of the multiplexer and the further time it takes for the buffered multi-

plexer output voltage to settle to within ½ LSB (12 bits) of the analog input range.

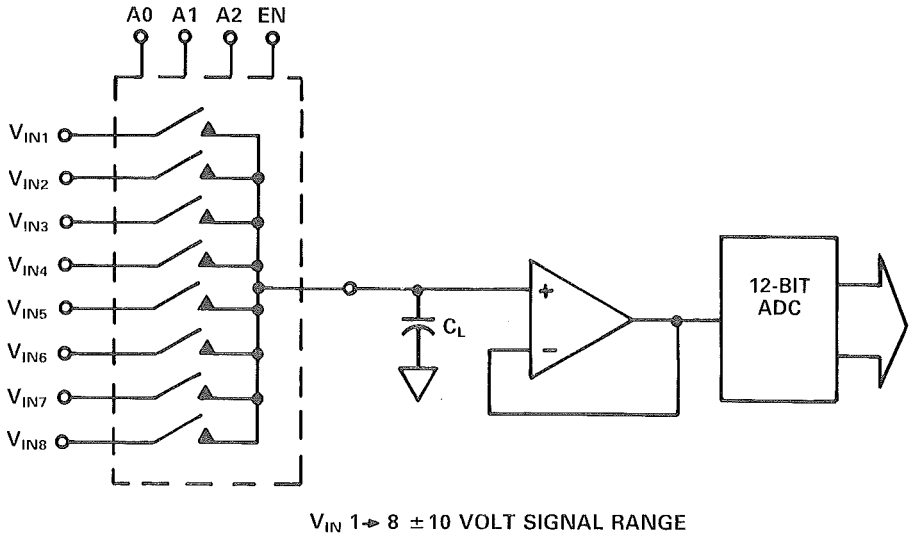

Figure 19.16. Multiplexer application.

19.3.1 LOW-LEVEL SIGNAL MULTIPLEXING

In high-resolution data-acquisition systems, electrically noisy environments can preclude the use of conventional single-ended multiplexing techniques. If a high-CMRR amplifier-per-channel, to eliminate common-mode voltage and amplify the signals substantially before multiplexing, is not feasible, differential multiplexing offers improved common-mode rejection and is particularly useful for switching low-level signals before amplification.

The advantages of differential multiplexing are threefold.

- The data-acquisition system remains differential, from the transducers through to the high CMRR instrumentation amplifier(s).
- Noise induced due to charge injection of the multiplexer switches, being a common-mode effect, tends to cancel if differential multiplexing is used.
- Low-cost twisted-pair wiring can sometimes be used as an alternative to expensive shielded conductors.

Generally, the lower the signal level, the more important the need for differential multiplexing. Figure 19.17 shows a typical differentially multiplexed front end of a low-signal-level data-acquisition system.

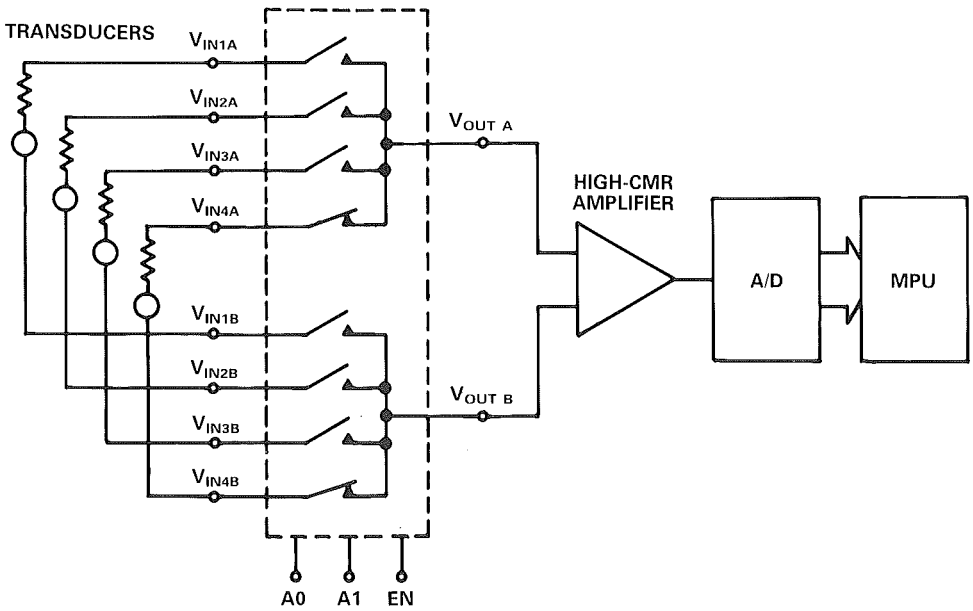

Figure 19.17. Four-channel differentially multiplexed data acquisition.

19.3.2 MULTI-RANK MULTIPLEXING

Figure 19.18 shows a two-level single-ended multiplexing system, used for channel expansion. In this example, eight 8-channel multiplexers are multiplexed into a single 8-channel multiplexer to handle 64 channels of single-ended input. Although it uses an additional multiplexer (the system could have consisted of eight 8-channel multiplexers with their outputs paralleled), it has significant advantages over the single-level n^2-channel system. For example, the n^2-channel system would have n multiplexer outputs connected in parallel, while the 2-rank system only looks at one output at a time, greatly reducing output capacitance and leakage effects.

The error of this circuit due to leakage current is

$$2 \times I_{OUT\ ON} \times (2\,R_{ON} + R_S)$$

where R_S is the analog input source resistance.

The error due to leakage of an equivalent single-level multiplexer system is:

$$(7 \times I_{OUT\ OFF} + I_{OUT\ ON}) \times (R_{ON} + R_S)$$

where $I_{OUT\ ON}$ and $I_{OUT\ OFF}$ are the output leakages of an 8-channel multiplexer (e.g., AD7501) in the Enabled and Disabled modes, respectively.

For source resistances which are significantly greater than R_{ON}, the leakage error of the single-level system approaches 4× that of the two-level multiplexer circuit. This improvement factor is reduced to 2× for source impedances that are low relative to R_{ON}. ($I_{OUT\ OFF} \simeq I_{OUT\ ON}$)

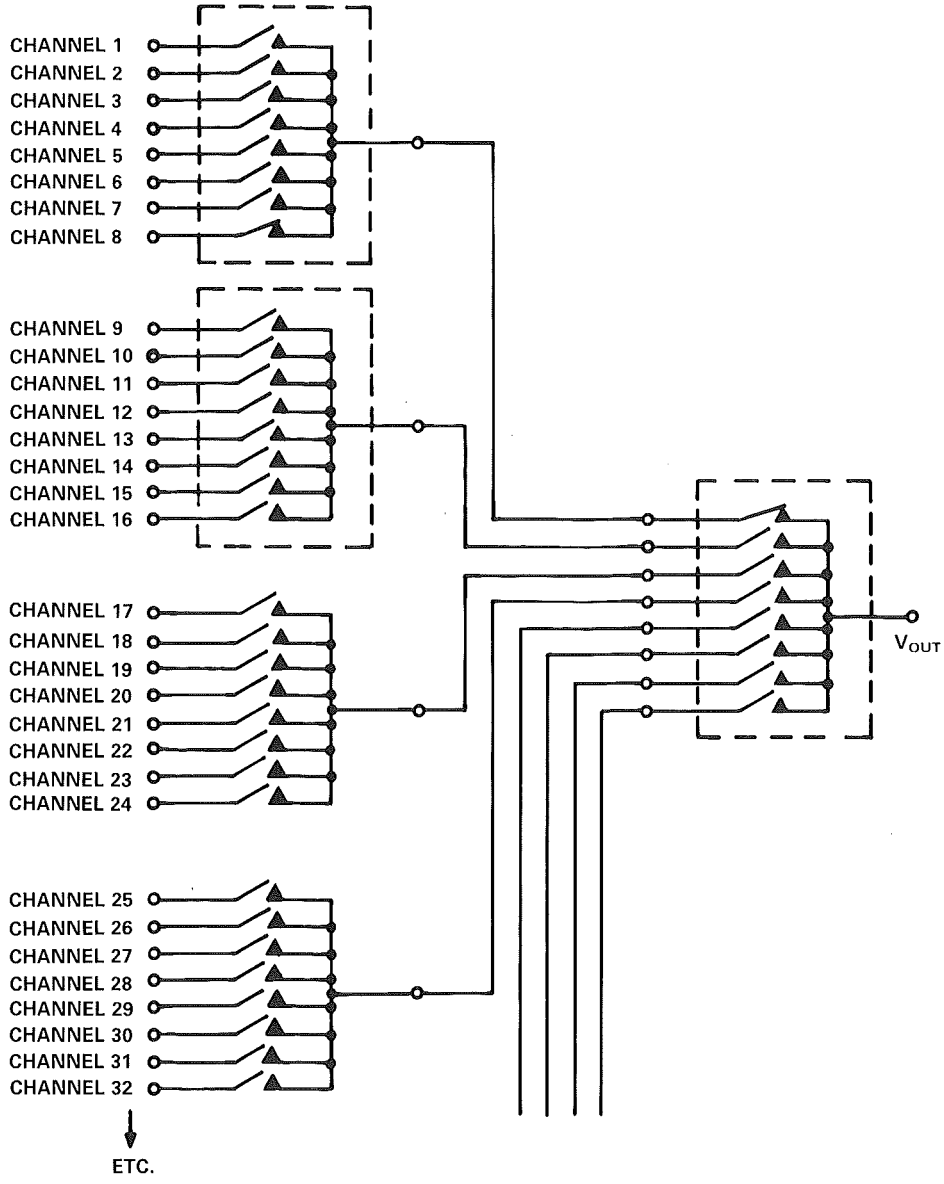

Figure 19.18. Two-level multiplexing. Channel 8 is selected.

A further advantage of two-level multiplexing is potentially faster system switching speed. The maximum toggling rate of the single-level system is limited by the maximum switching speed of each multiplexer. For example, if the multiplexer's minimum t_{ON} and t_{OFF} are 1 microsecond, the maximum toggling rate is 500 kHz. The two-level circuit, however, can be configured using standard 1-μs t_{ON} multiplexers in the first level and fast (e.g., less-than 200-ns t_{ON}) switches in the second level.

Suppose channel 1 of all the first-level multiplexers has been selected; the second-level switches can strobe each multiplexer in sequence. The toggling rate is thus determined by the switching speed of the second-level switches. Approximately half-way through the sequential selection of the eight multiplexers, the first four multiplexers are switched to Channel 2.

By the time that the strobing of channel 1 of the last four multiplexers has been completed, channel 2 of the first four multiplexers is already selected. The maximum system toggling rate has become four times that of the single-level system. This maximum toggling rate is achieved if the second level switches are at least four times as fast as those of the first-level multiplexers.

19.3.3 MULTIPLE-SWITCH CONFIGURATIONS

Combinations of switches can offer improvements in performance over that available from a single switch.

Switches in Parallel

By paralleling two or more switch channels, as shown in Figure 19.19, the ON resistance of the switch is reduced, and the current handling capability is increased.

$$R_{ONT} = \frac{R_{ON}}{N}$$

$$I_{DT} = N \times I_D$$

Where R_{ONT} = ON resistance of switch combination.
R_{ON} = ON resistance of separate switches.
N = No of switches in parallel.
I_{DT} = Maximum drain current of switch combination.
I_D = Maximum drain current of separate switches.

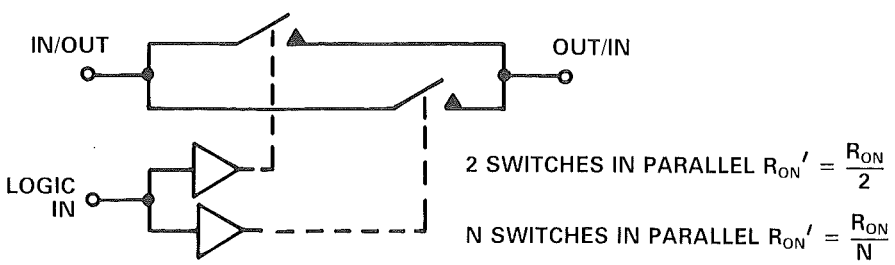

Figure 19.19. Paralleling analog switches.

T-Switch

Figure 19.20 shows three switches configured in T-formation. This configuration provides an extra stage of attenuation, offering superior OFF-isolation performance at both low and high frequencies.

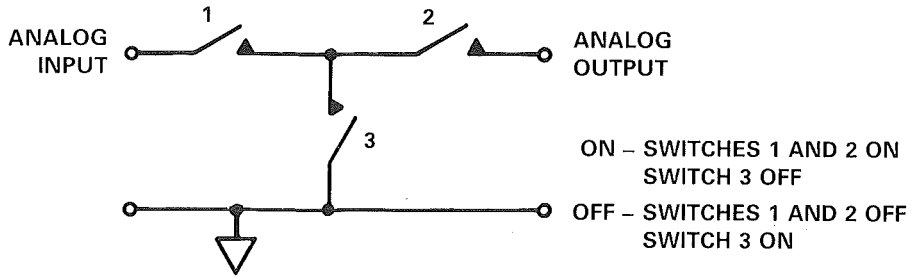

Figure 19.20. T-switch configuration.

19.4 APPLICATIONS

The applications of analog switches and multiplexers can be categorized under the following two broad headings:

Signal Routing
Signal Processing

Signal Routing

Examples of signal routing applications include multiplexed-input data-acquisition systems, automatic test equipment, force and sense selection, and message switching.

Signal Processing

Examples of signal processing applications include programmable-gain amplifiers and programmable filters.

19.4.1 SAMPLE-AND-HOLD

Sample-hold circuits form a significant subset of the applications of analog switches. It is worthwhile to discuss design aspects of such circuits.

Figure 19.21 shows a basic sample-hold circuit suitable for interfacing to 8-bit a/d converters. The time taken for the capacitor to charge to within ½ LSB (< 0.2%) of its final value is ln(512) = 6.24 time constants, or $6.24 \times (R_S + R_{ON})\, C_H$. For zero source impedance, C_H = 10 nF, and R_{ON} = 100Ω, 8-bit ½-LSB switch settling time will be about 6.2 μs.

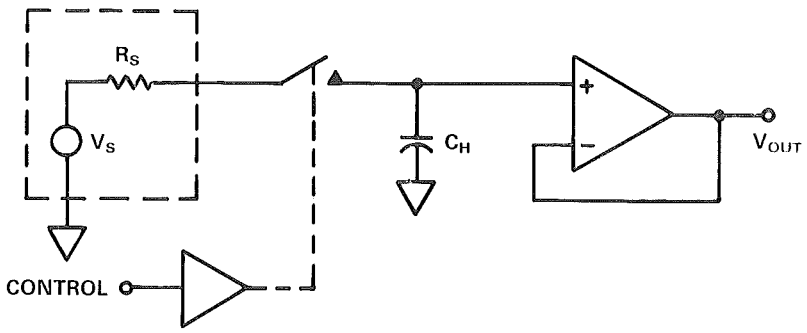

Figure 19.21. Sample-hold circuit.

The *acquisition time* (the time needed to acquire a sample of data), from the time the logic input transition is applied to the switch, is the switch-circuit settling time plus the switch propagation delay. A typical analog switch may have a propagation delay of 300ns; therefore about 6.5 microseconds total sampling time is needed.

Voltage *droop* on the hold capacitor, during HOLD, is caused by buffer amplifier bias-current flow, switch leakage and capacitor leakage. The voltage on the hold capacitor must stay within ½ LSB, at most, of the signal originally applied to the input of the sample/hold, until the a/d converter has completed its conversion.

A design trade off must be made between the following factors:

Magnitude of hold capacitance—C_H.
Settling Time of Sample/Hold—T_S.
Droop Rate of hold-capacitor voltage—$\Delta V_H/\Delta T$.
Charge Injection Error due to switch—ΔV_Q.

Reducing C_H will generally reduce T_S, but lower values of C_H will increase $\Delta V_H/\Delta T$ and ΔV_Q (charge injection is related to the switch capacitance).

The salient parameters affecting the choice of an analog switch for a sample/hold are:

Switch Leakage—I_{DOFF}
Charge Injection—Q_{INJ}
R_{ON}
Switch Propagation Delay—t_{ON} and t_{OFF}

Low charge injection can be achieved by selecting high ON-resistance switches, which exhibit low switch capacitance. Furthermore, switches having higher ON resistance tend to exhibit lower leakage. However, higher ON-resistance increases settling time. The first three listed parameters can be traded off to achieve a suitable system solution.

Propagation delays are generally independent of other analog constraints; the choice is for t_{ON} and t_{OFF} as low as required, based on what is available at

reasonable cost. Generally, the shaving of a few hundred nanoseconds off propagation delays will ease the tradeoffs between switch leakage, charge injection and R_{ON}.

19.4.2 DAC DEGLITCHING

Figure 19.22 shows a single analog switch channel being used in a sample-hold circuit to deglitch the output of an 8-bit d/a converter. During the DAC code transition, when the glitch is generated, the switch remains in the HOLD mode, ignoring changes in the DAC output. After the transition is over, the SAMPLE mode is initiated, recharging the hold capacitor to the new value of DAC output voltage.

The resulting performance is determined by the charge injection, glitch energy, and feedthrough of the analog switch.

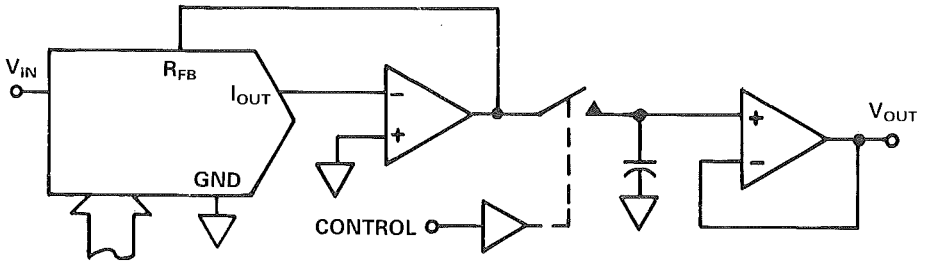

Figure 19.22. DAC deglitching.

19.4.3 DAC OUTPUT DEMULTIPLEXING

An extension of the basic deglitching function is the distributive multiplexing (demultiplexing) of the DAC output to provide a number of independent glitch-free voltage outputs. Figure 19.23 shows a quad analog switch in an arrangement that provides four separate digitally programmable voltages from the same DAC.

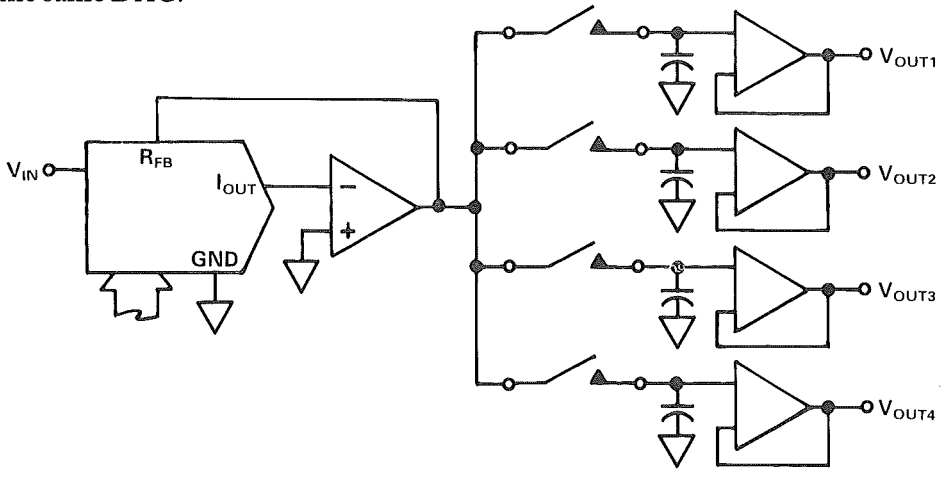

Figure 19.23. DAC output demultiplexing (data distribution).

19.4.4 OTHER SIGNIFICANT APPLICATIONS (Figures 19.24 and 19.25)

Programmable-Gain Amplifier

In Figure 19.24, a quad switch chooses the appropriate output of a tapped feedback voltage divider to set the gain of an operational amplifier. Similar circuitry can be used with some instrumentation amplifiers (e.g., the Analog Devices AD625) to set gains digitally, in such applications as automatic ranging, where dynamic range must be extended.

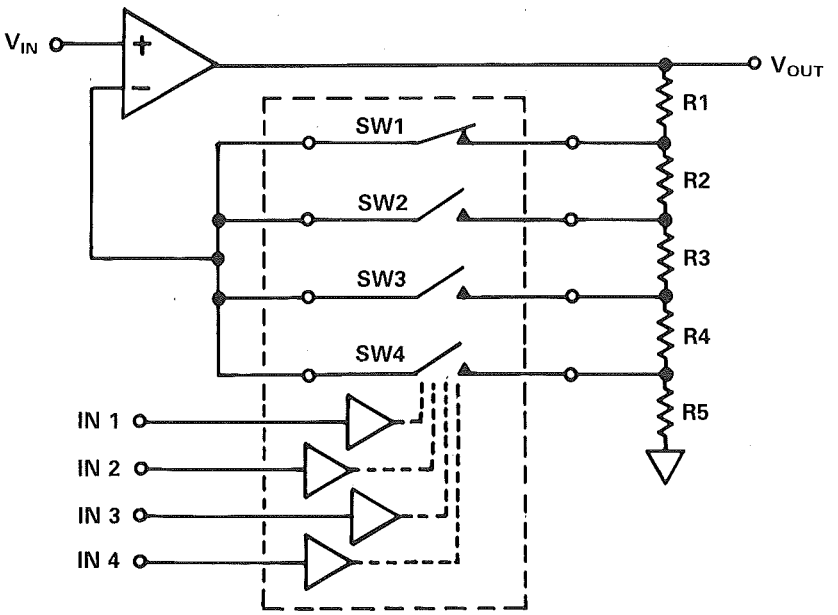

Figure 19.24. Digitally programmable-gain amplifier.

A/D Converter Input Autoranging

In Figure 19.25, a quad switch chooses a digitally determined attenuation level to establish the input range to an a/d converter having a buffered input. The result of the conversion can be used to readjust the attenuation for a smaller or larger input signal, thus achieving a limited degree of floating-point operation.

The above are but a small sampling of the uses for CMOS analog switches, limited to applications related to analog-digital conversion, within the charter of this book. However, there are a great many more potential applications in measurement and control and general electronics.

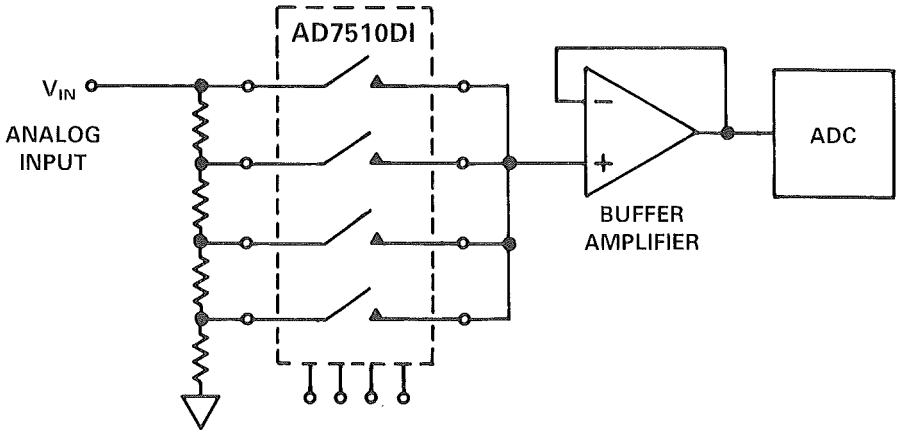

Figure 19.25. Digital scale-factor ranging for converter input.

Just one highly suggestive example: Figure 19.26 shows a circuit exemplifying the use of digitally controlled analog switches in remote control. Here they make possible a remotely programmable cutoff-frequency low-pass filter. The remote source of control can be either automatic digital equipment or a manual switch.

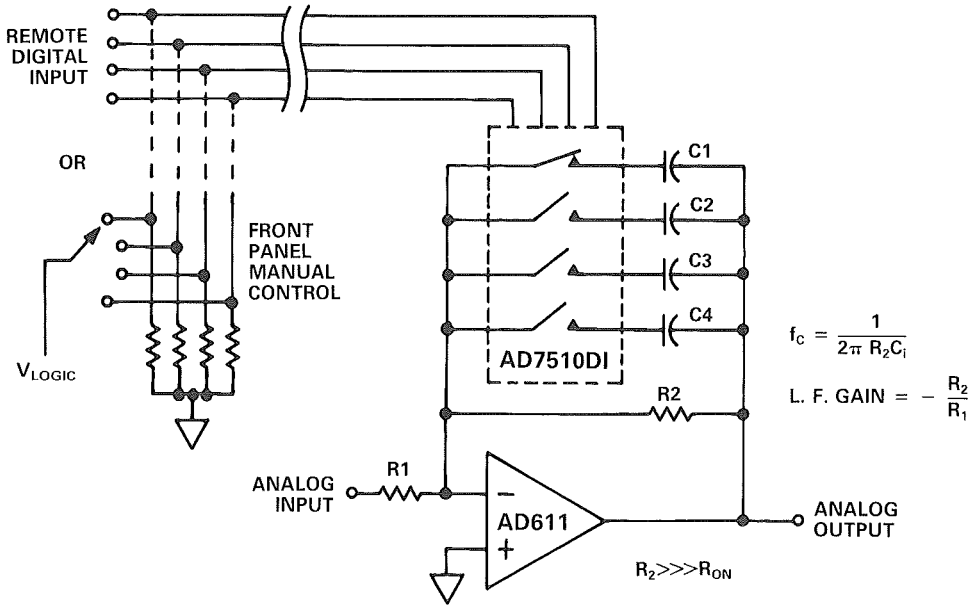

Figure 19.26 Remotely controlled analog filter circuit.

In such circuits, the only long leads are the digital control lines. The switches can be located in close proximity to the circuits they control. optoisolator circuitry can be employed if necessary to deal with large common-mode voltages in the digital circuitry.

Typical applications of the technique range from autocalibration routines in automatic test equipment to "soft keys" in portable laboratory instrumentation.

Chapter Twenty

Voltage References

Voltage references provide accurately known voltage for use in circuits or systems. Measurement systems rely on precision references in order to establish a basis for absolute measurement accuracy. Any reference error undermines the overall system accuracy, thus ideal references are characterized by accurately set constant output voltage, independent of load changes, temperature, input (supply voltage) and time.

20.1 WHERE REFERENCES ARE USED

A digital multimeter, digital communication system, portable instrument for precision measurement or calibration, electronic thermometer, precision switching regulator, or—in fact—any digital system, subsystem, or device with analog inputs or outputs needs at least one accurate reference. Such a reference might be a dc voltage source, with an accurate output somewhere in the range from 1.0 V to 10 V, stable to within a few parts per million per degree C and per month.

For example, in a/d converters, the digital output number depends on the ratio of the quantized input to the "full scale" reference. If the reference is allowed to change in response to a second analog input, the digital output will be proportional to the ratio of the analog input signal to the reference signal (i.e., a ratiometric ADC is equivalent to an analog divider with digital output). Often, the change in response to any "second analog input" may be unwanted, especially if that input is temperature, time, or upstream voltage; in such cases, an insensitive precision reference is essential.

A d/a converter is essentially a digitally controlled potentiometer that produces an analog output (voltage or current) that is a normalized fraction of

its "full scale" setting. The full-scale setting is determined by the reference value chosen. Again, if the reference changes, the output will change; if the DAC output must be fixed at any digitally set value, a precision reference with appropriate characteristics must be used.

Desirable voltage reference characteristics include:

- *Accurate output voltage*
- *Low temperature drift*
- *Good load regulation*
- *Good line regulation*
- *Good long-term stability*

20.2 TYPES OF REFERENCES

References come in a variety of shapes and forms, depending on their projected use, ranging from banks of saturated standard cells to Josephson junctions to integrated-circuit voltage regulators. There are a number of ways of classifying references; for example, there are primary standards, secondary standards, transfer standards, etc. For our purpose, we may divide the family of references into two principal categories: *depletable* types—which can be used without a primary power source but cannot deliver power over long periods of time without degradation (i.e., batteries maintaining an accurate voltage with low drift)—and essentially *undepletable* types, which are based on fundamental relationships that give constant or predictable voltage but require a primary power source.

In this chapter, we will consider only the kinds of references commonly employed as components in converter circuits, using available power, with the understanding that their calibration may be traceable to international voltage standards.

20.2.1 ZENER-DIODE REFERENCES

The most widely used reference device is the temperature-compensated Zener* diode. Zeners have constant voltage drop in a circuit when provided with a fairly constant current derived from a higher voltage elsewhere within the circuit. They are often used with operational amplifiers in circuits that stabilize operating points, unload the diode's output impedance, and/or transduce voltage to current.

How The Zener Diode Works

The active portion of a Zener diode is a reverse-biased semiconductor P-N junction. When the diode is forward biased (the P region is made more positive), there is very little resistance to current flow. Actually, in the forward-

*Whatever the reverse breakdown mechanism—avalanche, Zener, or mixed—all diodes used in the breakdown mode have come to be called "Zener diodes," even though many of the diodes purchased for specified breakdown characteristics are in fact purely avalanche diodes. We will follow the custom in this chapter.

biased region, a Zener diode looks very much like a normal high-conductance silicon diode (Figure 20.1). Values of V_F greater than about 0.7V will produce substantial amounts of current.

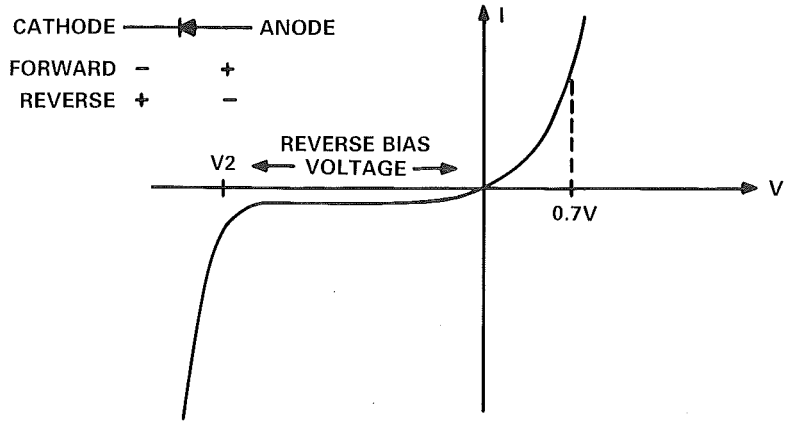

Figure 20.1. Generalized I-V characteristic for Zener diode.

In the reverse-bias region, very little current flows if V is less than V_Z, the *breakdown voltage*. The small amount of leakage current that flows—the reverse saturation current—is relatively insensitive to the actual magnitude of reverse voltage for fixed temperatures.

As the reverse voltage across the diode approaches the breakdown voltage, the reverse current increases more rapidly and will even run away if the applied voltage is sufficient. For this reason, diodes operated in the breakdown region are always used in series with resistsors or current sources. The sharpness of the transition depends on the relative value of the breakdown voltage and the manufacturing process used to make the diode. The most common circuit in which Zener diodes are used employs a series resistor to limit the current, as shown in Figure 20.2.

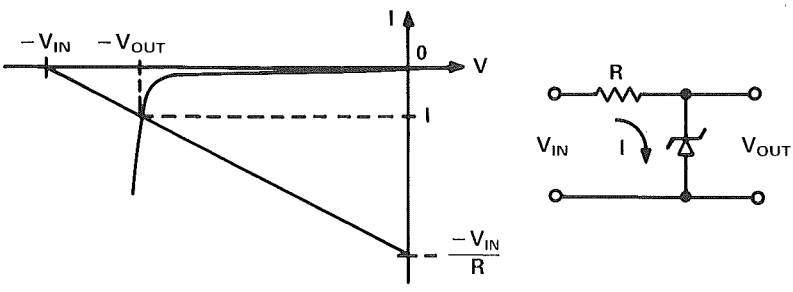

Figure 20.2. Breakdown diode characteristics and application.

It is worth knowing that there are two significantly different reverse-bias voltage breakdown phenomena, Zener breakdown and avalanche breakdown; the difference can be important to both the user and the designer of voltage-refer-

ence ICs. In Zener breakdown (Figure 20.3), a low-voltage phenomenon, the breakdown voltage *decreases* as the junction temperature rises. In avalanche breakdown the breakdown voltage *increases* as the junction temperature rises.

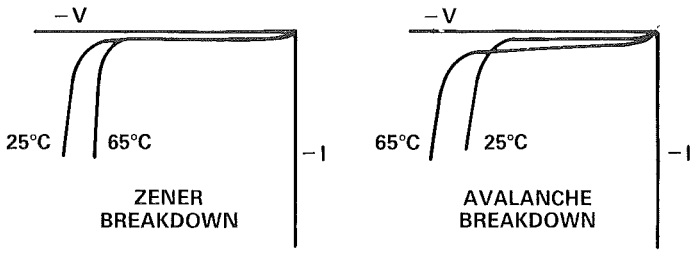

Figure 20.3. Avalanche breakdown vs. Zener breakdown.

The factor determining which manner of breakdown occurs is the relative concentration of impurities in the P- and N-type materials forming the P-N junction. A junction that has a narrow depletion region at a given voltage will develop a high field intensity and will break down by the Zener mechanism at relatively low voltages, as electrons and holes are stripped away from atoms, thus providing the carriers for conduction.

A high electric field supplies the energy required to traverse the "energy gap", i.e., electrons are excited from the valence band to the conduction band. Since conduction is a function of the energy gap (the band gap, or forbidden energy gap), an increase in temperature reduces the band gap and thus the breakdown, or Zener, voltage by increasing the energy of the valence electrons.

In high-resistivity materials, the depletion region is of sufficient width to avoid Zener breakdown; carriers will experience collisions before crossing the region completely at a given temperature. But, as the temperature increases, the kinetic energy of the valence electrons increases to the point that, because of the high energy level, a collision may rip off more carriers from an atom. The newly released carriers now gain sufficient energy from the field to begin collisions of their own, producing the "avalanche" effect.

Because avalanche breakdown is a kinetic process, it is more temperature-sensitive than Zener breakdown. Changes in temperature will cause a change in the mean kinetic energy of a particle and thus affect the avalanche breakdown mechanism. Zener breakdown, on the other hand, is dependent on the strength of the electric field and is theoretically insensitive to changes of temperature.

Zener-Diode Performance

"Zeners" are available with voltages from about 2 to 200V, tolerances of 10% to 20%, and power dissipation from a fraction of a watt to 40 or 50 watts.

Attractive as they might seem for use by themselves as general-purpose voltage references, they have many shortcomings. For use without additional circuitry, it would be necessary to stock a range of values; the voltage tolerance is generally poor, except in high-priced versions; and they are noisy and very sensitive to changes in current and temperature. For example, a 1N5221 (27-volt Zener) has a temperature coefficient of +0.1%/°C and will change by 1% for a variation in current of from 10% to 50% of its maximum rating.

Zener breakdown has a negative temperature coefficient, while avalanche breakdown has a positive temperature coefficient. Both are relatively independent of current, if self-heating effects are ignored. Zener diodes which are in the 6-V range exhibit both avalanche and Zener breakdown and have either positive or negative temperature coefficient, depending on which effect predominates.

For this reason, performance is best for diodes that break down in the neighborhood of 6 volts; they achieve very low temperature coefficients and become relatively stiff against changes in current (Figure 20.4) because the positive and negative temperature coefficients tend to cancel one another.

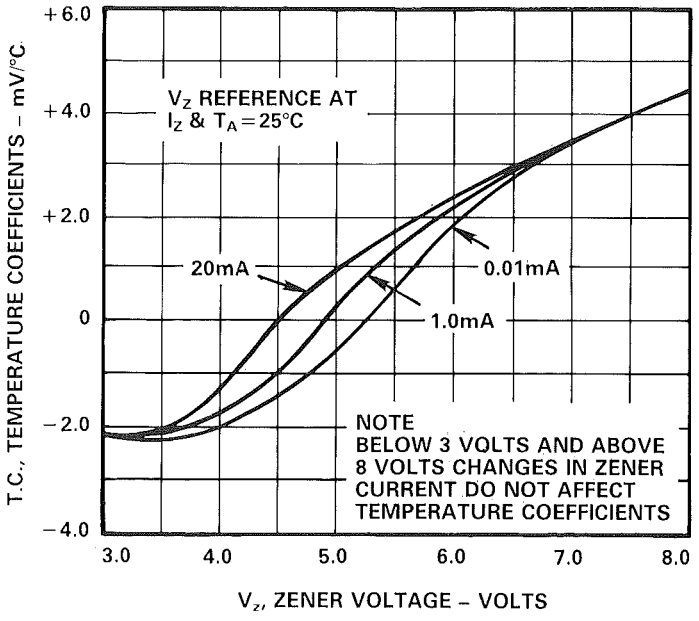

(Courtesy of Motorola, Inc.)

Figure 20.4. Temperature coefficient as a function of Zener voltage and current for a typical device.

Temperature Compensation

For a given Zener voltage, at low current levels, the Zener effect is stronger and the temperature coefficient is negative. At higher current levels, the

avalanche effect takes over and the temperature coefficient becomes positive. Because both Zener and avalanche effect are occurring simultaneously and are controlled by the current level, the temperature coefficient varies as the current level changes. At some specific level of current the negative TC of the Zener effect is equal to the positive TC of the avalanche effect, and the net temperature coefficient is theoretically zero. Therefore, as can be seen in Figure 20.4, by proper choice of the diode's reverse current, the temperature coefficient can be adjusted for breakdown at a given voltage. Alternatively, a compensated Zener reference can be built using a Zener diode with a positive temperature coefficient in series with a forward-biased diode. The Zener diode's voltage is chosen to cancel the forward diode's temperature coefficient.

If you need a Zener for use in an application where stability is the prime concern, and you don't care what the exact voltage value is, the best choice would be a compensated Zener reference, made from a 5.6-volt Zener in series with a forward-biased diode.

The temperature coefficient of the forward-biased diode (Figure 20.5) is often more sensitive to current than that of the reverse-biased Zener diode; this can be detrimental to good performance unless the bias current is kept very close to the level specified by the manufacturer.

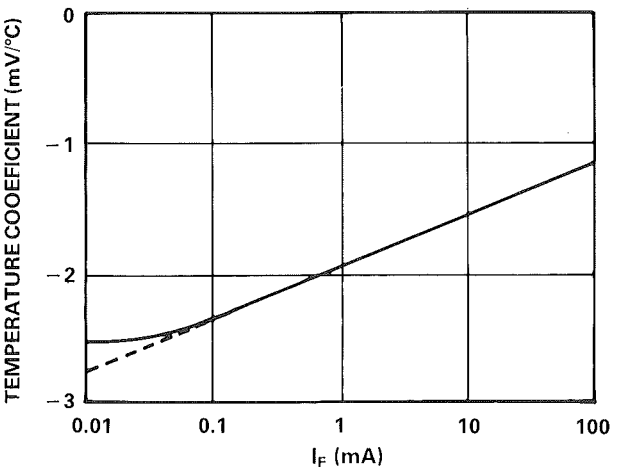

Figure 20.5. Temperature coefficient of a typical forward-biased junction diode as a function of current.

IC Zeners

Successful compensation to achieve near-zero temperature coefficients requires close tracking of the junction temperatures of the Zener and the forward-biased diodes. It is difficult to do this with diodes in different packages—the best arrangement is to place the diodes in direct contact with each other (Figure 20.6).

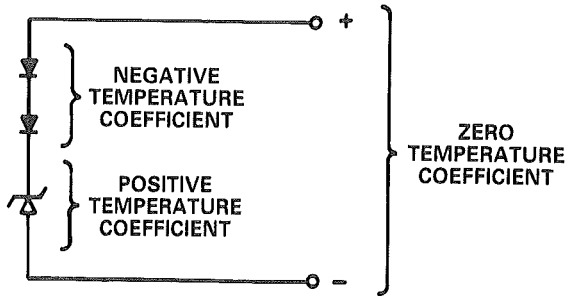

Figure 20.6. Using forward-biased diodes in temperature compensation.

Reference Circuitry Using Zener Diodes

These compensated Zeners can be used as stable voltage references within a circuit, but they must be supplied with constant current. A choice of technologies exists for performing this function; the appropriate choice, as always, depends on the circuit application.

High-performance hybrid references usually employ a temperature-compensated Zener diode and a feedback amplifier. The amplifier serves several purposes, providing constant bias current to the Zener, buffering it, allowing the reference circuit to source and sink current, and providing gain to boost the nominal 6.2-V diode voltage to the desired level (Figure 20.7). A low offset-drift amplifier, such as the AD510, or the AD OP-07, is essential, as the amplifier's offset drift will add to the drift of the Zener diode.

Here's how it works: Assuming that the circuit has started up and is operating properly (R5 provides the diode's startup current), current will flow through R3 and the diode, thus setting the plus input of the op amp at $+V_z$. However, the minus input must follow, and it can do this only if

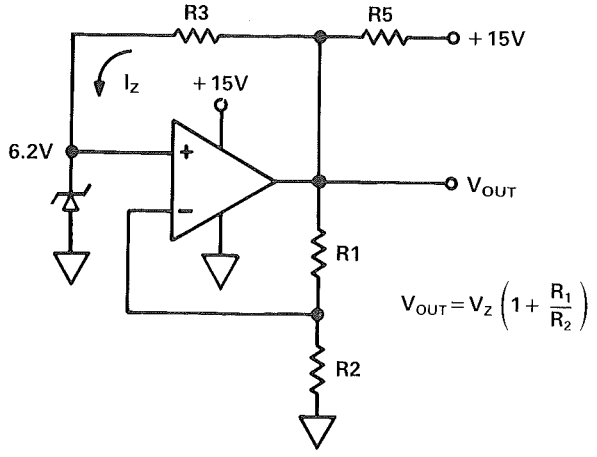

Figure 20.7. Hybrid 10-volt reference, using temperature-compensated Zener.

$$V_{OUT} = \left(1 + \frac{R_1}{R_2}\right) V_z \qquad 20.1$$

The current through R3 must then be equal to $(V_{OUT} - V_z)/R_3$, and is maintained at that value, independently of the supply voltage, amplifier loading, and—to a first order—temperature.

R3 is chosen to provide the desired value of bias current, which is determined by the stable difference between V_z and $(1 + R_1/R_2)V_z$, i.e.,

$$I = \frac{R_1}{R_2 R_3} V_z \qquad 20.2$$

The R_1/R_2 resistance ratio can be trimmed to set the output voltage to the desired degree of accuracy, and R_3 can be trimmed to set the current to the proper value to minimize the drift of the Zener and thus decrease the output temperature coefficient. Hybrid references, with their freedom to mix technology, can usually be obtained with tight initial accuracy, using lasers to trim thin-film resistor networks. An example of a commercially available precision reference using this kind of circuit is the Analog Devices AD2710, which has 1 millivolt of initial error and a temperature coefficient of one ppm/°C.

A major use for references of this kind is with d/a and a/d converter types that require an external reference. Figure 20.8 shows how the reference is used with a high-speed 12-bit DAC. The AD566KD is laser-trimmed for ± ¼-LSB maximum nonlinearity, and exhibits a gain temperature coefficient of 3 ppm/°C. When used with the AD2710LN reference, the worst-case total gain temperature coefficient is 4 ppm/°C. After initial calibration of the DAC scale fac-

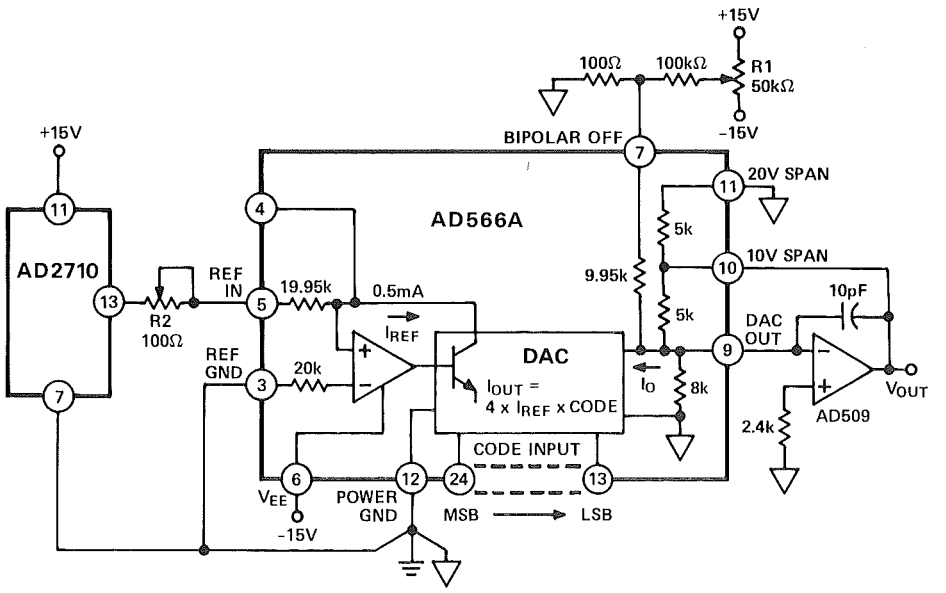

Figure 20.8. Reference for a precision d/a converter.

tor at room temperature, 12-bit accuracy can be maintained over the +15°C to +70°C temperature range. The AD2710's output-current capability makes it possible for it to serve as a reference for up to 10 such converters in a system.

Another approach to providing references for multiple DACs is to use the reference with single-package multiple DACs, such as the Analog Devices AD390. The combination of the AD2710LN and AD390KD (quad 12 bit d/a converter) will yield a compact multiple-DAC system (Figure 20.9) with a maximum full scale drift of ±6ppm/°C and excellent tracking.

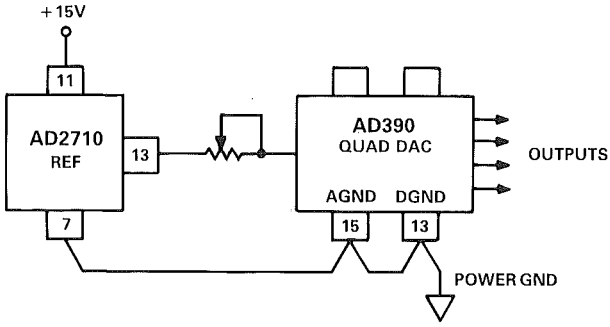

Figure 20.9. Furnishing a reference for four d/a converters at once in a quad DAC.

Buried-Zener References (BZRs)

A major category of precision *monolithic* voltage reference sources (other than band-gap devices) are subsurface, or *buried* Zener diodes. Development of the buried Zener diode was all but essential to the manufacture of accurate, high-resolution complete monolithic DACs and ADCs.

In ICs, Zener references have traditionally been produced using the reverse breakdown of the base-emitter junction of a vertical NPN transistor. This breakdown occurs right at the surface of the device, where the voltage is af-

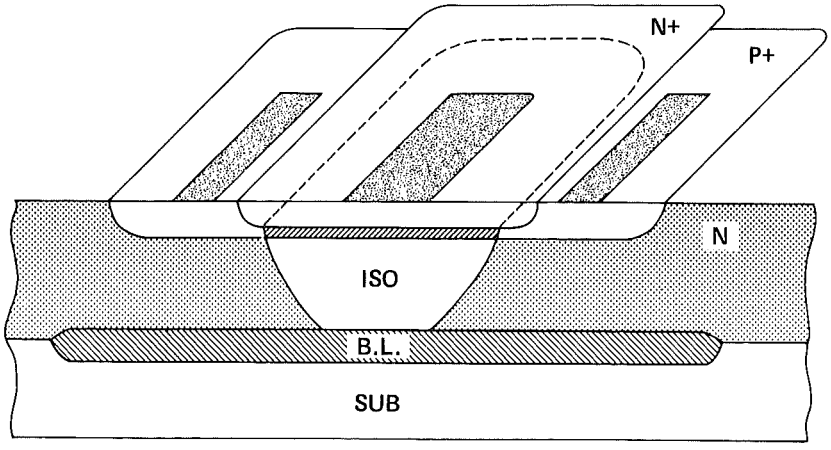

Figure 20.10. Subsurface Zener.

fected by crystal imperfections, mobile charges in the oxide, and other forms of contamination. These effects cause the surface-breakdown diode's noise and long-term stability to be unsuitable for 10- or 12-bit applications. The modified diffusions of the buried Zener (Figure 20.10) cause breakdown to occur well below the surface, thus avoiding surface effects. Long-term stability of 50 ppm/year can be achieved with this IC-process-compatible device.

However, because the diffusion is less controlled below the surface than on the surface, there is greater spread in the absolute values and temperature coefficients of subsurface Zener diodes, often exceeding acceptable tolerances. The circuitry used in conjunction with the diode must be designed to allow trimming of both the absolute value and the temperature coefficient of the whole reference circuit; A typical approach is shown in Figure 20.11. Typically, absolute accuracy is to within 0.1% and temperature coefficient is ± 10 ppm/°C. Buried-Zener reference circuitry is used to great advantage in complete monolithic high-resolution converters, such as the Analog Devices 12-bit AD667 DAC, since the overall *output* voltage and its tempco can be automatically laser-trimmed at the wafer stage—at the same time that other DAC parameters are trimmed.

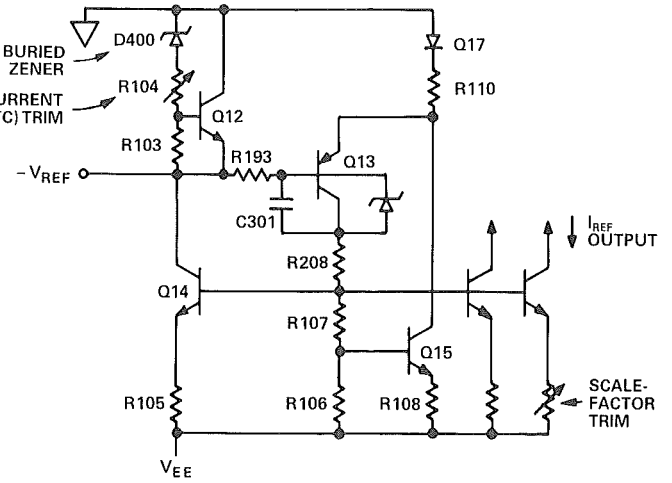

Figure 20.11. Zener-circuitry portion of IC chip.

20.2.2 TEMPERATURE-STABILIZED REFERENCES

Another way to improve temperature stability in IC references is to hold the reference at a constant elevated temperature. Temperature-stabilized (not "compensated") devices include on a single chip a temperature stabilizer (heater) and the reference circuit. The heater maintains the reference at a constant temperature that is independent of ambient temperatures up to the preset stabilizer temperature, but no cooling is available at higher temperatures. For this reason, it is common to see substantially differing temperature coeffi-

cients for the two operating ambient temperature ranges (ambients above and below heater temperature) for a device of this type.

For example, an LM199 temperature-stabilized reference is specified as having a temperature coefficient of ± ½ ppm/°C from −55°C to +85°C, and 10 ppm/°C from +85°C to +125°C. This type of performance is most suitable for laboratory equipment in controlled environmental conditions where ample power is available. Since the heater can draw nearly 200mA on a cold start at low temperature, equipment using these devices should be designed to furnish enough power for startup; it may be incompatible with low-power systems. It should also be noted that, if this type of device is intended to operate over a wide temperature range—say, +25°C to +125°C—the average *overall* temperature coefficient is

$$\frac{0.5\,\text{ppm/°C}(85°\text{C}-25°\text{C}) + 10\,\text{ppm/°C}(125°\text{C}-85°\text{C})}{100°\text{C}} = 4.3\,\text{ppm/°C}.$$

20.2.3 BANDGAP REFERENCES

Another popular design technique for voltage references in monolithic circuitry is the "bandgap" approach, based on an underlying physical property of base-emitter voltage in a forward-biased silicon transistor.

Most designers are familiar with the approximate −2 mV/°C temperature coefficient of the base-emittter voltage when biased into the active region. The exact value depends on the value of base-emitter voltage, but it is so repeatable for a given transistor that it can be used for linear temperature sensing—if the emitter current is made proportional to temperature, the temperature coefficient of V_{BE} is nearly constant over a wide temperature range.

If, as in Figure 20.12, the V_{BE} values for several devices are plotted as a function of temperature and extrapolated to absolute zero (−273.2°C), the straight lines would have different slopes—but all would intersect at the same voltage value: 1.205 V. This is the *bandgap voltage of silicon* at 0 kelvin. If one could generate a voltage that *increases* proportionally with temperature at the same rate V_{BE} decreases (for a given transistor) the sum of the two voltages would be a constant 1.205V at any temperature. This voltage can be obtained by amplifying the difference between the V_{BE}'s of similar transistors operating at different values of current density (J_1 and J_2, equal to I_1/A_1 and I_2/A_2, in amperes/meter^2). Then

$$\Delta V_{BE} = \frac{kT}{q} \ln \frac{J_1}{J_2} = \frac{I_1}{I_2}\frac{A_2}{A_1} \tag{20.3}$$

k/q is the ratio of Boltzmann's constant to the unit of electronic charge—86.14μV/K. T is absolute temperature, in kelvins, and ΔV_{BE} is the difference of the base emitter voltages, which is proportional to absolute temperature if J_1/J_2 is constant. It is then scaled up to a value that, when summed with V_{BE} at the same temperature, equals 1.205V and is theoretically independent of temperature.

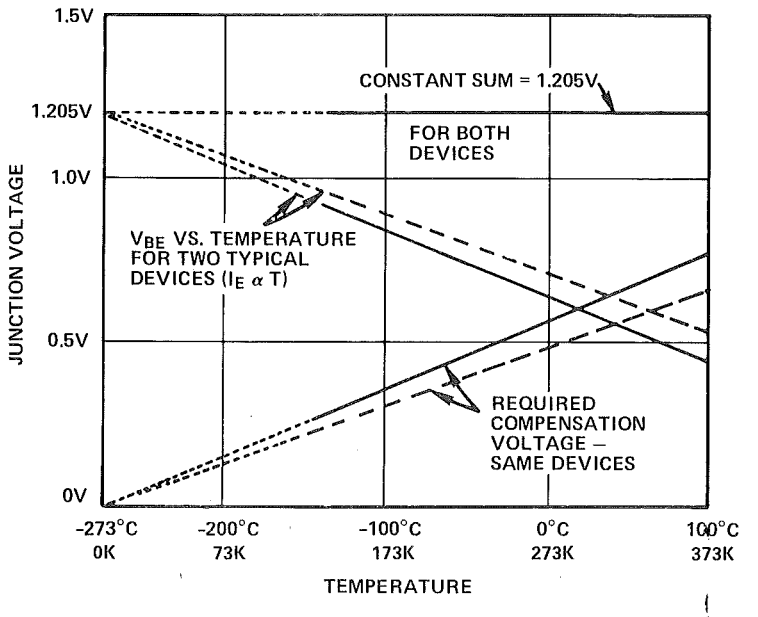

Figure 20.12. Extrapolated variation of base-emitter voltage with temperature (I_E proportional to absolute temperature), and required compensation, shown for two different devices.

Circuit Operation

Bandgap reference circuits are based on the circuit of Figure 20.13. If the amplifier is a high-gain op amp, its two inputs are kept at the same voltage by feedback to the base of Q1 through voltage divider, R4-R5. If R8 and R7 have equal resistances, then equal currents must flow through them and therefore through the collector and emitter circuits of high-β transistors, Q1 and Q2.

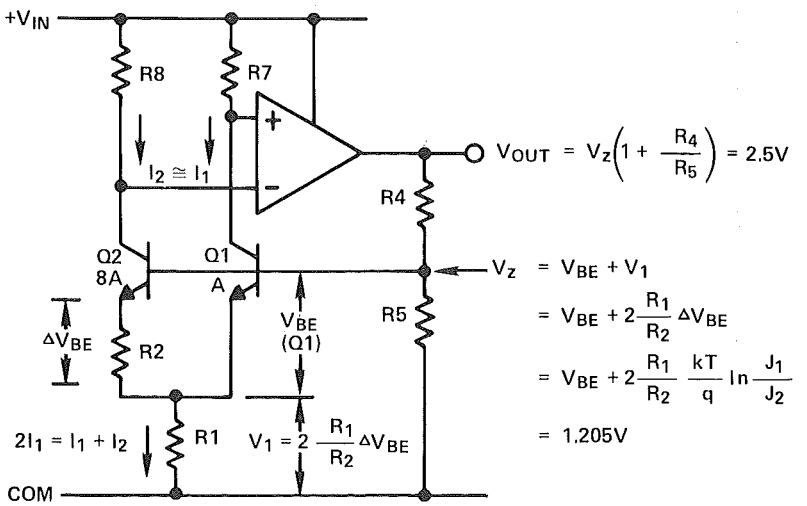

Figure 20.13. Basic bandgap-reference regulator circuit.

The emitter area of Q2, in this example, is eight times as great as the area of Q1, therefore the current density, J_2, is ⅛ J_1, and ΔV_{BE} is (kT/q) ln 8, or 179.2 T microvolts. Since R2 is connected between the emitters, the voltage across R2 is also equal to ΔV_{BE}. The design value of R_2 is determined by the desired level of current through Q1 and Q2; that current, equal to $\Delta V_{BE}/R_2$, also flows through R1, and, since $I_1 = I_2$, the total current is 2 $\Delta V_{BE}/R_2$, and the voltage across R1 is

$$V_1 = 2\frac{R_1}{R_2}\frac{kT}{q}\ln 8 \tag{20.4}$$

For the proper choice of the ratio, R_1/R_2, the sum of the voltages, $V_{BE1} + V_1$, will be equal to the bandgap voltage, 1.205 V, which in turn will be amplified by the ratio, $R_4/R_5 + 1$, to give the desired value of output voltage.

Integrated-circuit process control makes V_{BE} predictable, so that R_1/R_2 can be predetermined and implemented with stable, low-tracking-tempco, thin-film resistors, deposited on the silicon chip and laser trimmed for increased accuracy.

Monolithic bandgap-reference circuits are available in several common forms: some can be used as three-terminal voltage-output regulator-amplifier circuits, others can be used as two-terminal synthetic Zener diodes, and many types can serve in either application. Because it is an integrated circuit, a bandgap reference may also be integrated on-chip with the converter for which it serves as the reference; an example is the Analog Devices AD558 8-bit DAC.

Using Bandgap References

Bandgap references are often useful substitutes for Zener diodes and circuits employing Zener diodes. Bandgap devices will operate from low voltage supplies (typically $V_{out} + 2V$), compared to the breakdown voltage-plus-current-generation "headroom" required to set the proper operating current for the Zener. Bandgap references typically have ten times lower output impedance than low-voltage Zener diodes and can be obtained in a variety of nominal output voltages, ranging from 1.2 to 10 V.

Integrated-circuit designers are able to incorporate additional features in bandgap circuits to make them even more useful. For example, the feedback circuit of the Analog Devices AD584 has a multi-tap voltage divider, with the terminals brought out, so that external jumpers can be used to program a variety of fixed calibrated voltages (2.5 V, 5.0 V, 7.5 V, and 10 V), at currents up to +10 mA at +25°C, or +5 mA over the temperature range, and external precision resistors can be used to set arbitrary voltage values.

The AD584 also has a "strobe" terminal, which permits the device to be turned on or off. When used in this mode, the reference output can be switched to zero with an external signal. When the AD584 is used as a

reference for a power supply, the supply can be switched off to conserve energy, via the AD584's strobe terminal.

A typical application of a reference with a multiplying d/a converter is shown in Figure 20.14. The reference output (shown here programmed for +10V)

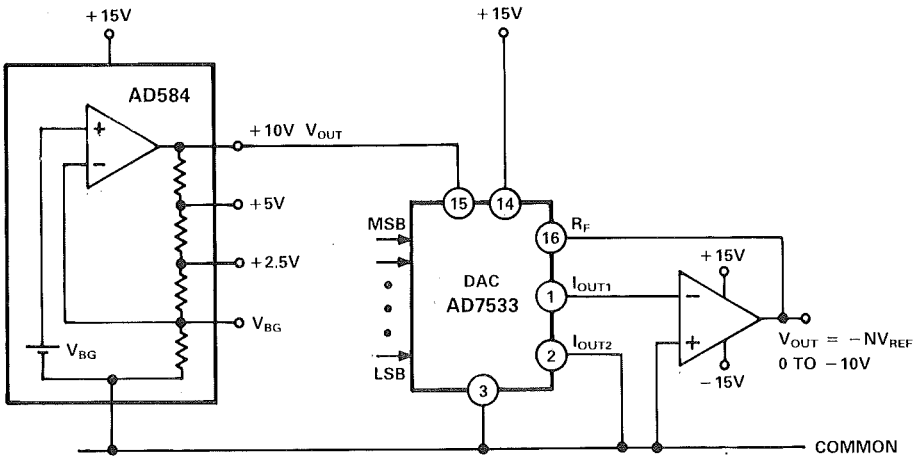

Figure 20.14. Low-power d/a converter, using bandgap reference.

is connected, in the regulator mode, to the reference-input terminal of the DAC. The digital input to the DAC sets the gain, and the op amp inverts the converted output; a +10-volt reference produces a 0 to −10-volt output range. Key specs for the op amp may include speed, dc stability, and/or low power drain.

By using a bandgap reference, such as the AD584, in the two-terminal "Zener" mode, a variety of useful circuits can be derived:

Figure 20.15 shows the device connected as a −5-volt reference. When it is

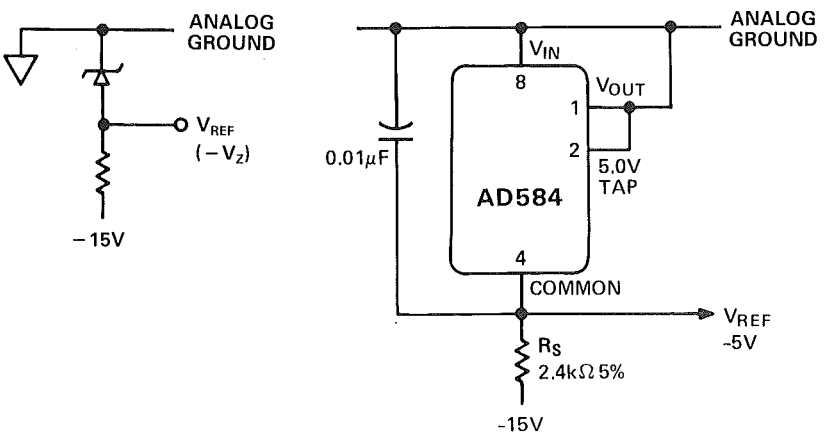

Figure 20.15. Two-terminal −5-volt reference.

used with a multiplying d/a converter and an inverting op amp, the output voltage will be positive. The AD584 has better stability and lower output impedance than a Zener diode.

Figure 20.16 shows how a device of this type can be used as a precision current limiter. The current drawn through the AD584 is equal to the load current (V_{OUT}/R) plus a fixed quiescent current of the order of 0.75 mA (1 mA max). In practice, R_{LOAD} could be adjusted to set the precise value of current.

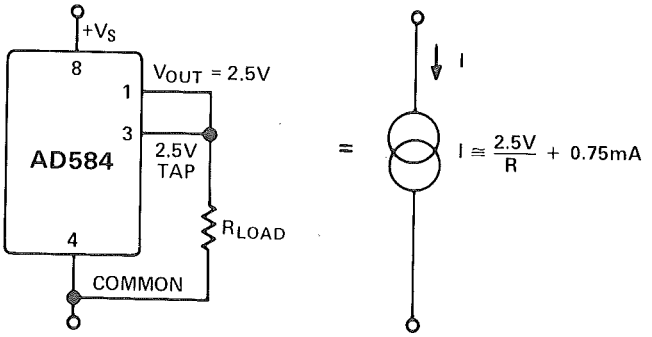

Figure 20.16. Two-component precision current limiter.

It is also possible to use the AD584, with a few external components, to provide a very low-level current source, ignoring the quiescent current. In Figure 20.17, the follower-connected op amp maintains the voltage at its negative input equal to common by driving the common terminal at the load voltage. Since the AD584's output is a constant 2.5V, the current through R_{SENSE} must be equal to $2.5V/R_{SENSE}$, and essentially all of it flows through the load (the FET-input AD547 has very low bias current). The quiescent current

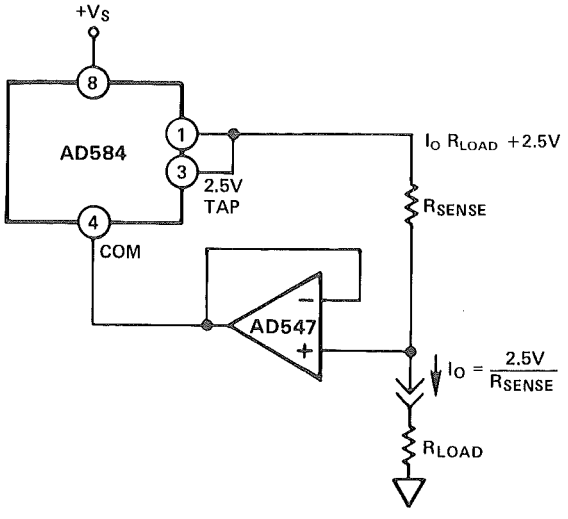

Figure 20.17. Precision low-level current source.

of the reference—whatever its value—is furnished by the op amp's output circuit.

If more than the rated 10-mA (25°C) output current is required, a simple inside-the-loop booster circuit may be added without significantly degrading performance. Figure 20.18 shows such a circuit, capable of furnishing 10 volts

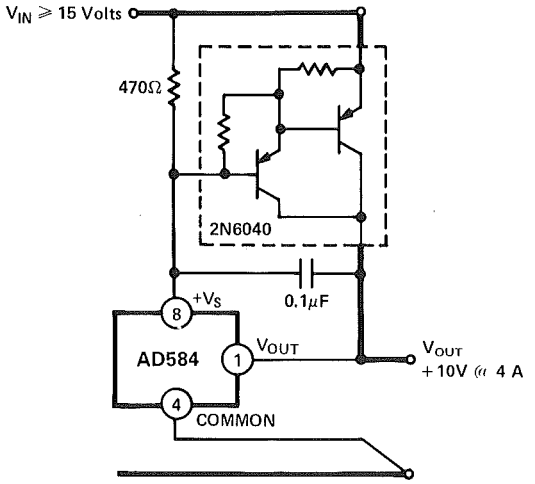

Figure 20.18. High-current precision supply.

at up to 4 amperes. The 2N6040 should be properly heat-sinked, and leads should be short—and of reasonable gauge.

When used with an additional op amp, and no external resistors, a tracking ±5-volt reference circuit can be created, as shown in Figure 20.19. In this circuit, the op amp must keep the +5-volt tap at output common by driving the regulator's COM terminal at −5 V. The 10-volt tap must then be at a voltage equal to +5 volts, with respect to output ground.

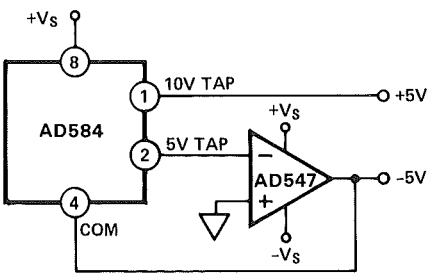

Figure 20.19. ±5-volt tracking reference.

20.3 WHAT THE SPECIFICATIONS MEAN

How do you choose the right reference for your needs? Ideally, a precision reference provides an output voltage that is stable with temperature, input voltage, varying load, and time. If the reference is in error, it will create errors in the device, instrument, or system it is connected to.

In A/D converters, the digital output number depends on the ratio of the quantized input to the "full-scale" reference. If the reference is allowed to (say) increase, the digital output, which is proportional to the ratio of the analog input signal to the reference signal, will decrease by the same percentage.

A D/A converter can be thought of as a digitally controlled potentiometer that produces an analog output (voltage or current) proportional to the product of the reference and the digital input. If the reference increases by 1%, the output will increase by 1%. If the reference is used with a 12-bit converter, that change will create a gain error of about 41 least-significant bits.

Measurement systems of any kind rely on precision references in order to establish a basis for absolute measurement accuracy. Reference errors translate into system errors.

Key sources of reference error include the initial error, and changes of error with line, load, time, and temperature. They lead to the five basic accuracy specifications for any IC or hybrid Zener or bandgap reference used in data-acquisition applications, highlighted in Figure 20.20.

SPECIFICATIONS (typical @ V_S ±15V after a 5 minute warm-up at +25°C, no load condition unless otherwise specified)

Model	AD2710KN	AD2710LN
ABSOLUTE MAXIMUM RATINGS		
Input Voltage (for applicable supply)	±18V	*
Power Dissipation @ +25°C	300mW	*
Operating Temperature Range	0 to +70°C	*
Storage Temperature Range	−55°C to +100°C	*
Lead Temperature (soldering, 20s)	+260°C	*
Short Circuit Protection (to GND)	Continuous	*
OUTPUT VOLTAGE ERROR		
+25°C	±1.0mV max	*
OUTPUT VOLTAGE TEMPERATURE COEFFICIENT		
+10V Output +25°C to +70°C	±2ppm/°C max	±1ppm/°C max
0 to +25°C	±5ppm/°C max	*
LINE REGULATION		
V_S = ±13.5 to ±16.5	125μV/V(200μV/V max)	*
OUTPUT CURRENT	10mA	*
LOAD REGULATION		
I_O = 0 to ±5mA	50μV/mA(100μV/mA max)	*
OUTPUT RESISTANCE	0.05Ω	*
INPUT VOLTAGE		
Operating Range	±13V to ±18V	*
Specified Performance	±13.5V to ±16.5V	*
QUIESCENT SUPPLY CURRENT		
V_S+	9mA(14mA max)	*
NOISE		
0.1 to 10Hz	30μV p-p	*
LONG TERM STABILITY		
T_A = +25°C	25ppm/1000 Hours	*

NOTES
*Same as AD2710KN.

1. INITIAL ACCURACY, UNTRIMMED
2. ACCURACY OVER TEMPERATURE
3. ACCURACY WITH VARYING INPUT VOLTAGE
4. ACCURACY UNDER VARYING LOAD CONDITIONS
5. ACCURACY OVER TIME

Figure 20.20. Key reference specifications.

20.3.1 INITIAL ACCURACY

Initial accuracy, or output voltage error, or output voltage tolerance, is the deviation from the nominal output voltage at 25°C and specified input voltage. It should be measured by a device traceable to a fundamental voltage standard. Initial accuracy is trimmable on some devices, such as the AD584 and the AD2700 series. Fine-adjust pins are provided for use with an external potentiometer (Figure 20.21). Considerable care must be taken, though, as these external components can affect the temperature coefficient.

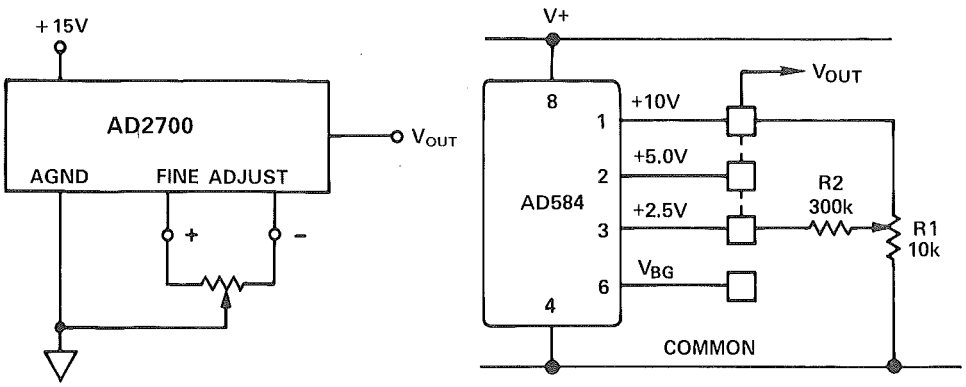

a. Hybrid precision reference.

b. Monolithic pin-programmable precision reference.

Figure 20.21. Trimming the reference output.

20.3.2 OUTPUT VOLTAGE CHANGE WITH TEMPERATURE

Output voltage change with temperature, or output voltage temperature coefficient, or voltage drift. This is the change in output voltage from the value at 25°C ambient. It is independent of variations in other operating conditions. There are two common methods of specifying drift, and the differences between the two must be understood in order to pick the right reference for the job.

The *Box Method* of specifying drift consists of specifying an error band and an equivalent temperature coefficient, usually in ppm/°C. The error band is graphically defined as a box (voltage on the vertical axis, temperature on the horizontal) whose diagonals extend from T_{LOW} to T_{HIGH} through 25°C. The slope of this diagonal is given as the stated temperature coefficient. Thus, the total absolute error for a reference over its specified temperature range is equal to the output voltage tolerance at 25°C plus the error band.

One difficulty with the box method is that manufacturers have defined T_{LOW} and T_{HIGH} in different ways. In some box specifications T_{LOW} is T_{MIN}, T_{HIGH} is T_{MAX} (0°C and 70°C for a "commercial" temperature range, or −55°C and +125°C for the extended "military" range). T_{LOW} has also been defined as 25°C ("ambient") and T_{HIGH} as T_{MAX}. The "box" specification defines the drift by the worst (or biggest) temperature excursion from the initial temperature. It is up to the user to be sure which method is used; for example, a "5-ppm/°C" box from −55°C to +125°C allows 900ppm drift (9 mV, for a +10V reference), while a "5-ppm/°C" box from +25°C to +125°C allows 500ppm (5 mV). If the user is careless, the 5-mV limit could be misinterpreted, in the latter case, to apply for −55°C to +125°C.

A second trap of the box method relates to the location of the voltage "origin"; i.e., is the box measured from $V_{nominal}$ at 25°C or V_{actual} at 25°C? In the above example (5ppm box, 25 to +125°), drift of 5 mV is theoretically allowed. To translate this to an absolute accuracy spec one must know if this drift is measured from $V_{nominal}$ or added to the initial accuracy error. If measured from V_{actual} (at 25°C), the box floats up or down depending on the initial value, and the device meets spec if V_{out} at T_{MIN}, $T_{ambient}$, and T_{MAX} all fall within the box. Figure 20.22 shows an extreme example of this case.

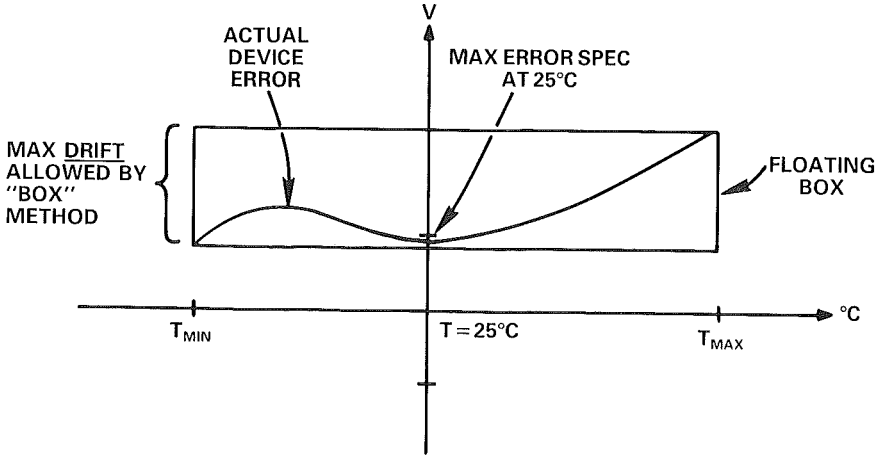

Figure 20.22. Floating box drift specification.

To remove some of the ambiguity of the box method, Analog Devices uses a *"modified" box* method (Figure 20.23) that fixes the box at $V_{nominal}$ and specifies only the absolute error at a temperature rather than initial error plus allowed drift. The AD2700 series is specified in this manner. The "modified" in modified box refers to the notch at 25°C, which tightens the initial error band allowed.

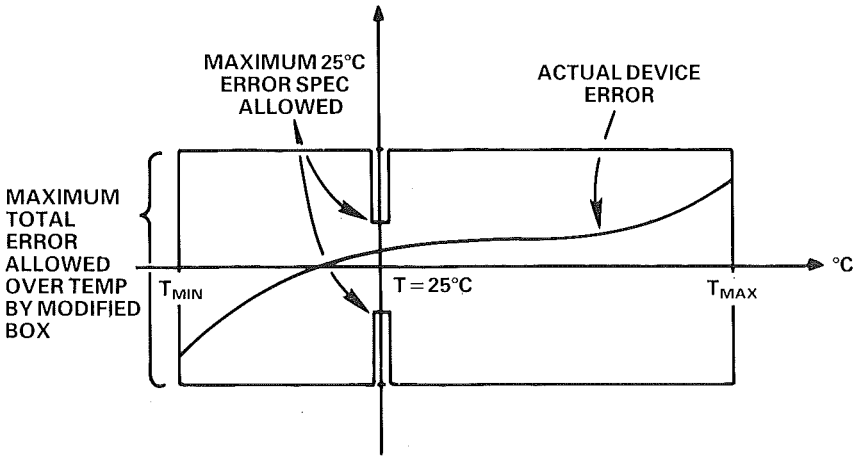

Figure 20.23. Modified box drift specification.

The third method of specifying output voltage change with temperature is the *butterfly* method. This method can be the tightest specification method, as it ties down the maximum excursion or change with temperature and can be extrapolated to any temperature within specification (rather than the end points only). Again, the danger is that the center of the butterfly may be either $V_{nominal}$ or V_{actual} at $T_{ambient}$ (25°C). Figure 20.24 demonstrates the butterfly method, using $V_{nominal}$.

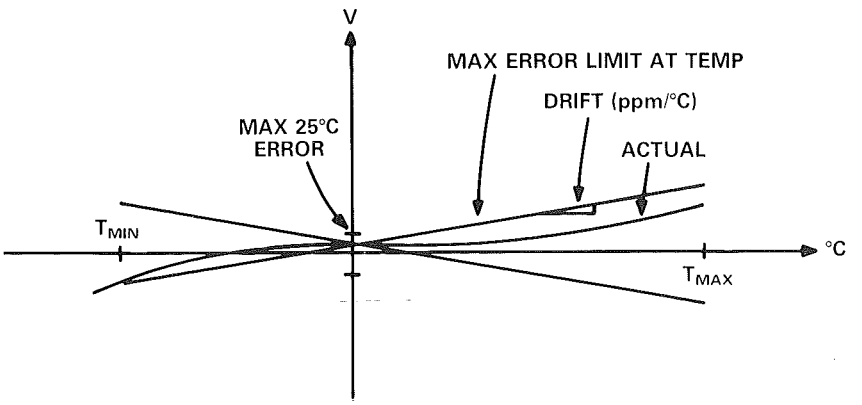

Figure 20.24. Butterfly method of specifying drift.

BOX METHOD	MODIFIED BOX	BUTTERFLY
Box Floats	Box fixed	Can float or be fixed
Limits absolute drift error by end points	Limits total absolute error	Limits absolute drift error over the range
Tends to ignore initial error	Includes initial error	Includes initial error

An example of a device with a tight butterfly drift specification is the AD2710/2712 series of precision references. Top grades are specified for ± 1-millivolt initial error (maximum) and ± 1ppm/°C maximum drift error (25° to 70°). Figure 20.25 shows the typical and maximum drift error of these references.

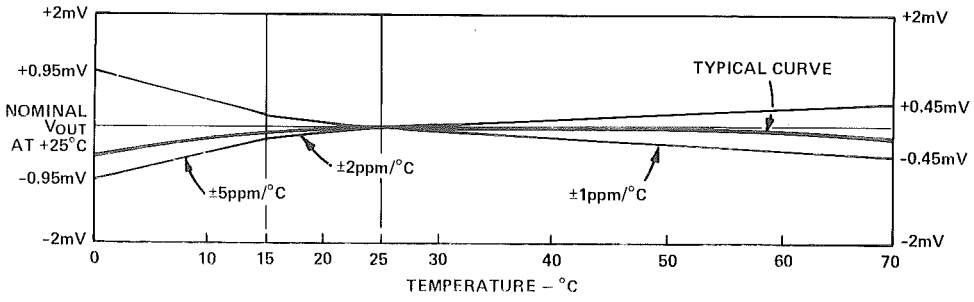

Figure 20.25. Maximum change of +10-volt output from +25°C value vs. temperature.

20.3.3 LINE REGULATION

Line regulation, or accuracy with varying input voltage is the change in output due to a specified change in input voltage, usually specified as % per volt or μV per volt of input change. It is a measure of power-supply rejection and is typically measured at dc. A related specification, ripple rejection, (ratio of ripple components in the output to residual upstream ripple) is rarely specified for IC references; it therefore behooves the careful reference user to choose an upstream power supply with minimal ripple.

20.3.4 LOAD REGULATION

Load regulation, or accuracy under varying load conditions, is the change in output voltage for a specified dc change in load current. It is generally expressed in μV/mA, and sometimes as ohms of dc output resistance. It includes the effect of self-heating due to increased power dissipation at high load currents.

20.3.5 LONG-TERM STABILITY

Long-term stability, or accuracy over time, is usually specified as parts-per-million per 1000 hours at a specified temperature. This is a difficult spec to verify, and is generally given as "typical", based on characterization data.

Zener diodes undergo the major portion of their long-term drift during the first part of their life. They tend to settle down with age, finally reaching a point where only small random variations compatible with 1/f noise occur. Unfortunately, it may take years to arrive at this condition (Figure 20.26). For this reason, most Zeners used in precision references are aged ("burned in") at an elevated temperature to accelerate the process.

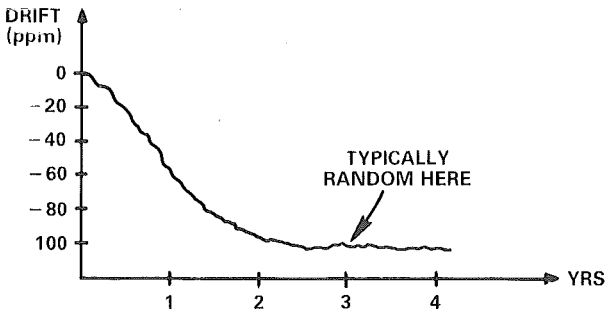

Figure 20.26. Typical 1N829 drift with time.

Chapter Twenty-One

Digital Signal Processing

Digital signal processing (DSP) means the processing of (analog)* signals using digital techniques (i.e., digital hardware and software). While *any* digital processing of signals that originate in the analog world, and have at some point been converted to digital, would qualify under this broad definition, the term has come to be used in a much more specific way—as it will be in this chapter: DSP is the application of fast, specialized hardware, sophisticated algorithms, and the appropriate software for the purpose of manipulating large amounts of data associated with extracting and processing analog-based information in essentially "real time."

The emergence of DSP hardware is changing the role of analog-to-digital conversion in today's signal processing systems. In early days, all processing of a signal, with the goal of obtaining results with sufficient speed to be useful in real time, was of necessity handled by analog components. The principal destinations for analog signals converted to digital format, after substantial analog processing, were off-line computation, data storage, and hard-copy tabulation, rather than real-time instrumentation, computation, and control.

Now, however, system designers have an incentive to perform the signal conversion as early in the loop as possible (see Figure 21.1). The reason for this is that much or all of the required signal processing can be handled by fast, flexible digital components that allow high-performance DSP routines to be

*It should be noted that any electrical signal is by nature an analog signal, even if it represents a digital "1" or "0". This can be understood if one pictures what happens to a chain of 1s and 0s returned from beyond the orbit of Jupiter, buried in cosmic noise: the signal must be received, amplified, converted, and processed (digitally) to reconstruct the original digital information; but until that digital information has been identified, the signal—to all intents and purposes—is purely analog in origin.

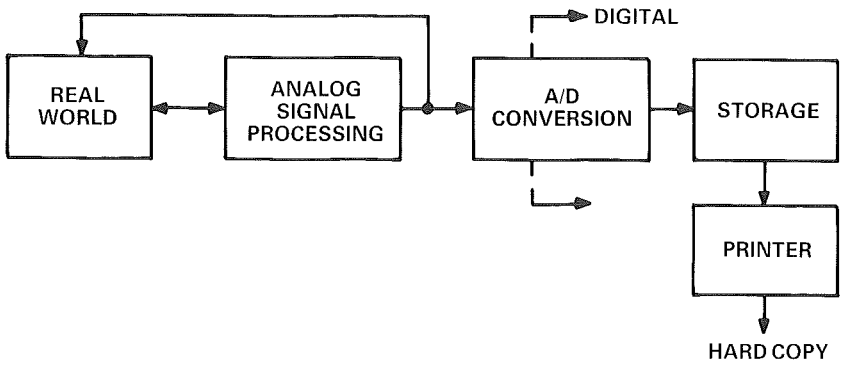

a. Before DSP: digital used principally for record-keeping.

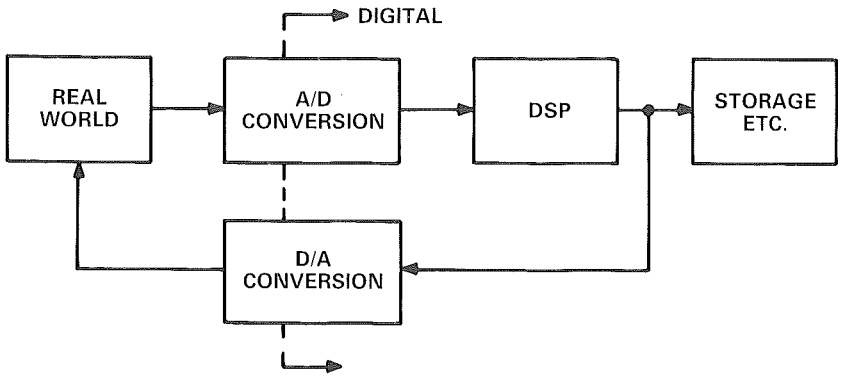

b. DSP era: digital responds to and operates on real-world phenomena.

Figure 21.1. Analog and digital signal processing.

implemented more accurately, reliably, and flexibly than with analog circuitry, yet, in many cases, with sufficient speed to interact in real time.

This chapter first reviews basic signal processing tasks, giving emphasis to the general role played by DSP. Two key DSP algorithms are examined in some detail—digital filters and spectral analysis. The basic hardware required to perform DSP is described. Finally, some applications that exemplify DSP's advantages are reviewed.

21.1. SIGNAL-PROCESSING BASICS

Signal processing revolves around two basic tasks—digital filtering and spectral analysis. *Filtering* smoothes, removes noise from, selects particular signal components from, or predicts future values of an incoming signal. A time-domain signal can be interpreted as a weighted combination of purely sinusoidal spectral components; *spectral analysis* determines the weights corresponding to each frequency in the spectrum.

Signal-processing applications span many areas, including speech analysis and synthesis, telecommunications, instrumentation, radar and sonar, and—using multi-dimensional techniques—graphics and imaging. For example, filtering is used to minimize high-frequency noise and the low-frequency hum in telephone-line transmission. Spectral analysis is used to determine the formant content of incoming speech for recognition. Two-dimensional filtering improves the clarity of a satellite image.

Filtering and spectral analysis have traditionally been implemented with analog components. Filtering is carried out by passing the signal through a circuit consisting of resistors, capacitors, op amps, and/or inductors; the precise configuration of these components and the relationship of the magnitudes of their parameters determine the filter's characteristics. Multiple analog filters—each passing energy in a narrow band—can be cascaded for sharpness and banked together to perform spectrum analysis.

Analog-based signal processing has numerous advantages, including low component cost, the ability to handle wide bandwidths in real time, the availability of pre-packaged modules and ICs, and a large existing base of knowledge. However, analog components introduce noise at each stage; and filter characteristics—requiring effort to tune initially—are sensitive to the effects of temperature and aging. In addition, multi-stage filters pose subtle design challenges. Because coefficients and configurations—once established—tend to be inflexible, signal-processing hardware using analog parts generally is restricted to performing a narrow, dedicated task.

In response to the limitations of analog-based processing, the digital processing of signals has emerged as an alternative. The next section demonstrates how signal-processing tasks—including filtering, spectral analysis, and a host of others—can be carried out with digital arithmetic operating on digitized data. Recent advances in VLSI (very large-scale integration) now make it feasible to perform real-time digital signal processing with just a handful of ICs. The advantages conferred upon a system by such DSP hardware are dramatic—substantially improved performance, stability, and flexibility. Just as digital computers supplanted analog computers two decades ago in general-purpose computing applications, DSP is strongly challenging analog circuit configurations in real-time processing.

Our discussion of spectral analysis and digital filtering will benefit from a brief discussion of DSP nomenclature (there is also a brief glossary at the end of this chapter). Following Figure 21.2, an incoming analog signal is digitized, with the sampled data output points denoted x_i, or x(i). The index, i, corresponds to the discrete sampling time. This sampled data is stored in a buffer and operated on by DSP hardware. The DSP algorithm determines the sequence in which data and coefficients are accessed and how they are processed. In the cases below, the computational outputs are spectral weights or filtered sampled data.

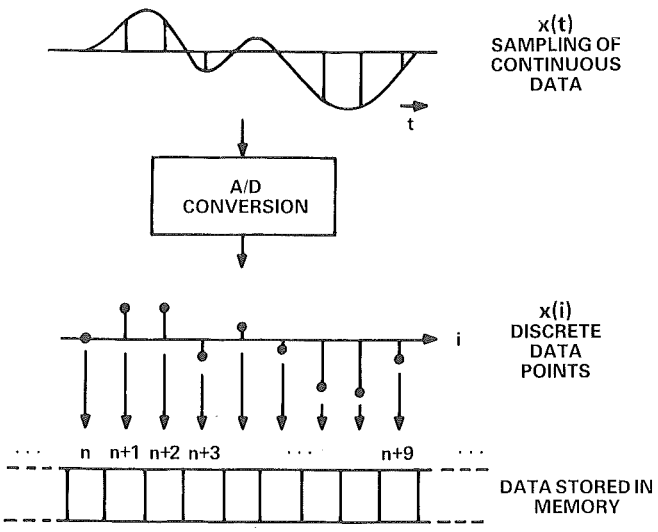

Figure 21.2. In DSP, continuous data is replaced by sampled data and continuous time by discrete time.

21.1.1. SPECTRAL ANALYSIS

The departure point for a discussion of spectral analysis is the Fourier transform equation pair:

$$x(t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} X(\omega)\, e^{j\omega t}\, d\omega \tag{21.1}$$

$$X(\omega) = \int_{-\infty}^{\infty} x(t)\, e^{-j\omega t}\, dt$$

where t is time, ω is angular frequency, $(2\pi f)$, x(t) is the signal—a function of time—and $X(\omega)$ is its counterpart in the frequency (spectral) domain. These equations give us, at least formally, the mechanics for taking a signal's time-domain representation and resolving it into its spectral weights—called Fourier coefficients.

Since the Fourier equations require continuous integrals, they have only indirect bearing on digital processors. However, under certain circumstances, a sampled (digitized) signal can be related faithfully to its Fourier coefficients through the discrete Fourier transform (DFT):

$$x(n) = \frac{1}{N}\sum_{k=0}^{N-1} X(k)\, W_N^{-kn} \tag{21.2}$$

$$X(k) = \sum_{n=0}^{N-1} x(n)\, W_N^{kn}$$

Provided that the signal is sampled frequently enough (at a rate ≧ twice its highest frequency component), and assuming that the signal is periodic, the above DFT equations hold exactly. What is most interesting from the standpoint of DSP is that the DFT equation provides us with a means to estimate spectral content by digitizing an incoming signal and simply performing a series of multiply/accumulate operations.

To see qualitatively why the DFT equation yields spectral information, consider Figure 21.3. A time signal is superposed on a spectral "template" at various frequencies. In the first case, for frequency ω_1, the input signal and the spectral template have little relationship; as a result, the positive products are more or less cancelled out by negative products. The net effect is that the summation in the DFT equation indicates little spectral energy of frequency ω_1 in the input signal. In the second case, however, a reinforcing pattern emerges; the signal and the template tend to be positive or negative concurrently—producing a positive product nearly everywhere. Thus, the sum of the products will be a large positive number, indicating that the incoming signal has significant energy of frequency, ω_2.

Unfortunately, the large number of multiplications required by the DFT limits its use in real-time signal processing. The computational complexity of

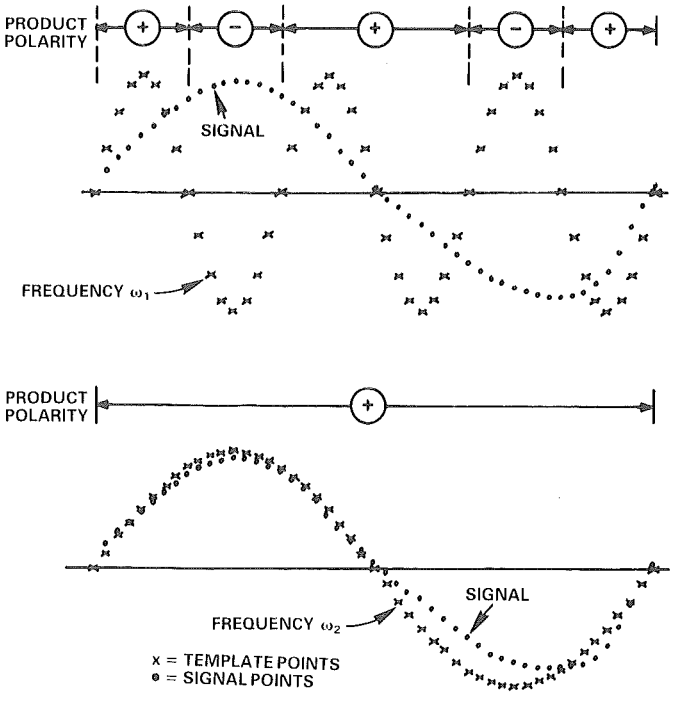

Figure 21.3. Input signal compared with sinusoidal "template" to measure frequency content of input signal. ω_1 is at a much different frequency, ω_2 is at very nearly the same frequency as the signal's fundamental.

the DFT grows with the square of the number of input points; to resolve a signal of length N into N spectral components requires N^2 complex multiplications ($4N^2$ real multiplications*). Given the large number of input points needed to provide acceptable spectral resolution, the computational requirements of the DFT are prohibitive for most applications.

The fast Fourier-transform (FFT) algorithm produces results identical to those of the DFT but reduces computation requirements by several orders of magnitude. The FFT achieves its economies by exploiting computational symmetries and redundancies that exist in computing the DFT. The availability of the FFT makes spectral analysis feasible, at virtually real-time rates.

The basic evolutionary process from DFT to FFT is demonstrated by simply segregating the odd and even terms to break an N-point DFT equation's summation into the sum of two (N/2)-point sub-series:

$$\begin{aligned} X(k) &= \sum_{n\,\text{even}} x(n)\, W_N^{nk} + \sum_{n\,\text{odd}} x(n)\, W_N^{nk} \\ &= \sum_r x(2r)\, W_{N/2}^{rk} + W_N^{k} \sum_r x(2r+1)\, W_{N/2}^{rk} \end{aligned} \tag{21.3}$$

N/2–point DFT N/2–point DFT

In this way, one N-point DFT has been reduced to two N/2-point DFTs. While this decomposition may not seem material, the key is that computational requirements for the DFT grow with the square of the number of points involved. Breaking the DFT in half decreases the number of complex multiplications from N^2 to $2x(N/2)^2$. Each of the above (N/2)-point DFTs can in turn be partitioned into two (N/4)-point DFTs, again reducing computation.

Proceeding in this fashion, an entire N-point DFT can be reduced (for N a power of two) to a series of elementary operations, called butterflies. A butterfly is a two-point DFT, along with a multiplication by a complex rotation factor (see Figure 21.4). The net effect of this systematic decomposition is to

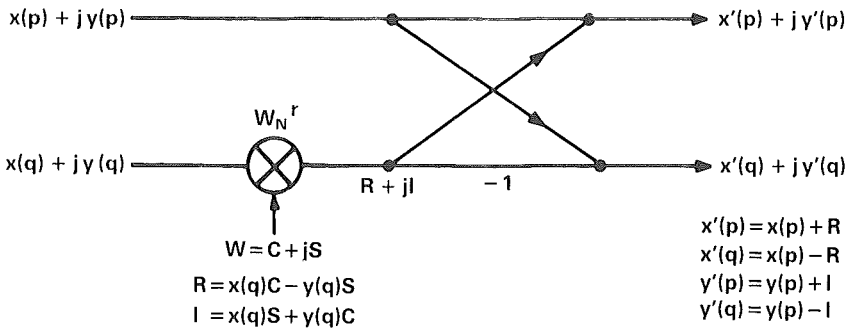

Figure 21.4. Fast Fourier transform radix-2 "butterfly." An N-point transform contains (N/2) $\log_2$ (N) of these operations.

*$(A + jB)(C + jD) = (AC - BD) + j(AD + BC)$; $j = \sqrt{-1}$

reduce the total number of complex multiplications from N^2 to $N/2 \log_2 N$ for executing an N-point DFT. For example, a 1,024-point DFT would require more than 1,000,000 complex multiplications, while the corresponding FFT requires only $512 \times 10 = 5{,}120$ complex multiplications.

Figure 21.5 illustrates how an FFT resolves a signal into its spectral components—and the effect of FFT length on spectral resolution. In all three cases, the same input signal is examined. In the first case, we perform a 64-point FFT on the first 64 sample points; the second and third cases perform 256- and 1,024-point FFTs on the first 256 (1,024) data points. The differences observed in spectral resolution underscore a key principle—the longer the time period in which a signal is observed, the sharper the spectral resolution that can be attained.

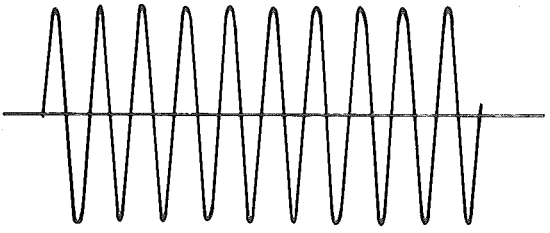

a. Input data a sine wave. Spectrum should be single-frequency line.

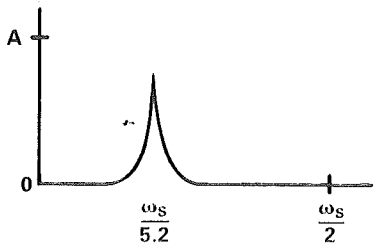

b. 64-point transform.

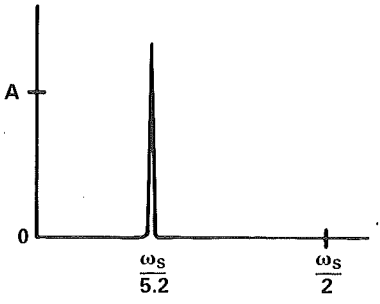

c. 256-point transform.

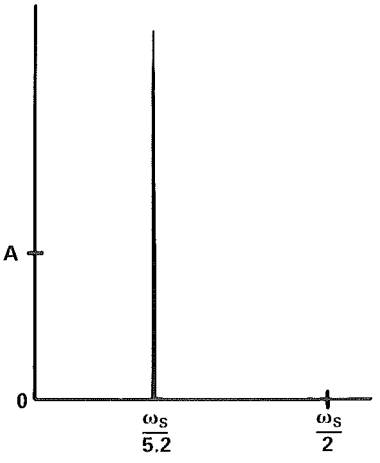

d. 1,024-point transform.

Figure 21.5. Fourier transform: effect of number of points on computed spectrum.

Bringing digital hardware to bear on a spectral-analysis task has numerous advantages. With a long-enough window of data, it can provide very precise spectral resolution. Moreover, the system can be flexibly programmed to vary the FFT size dynamically, according to the spectral resolution needed. Finally, once the data is digitized, it is possible to perform additional DSP tasks, such as spectrum averaging, to further improve FFT performance.

21.1.2. DIGITAL FILTERING

Digital filters have performance attributes similar to those of analog filters—ripple in the passband and attenuation in the stopband. What distinguishes digital filters is their ability to provide arbitrarily high performance. For example, the rolloff slope (i.e., the rate at which the filter makes a transition from the passband to the stopband) can be made virtually as steep as is desired. In general, it is straightforward to design a digital filter that easily out-performs the most complicated analog designs.

The fundamental digital filtering equation is:

$$y(n) = \sum_{i=0}^{N-1} h(i)\,x(n-i) \; + \sum_{j=1}^{M} b(j)\,y(n-j) \qquad (21.4a)$$

The coefficients, $h(i)$, are weighting factors applied to the most recent N sample points; the coefficients $b(j)$ correspond to terms feeding back the M most recent filtered output points. In the case where the feedback coefficients, b(j), are all equal to zero, the digital filter is termed a Finite Impulse Response (FIR) filter.

$$y(n) = \sum_{i=0}^{N-1} h(i)\,x(n-i) \qquad (21.4b)$$

If feedback terms are used, then the filter belongs to the Infinite Impulse Response (IIR) class. These two types of digital filter manifest important tradeoffs, which warrant further discussion.

Finite Impulse Response Filters

In equation 21.4b, the fundamental FIR filter equation, each output point, y(n), is obtained by convolving the past N input points, x(n – i), with a set of coefficients. An FIR filter can be viewed as a tapped delay line (see Figure 21.6); the parameter, N, corresponds to the number of taps of the FIR filter. The number of taps tells us the number of multiply/accumulate operations required to compute this convolution.

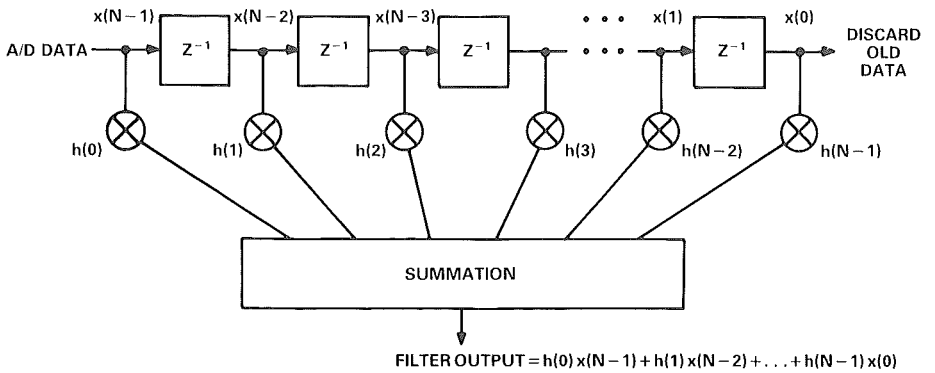

Figure 21.6. FIR filter as a tapped delay line.

The coefficients, h(i), represent the impulse response of the FIR filter. As Figure 21.7 demonstrates, an input of 1 at time 0 (x(0) = 1), and zero at all other times, results in output values equal to h(i) for the periods $i = 0, \ldots, N-1$. Note that the h(i) can be non-zero for only a finite number of time periods, hence the term "finite" impulse response. Since they use no feedback, FIR filters are unconditionally stable.

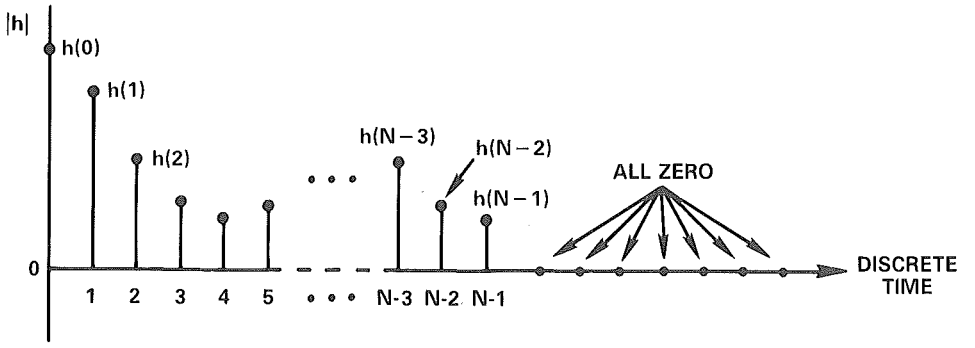

Figure 21.7. FIR-filter impulse response function, h(i).

FIR filters can best be understood in the context of two fundamental relationships. First, a filter's time-domain impulse response, h(i), and its frequency response, H(f), are related via the Fourier transform. Second (a key principle of DSP), multiplication in one domain is equivalent to convolution in the conjugate domain. With respect to FIR filters, this tells us that multiplying the input spectrum by the desired filter transfer function is equivalent to convolving the input time-function with the filter's impulse response in the time domain.

To further amplify the above point, consider Figure 21.8. Figure 21.8(a) illustrates an incoming signal that we wish to low-pass filter. It consists of the sum of two signals at frequencies, f_1 and f_2. Since there are just two frequencies present, its spectrum looks like (b). We'd like to design an FIR filter to filter out f_2, leaving just f_1, as shown vs. time in (e) and frequency in (f).

An ideal low pass filter is suggested in (d); note that multiplying it by the input spectrum in (b) will give the spectral domain representation of a low-pass filtered output, allowing f_1 to pass and completely attenuating f_2. Now, the Fourier transform of (d)'s ideal filter is the sinc function ($\sin x/x$) in (c). Consequently, if the input (a) is convolved with a discretized sinc function, (c), we can directly compute the filtered output signal, as a function of time (e).

More generally, an FIR filter boils down to simply convolving the digitized input signal with the filter's time-domain coefficients, h(i)—an action equivalent to multiplying the frequency representation of the input signal by the filter's transfer function.

Unfortunately, from the perspective of practical implementation, Figure 21.8(c)'s sinc function is infinite in duration. To obtain a filter that can be

implemented, we must somehow truncate the number of coefficients used to represent (c); this can be carried out by discarding the tails—or, more effectively—by multiplying the function by some window. This truncation/windowing, however, makes it impossible to realize (d)'s ideal low-pass filter transfer function, and ripple and rolloff are necessarily introduced (see Figure 21.10).

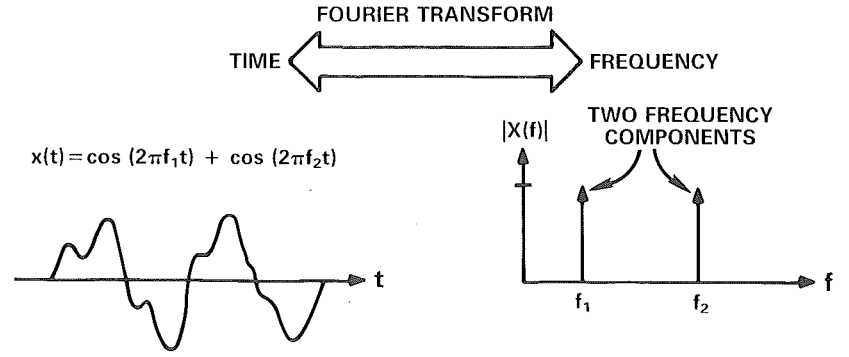

a. Input signal, sum of two sine waves at different frequencies.

b. Spectrum of signal in (a).

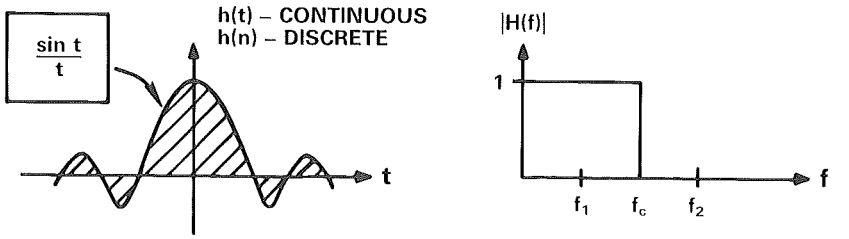

c. Impulse response of desired filter, (sin t)/t function—transform of (d).

d. Desired low-pass filter transfer function.

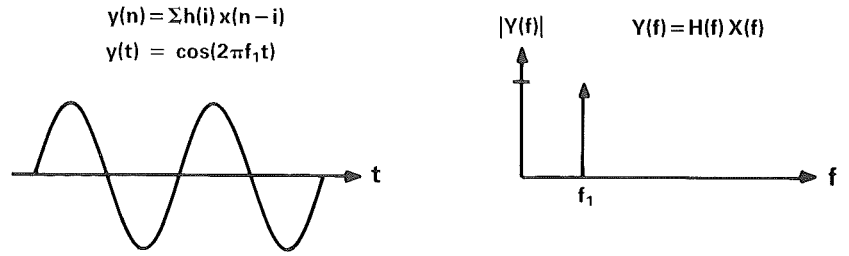

e. Time response of filtered signal, convolution of (a) and (c).

f. Frequency response of filtered signal, product of (b) and (d).

Figure 21.8. Basics of low-pass filter design using Discrete Fourier transform, (a), (c), (e) in time domain, (b), (d), (f) in frequency domain.

The essentials of FIR filters can be illustrated in a very simple example, something many readers may have already used. It is intuitively apparent that noisy data can be smoothed (or filtered) by taking a moving average. For example, noisy laboratory data might be plotted using a moving average of the last five points:

$$y_n = (1/5)[x(n) + x(n-1) + x(n-2) + x(n-3) + x(n-4)].$$

This simple scheme is nothing less than a five-tap FIR filter. Its impulse response is constant for five periods and then drops to and stays at zero. In the spectral domain, the Fourier transform of the impulse response, $h_i = 0.2$, $i = 0, \ldots, 4$, and zero elsewhere, is shown in Figure 21.9. This frequency-domain plot tells us, consistent with our intuition, that low-frequency components will largely be passed, while high-frequency noise will be relatively damped, but not uniformly.

The real value of digital filters is not apparent in an elementary example, such as Figure 21.9's rather sloppy transfer function. By taking an adequate

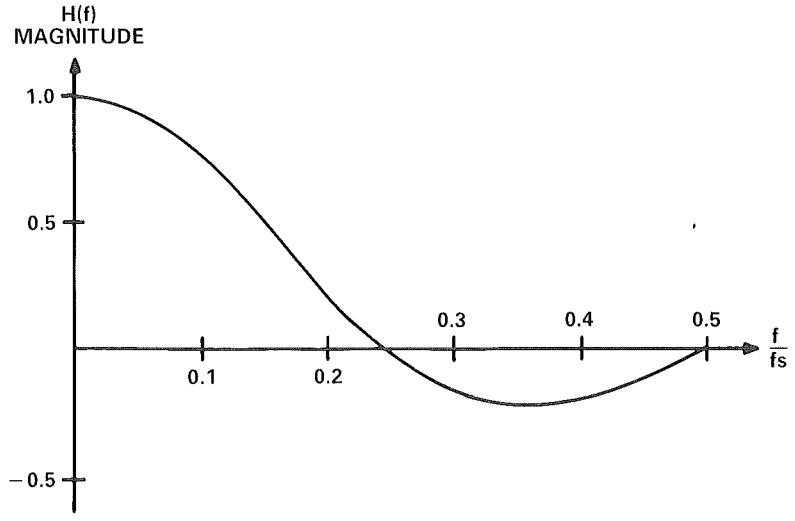

Figure 21.9. Discrete Fourier transform of simple running-average FIR filter with 5 taps.

number of taps and properly choosing the coefficients, an FIR filter can provide excellent discrimination, as the response spectra shown in Figure 21.10 illustrate, for various numbers of taps. In general, the greater the number of taps used in an FIR filter, the better the filter's performance, at the expense of reduced throughput.

In designing FIR filters, tradeoffs must be made among several attributes (i.e., ripple in the passband, ripple in the stopband, width of the transition band, phase distortion, and throughput). These tradeoffs are reflected in the number of coefficients used—and their particular values. This selection can be made directly in the time domain (for example, to implement a pure time

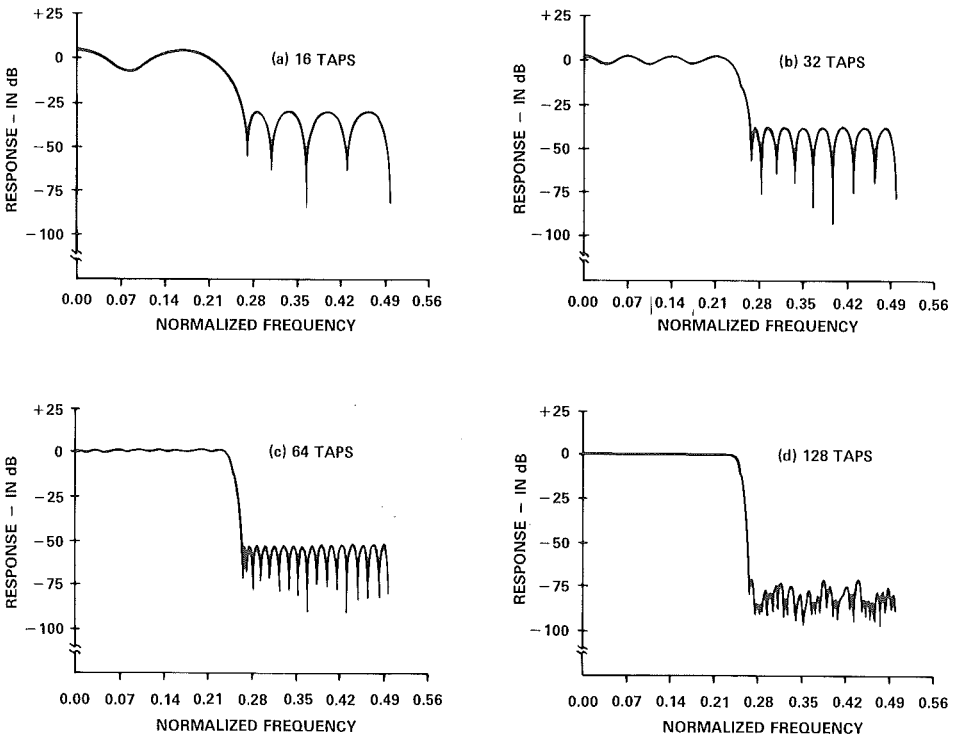

Figure 21.10. Comparison of FIR low-pass filters with (a) 16, (b) 32, (c) 64, and (d) 128 taps. Normalized frequency = $f_{actual}/f_{sampling}$.

delay or an N-point running average), but more commonly is made employing powerful and easy-to-use computer-aided-design (CAD) techniques to determine optimal parameter values for the desired filter performance.[1]

With an FIR filter, there is a direct relationship between the incoming sample rate, the speed of the digital hardware used in the system, and the performance of the FIR filter:

$$\frac{1}{\tau_m} = N \times S \tag{21.5}$$

where τ_m = multiply time
N = number of taps
S = sample rate

For example, a specified performance and incoming sampling rate determine the multiplication rate required of the hardware. Thus, if a filter is to have 100 taps, and the sampling rate is 100 kHz, the multipliers must perform their

[1]While a discussion of these approaches is beyond the scope of this chapter, the interested reader is referred to Rabiner, Lawrence R., and Gold, Bernard, *Theory and Application of Digital Signal Processing*, (Englewood Cliffs, N.J.: Prentice-Hall, 1975).

multiplications within 100 nanoseconds, a speed not difficult to achieve with modern CMOS multipliers.

The highest-performance filters of Figure 21.10 could not be matched by an analog-based implementation. Moreover, these digital filters are straightforward to design and implement in hardware. There are other advantages of digital FIR filters that further increase their desirability. Once designed, they are stable; performance is insensitive to the effects of temperature or aging. In addition, a key consideration is that the filter's performance can be changed simply, just by modifying the number of coefficients used and their values. For instance, a simple software modification would shift a filter's performance from (a)'s to (d)'s—with no change in hardware, except that slightly more memory is used.

Infinite Impulse Response Filters

Infinite Impulse-Response (IIR) filters are the other commonly used digital filter, differing from FIR filters in one fundamental respect: feedback. Because of feedback, the filter's impulse response can continue long after the initial impulse—indeed, for an infinite duration. The use of feedback allows an IIR filter to economize in the number of multiplications required to provide a given filter performance. But this efficiency is not without its costs. As in other recursive systems, input perturbations can ring indefinitely—in some cases causing the filter to be unstable. Also, the accumulated effects of fed-back round-off noise can noticeably degrade the filter's performance.

Equation 21.4a is the basic IIR filter equation, assuming one or more of the coefficients, b(i), is non-zero. Figure 21.11 plots an impulse response function for an IIR filter with a typical set of coefficients. In this case, and in gen-

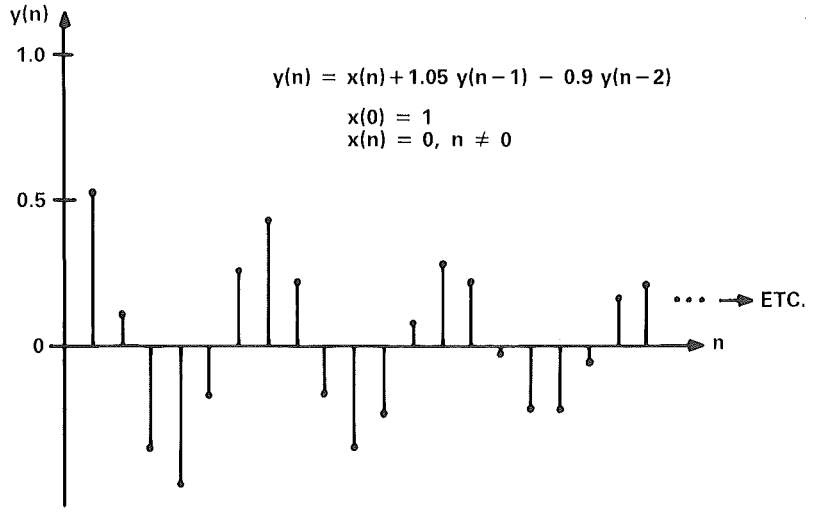

Figure 21.11. A portion of the impulse response of an IIR filter.

eral, the presence of feedback means that the impulse response of an IIR filter never converges to zero—may even diverge—hence *infinite* impulse-response. However, as a practical matter, noise, round-off error, and limited resolution do result in effective convergence when simulating analog filters that settle physically.*

One form of IIR filter that is widely used is the biquadratic, or biquad, form of the generalized IIR equation (equation 21.4):

$$y(n) = a_0 x(n) + a_1 x(n-1) + a_2 x(n-2) - b_1 y(n-1) - b_2 y(n-2) \quad (21.6)$$

corresponding to the continuous transfer function:

$$\frac{A s^2 + B s + C}{D s^2 + E s + F} \quad (21.7)$$

The biquad serves as a building block for IIR filter design. Generally, several biquad sections are cascaded together to obtain the desired filter performance.

The goal in IIR filter design, to match a given transfer-function requirement, is to determine the number of biquad sections to be used, and the values of their coefficients. Two principal IIR design techniques exist. The first considers the transfer functions of conventional analog filters, such as the Butterworth, Chebyshev, or Elliptic; a digital filter is then constructed that provides the same impulse response as its analog counterpart. The second relies on computer-aided-design techniques to arrive at an optimal IIR implementation. In this context, "optimal" means that the number of terms needed to meet a specified performance specification is minimized.[2]

An example that demonstrates the efficiency of an IIR implementation is a comparison of an FIR and IIR implementation of a 70-dB stopband attenuation filter. To achieve this performance, an FIR filter would require nearly three times as many multiplications-per-second as an IIR implementation. These performance advantages, however, require tradeoffs to be made in other key respects, as summarized below:

	IIR	FIR
Performance/Throughput	Higher	
Ease of Design		Easier
Filter Stability	Sensitive	Unconditional
Round-off Noise	Sensitive	Insensitive

*Even a simple single-time-constant R-C analog filter will theoretically take an infinite time to reach its asymptotic steady-state condition, but in practice it settles, for example, to 1 LSB of 32 bits within 23 time constants—and so does its IIR-filter equivalent.

[2]A more detailed discussion of IIR filter design can be found in Rabiner and Gold, *op. cit.*[1]

21.1.3 OTHER DSP ALGORITHMS

DSP is not limited to FFTs and digital filters. In fact, one of the prime advantages of DSP is that, once the data is digitized, fast digital hardware can perform a broad range of tasks. Commonly used DSP routines include modulation/demodulation (heterodyning), waveform generation, correlation, estimation, control, power spectrum calculations, and multi-dimensional transforms. While a discussion of these areas would take us far afield of this chapter's focus, their breadth points to an important advantage of DSP—system flexibility. By converting signals early and incorporating fast multiply/accumulate hardware to perform digital filtering and/or spectral analysis, the system can readily offer numerous enhancements.

21.2 DSP HARDWARE

As the previous sections have suggested, DSP algorithms require fast data transfers and, what is most important, a fast multiply/accumulate capability. For example, if a 30-tap FIR filter is applied to a signal sampled at a 100-kHz rate, the system must be able to perform 3 million multiply/accumulates per second (Equation 21.5). As we will see, DSP processors can easily meet this challenge. In contrast, it would take nearly twenty of one of the fastest microprocessors available (the 12.5-MHz Motorola 68000, with its 5.5-microsecond multiply time) in parallel to handle this task.

Although the number of DSP architectures that can be designed is quasi-infinite, they have in common the need for several functional elements—program sequencing, address generation, and number crunching—as well as memories that store program instructions, data, and coefficients.

Program instructions provide the controls that, on a cycle-by-cycle basis, govern the operation of all circuits in the DSP processor. By specifying the addresses to the program memory, the *sequencer* controls the system's instruction flow. It is also responsible for branching, subroutine jumps, interrupt handling, and overall system control.

A digital signal processor's *address-generator* logic determines read/write locations for coefficient and data memory. A fast, flexible addressing element is needed in DSP, since many algorithms have intricate addressing structures and require rapid data transfers to and from the number crunchers. System capabilities are seriously compromised by inadequate addressing logic.

A DSP system's *number crunchers*—multipliers, arithmetic/logic units (ALUs), and barrel shifters—perform the data computation for DSP algorithms. These elements are characterized by three attributes—precision, throughput, and cost. Precision is established by the size of arithmetic word used to handle the digitized signal, ranging from a 4-bit fixed-point number (one part in 16 resolution) to double-precision floating point (dynamic range of better than one part in 10^{300}). Throughput is a function of speed, internal

architecture, and port structure. The arithmetic characteristics of these devices must match the system's processing demands.

Two principal approaches are used today: microcoded systems—using building-block ICs—and single-chip signal processors. The advantages of a microcoded system are generally speed and flexibility. Single-chip processors, in contrast, generally benefit from compactness and ease of design.

21.2.1 SINGLE-CHIP PROCESSORS

Single-chip processors have program sequencing, addressing, and arithmetic logic on a single device. In addition, devices available today generally include a modest amount of program and data memory on the device. In the present state of commercial semiconductor technology, it isn't feasible to fabricate high-performance DSP functions and a large amount of memory on a single monolithic device. Consequently, compromises must be made that force today's single-chip processors to be optimized for specific applications (e.g., telecommunications, imaging). They offer adequate performance in such narrowly defined roles—but are precluded from general-purpose high-performance digital signal-processing.

21.2.2 MICROCODED SYSTEMS

A microcoded system employs building-block ICs to construct the digital signal-processor (see Figure 21.12). This gives the designer increased control

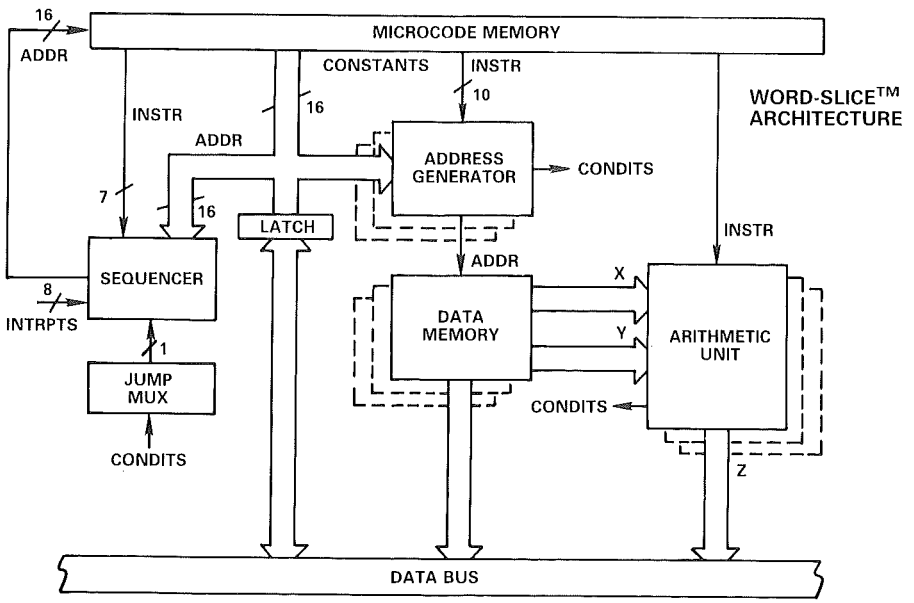

Figure 21.12. Building blocks of a microcoded DSP system. Word Slice™ is a trade mark of Analog Devices, Inc.

over the system's architecture, with sufficient latitude to meet very high throughput specifications. In the past, such designs traditionally were pieced together with numerous power-hungry bit-slice devices, which tended to increase the system's cost and complexity. A second-generation family of microcoded devices, which provide a comprehensive set of building blocks that enable high-performance, relatively compact digital signal processors to be constructed, is shown in figures 21.13 through 21.20.

The power of a microcoded system lies in the fact that, during each clock cycle, each component can execute an instruction, allowing the system to attain high throughput. The operation of each component is governed by the system's microcode memory, which contains the microcode controls (0's and 1's). On a cycle-by-cycle basis, these controls are fed to each component, telling it what instruction to execute. The microcode memory's width is proportional to the number of components in the system; its depth is proportional to the length of the overall microcode program (or microprogram).

The brain of a microcoded system is its program sequencer, which steps the system through the microcode memory. A representative sequencer is the device shown in Figure 21.13, a high-speed, 16-bit microprogram sequencer op-

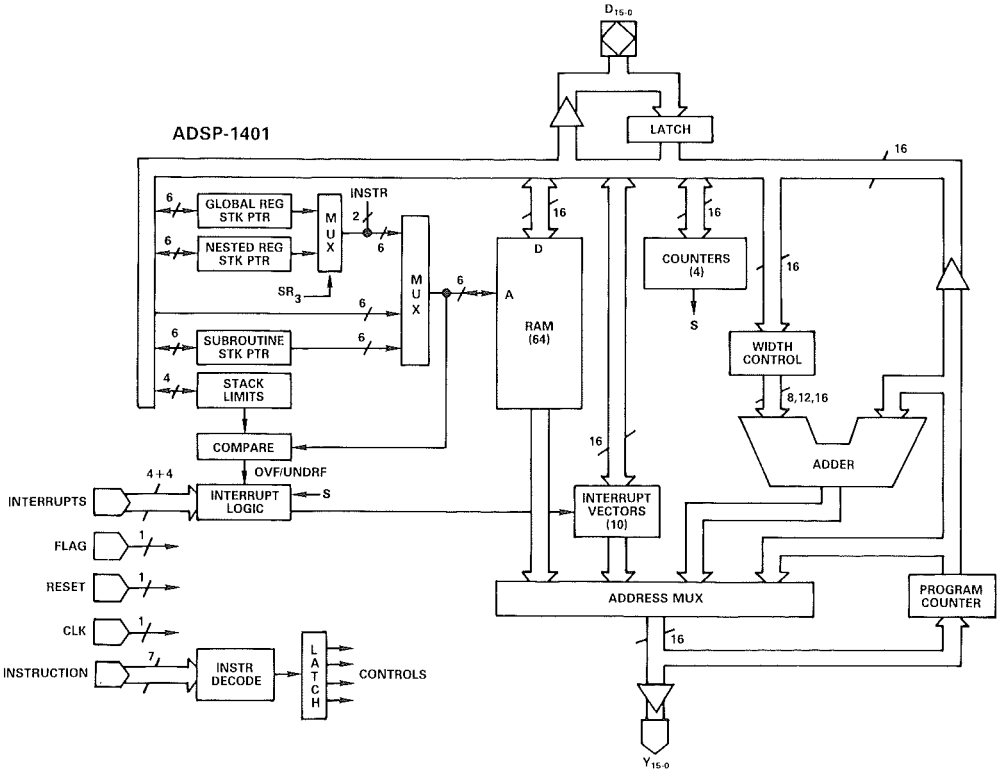

Figure 21.13. Program sequencer for use in microcoded DSP designs.

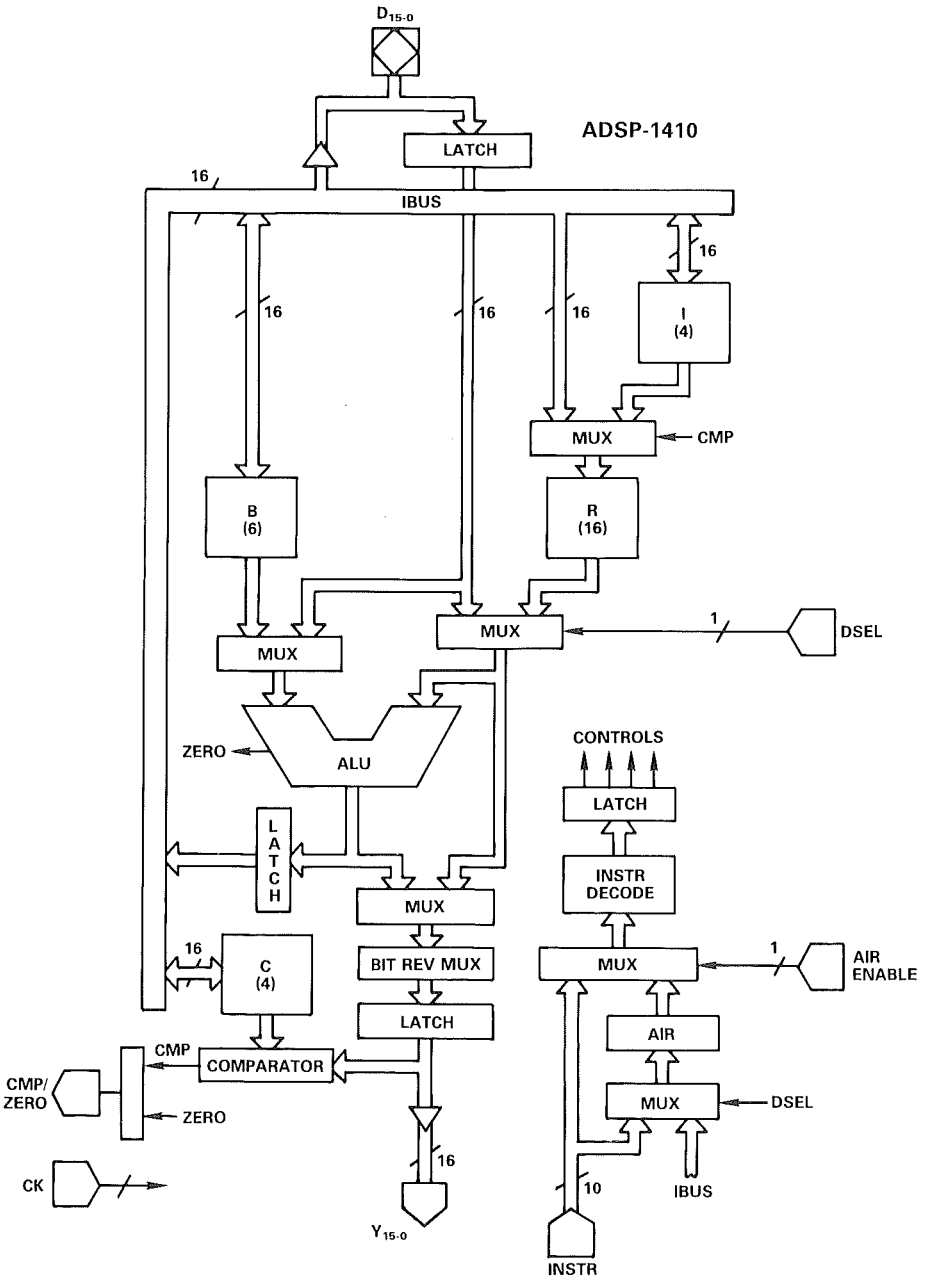

Figure 21.14. Address generator for microcoded DSP designs.

timized for DSP. This device usually performs an elementary function—it stores the current microcode address (instruction number) in the program counter, increments it, and outputs it. This activity corresponds to simply stepping through the instructions of a program. Obviously, though, limiting

any computer to executing straightline code is unduly restrictive. For this reason, microprogram sequencers must also have the flexibility to handle subroutine jumps, branches, interrupts, and indirect jump addresses. For example, the 48-pin ADSP-1401 determines the next program memory address from i) a simple increment of the current address; ii) an absolute or relative jump; iii) a jump address from the internal RAM; iv) a subroutine return; or, v) an internal interrupt vector.

For its addressing requirements, a microcoded system can splice together several 4-bit or 8-bit arithmetic and logic units (ALUs). Alternatively, it can use devices such as the high-speed, general-purpose data address generator, shown in Figure 21.14. It performs all addressing tasks required by DSP routines in a single cycle; specifically, it

Outputs an address pointer to data memory;
Modifies the pointer by an offset value to determine the next memory read/write address;
Compares the output pointer to a pre-set value, and, if equal, re-initializes the pointer with a value stored on-chip.

One address generator may be adequate for an application; however, if extremely high throughput is required, several address generators may be needed in the system.

The heart of a DSP microcoded system is its arithmetic units, the devices that perform the DSP number crunching. Until recently, designers of systems demanding high throughput had relatively few options; but numerous recently available chips increase the designer's choices, imparting the flexibility to optimize DSP hardware for the application's specific processing requirements.

21.2.3 FIXED-POINT OPTIONS

Initially introduced in 1976 by TRW, and now available from many manufacturers—Analog Devices among them—industry-standard fixed-point multipliers provide systems with fast multiplication capability. These VLSI devices dedicate a large amount of silicon to a parallel array multiplier that can deliver the product of two numbers at 10-MHz rates. To maximize throughput, these devices have dedicated ports for multiplier inputs and outputs. As a result, device I/O rates match multiply speed. The block diagram of a 16×16-bit example is shown in Figure 21.15.

Since many DSP algorithms require a chain of products to be continuously summed, high-speed multiplier/accumulators (MACs) are also available (Figure 21.16). These devices not only multiply two numbers together rapidly, they can optionally add each product to a value stored in an on-chip accumulator. In this fashion, a single device can perform the high-speed MACs required in DSP.

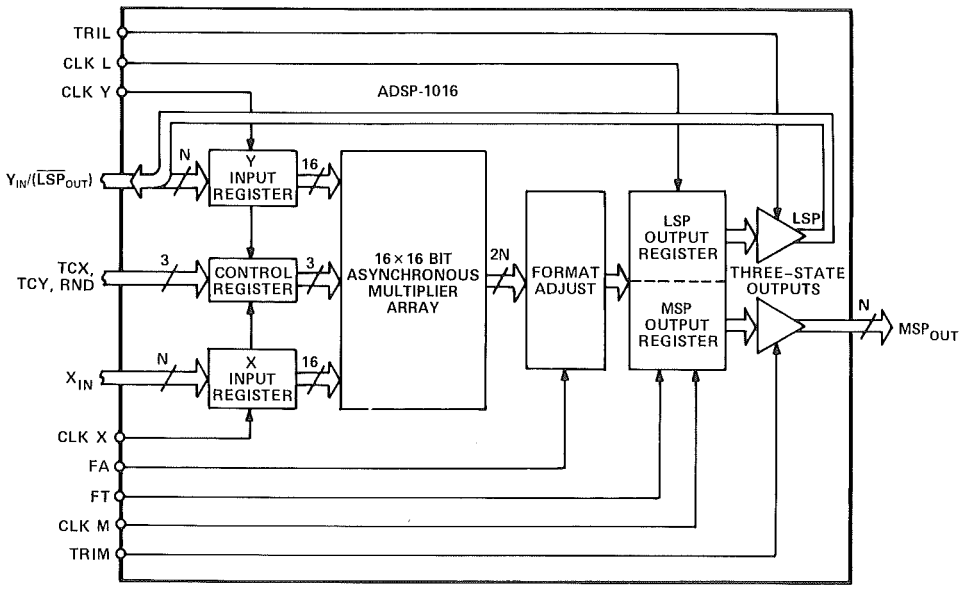

Figure 21.15. Typical 16 × 16 multiplier for DSP. Note that full 32-bit output product is available.

Several different versions of these industry-standard multipliers and MACs are available. Depending on the arithmetic precision required in an application, the designer can choose 8-bit, 12-bit, or 16-bit implementations. Number formats they will support include unsigned magnitude or twos com-

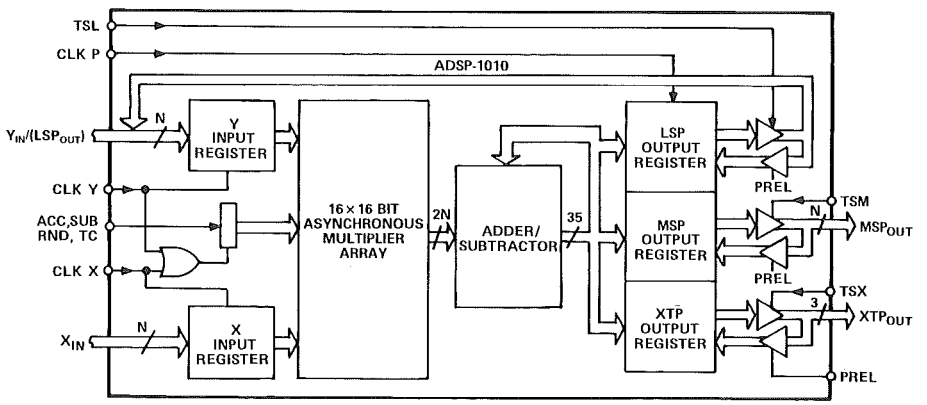

Figure 21.16. Typical 16 × 16 multiplier/accumulator (MAC).

plement. Tradeoffs must also be made among two key attributes of these devices—speed and power—although recent advances in high-speed CMOS are allowing the same device to dominate in both respects. Table 21.1 details some of the available high-speed, low-power CMOS multipliers and MACs.

Designation	Device Type	Resolution (Bits)	Cycle Time Commercial (ns)	Cycle Time Military (ns)
ADSP-1016A	MULTIPLIER	16 × 16	75	90
ADSP-1010A	MAC	16 × 16	95	110
ADSP-1012A	MULTIPLIER	12 × 12	65	75
ADSP-1009A	MAC	12 × 12	80	90
ADSP-1080A	MULTIPLIER	8 × 8	50	60
ADSP-1081	UNSIGNED MULT.	8 × 8	100	115
ADSP-1008A	MAC	8 × 8	55	85

Table 21.1 CMOS Multipliers and Multiplier/Accumulators

There are also devices that differ from the industry-standard fixed point multipliers and MACs to provide special application conveniences. For example, for applications where cost and board space must be reduced, a single-port MAC (Figure 21.17a) can be used with little sacrifice in throughput. Both inputs, and the output, communicate with the 16-bit computer bus via its single port, but a 6-bit instruction set makes possible a substantial amount of time-saving internal processing. Extra fixed-point precision can be obtained using devices such as a 24 × 24-bit multiplier (b). Extremely high speed can be attained with high-speed multipliers employing *pipelining* (see Glossary); for example, the Analog Devices ADSP-1102 and 1103, 16-bit fixed point devices can, in pipelined mode, deliver products at a 30-MHz rate.

Figure 21.18 shows the architecture of a 16-bit "enhanced" MAC (EMAC), that provides high-speed DSP capabilities, while eliminating considerable external circuitry. A monolithic integrated circuit, packaged in a one-square-inch pin-grid array, the device features a 16 × 16-bit array multiplier, two addressable input registers on each input port, two 40-bit-wide accumulators, and an internal feedback path, which simplifies polynomial expansions.

21.2.4 FLOATING-POINT OPTIONS

Floating-point arithmetic differs from fixed point in that each number has its own exponent, as well as a mantissa. As a result, systems using floating-point arithmetic handle very wide dynamic ranges. Floating point is valuable to a system any time a weak signal must be detected in the presence of noise. Also, if a large number of arithmetic operations are being performed on the data, floating-point largely eliminates the distorting effects of round-off noise. Applications demanding floating point include instrumentation, graphics and image processing, engineering workstations, and general-purpose array processors.

In the past, floating-point capability required considerable LSI, MSI, and SSI logic, configured on one or more circuit boards. Recently, however, monolithic floating-point implementations have become available. Initial offerings included limited-precision multipliers and ALUs—22-bit devices.

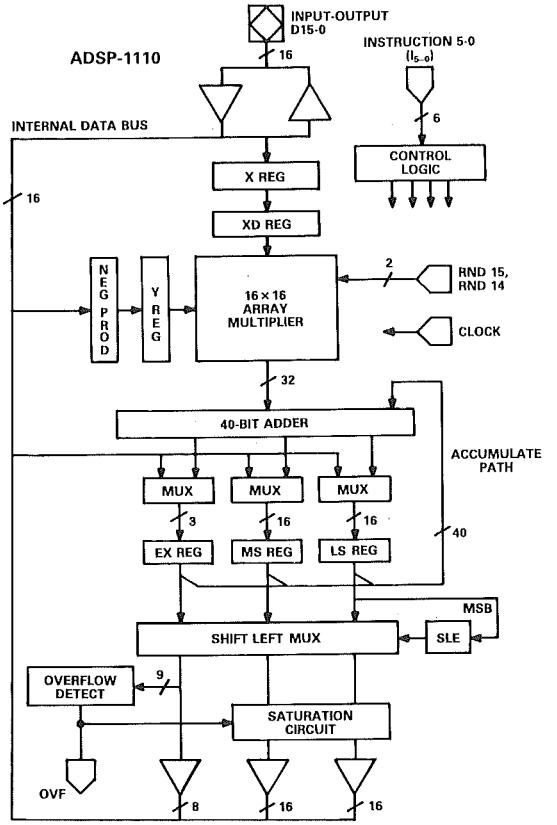

a. Single-port 16 × 16 multiplier/accumulator architecture. 40-bit adder provides extra accumulator capacity, reducing tendency to overflow.

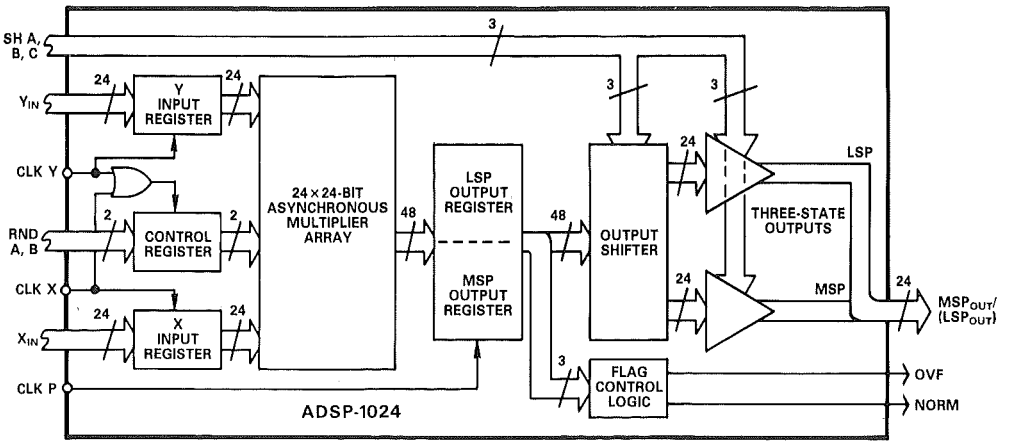

b. 24-bit multiplier with full 48-bit output in two bytes.

Figure 21.17. Special-purpose multipliers.

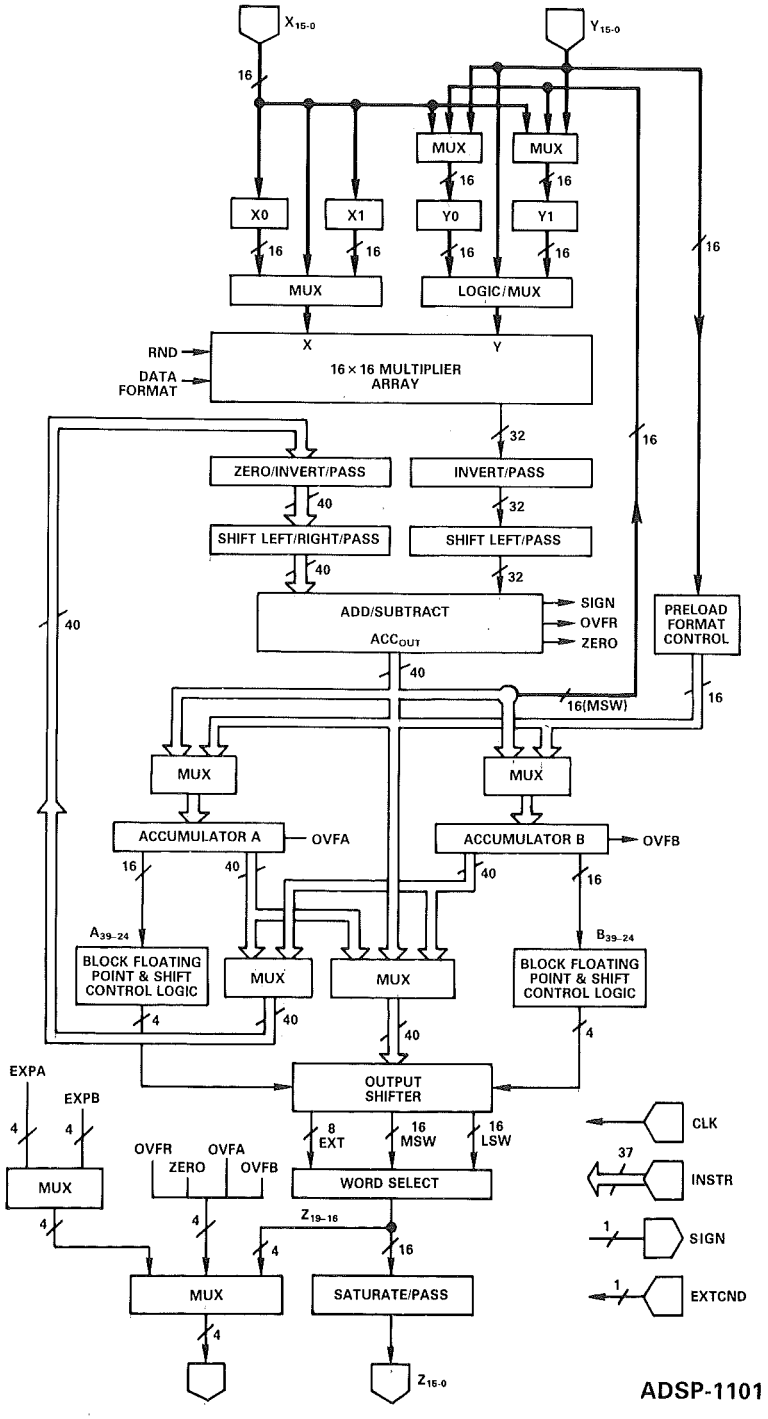

Figure 21.18. Enhanced multiplier/accumulator (EMAC).

The IEEE has issued a 32-bit floating point standard, and monolithic DSP ICs now exist that implement its specifications. Figure 21.19a shows an IEEE-compatible floating-point multiplier, capable of performing single- and double-precision floating-point operations, as well as 32-bit fixed point multiplications. A companion device (b) is an ALU, which handles single- and double-precision IEEE floating-point, as well as 32-bit fixed-point operations.

21.2.5 REPRESENTATIVE SYSTEMS

Figure 21.20 shows several alternative DSP architectures, each with its own characteristic throughput, precision, and cost. These compact, low-power CMOS-based systems can be microprogrammed to carry out any DSP algorithm. A key benchmark in DSP is the 1,024-point complex FFT. This al-

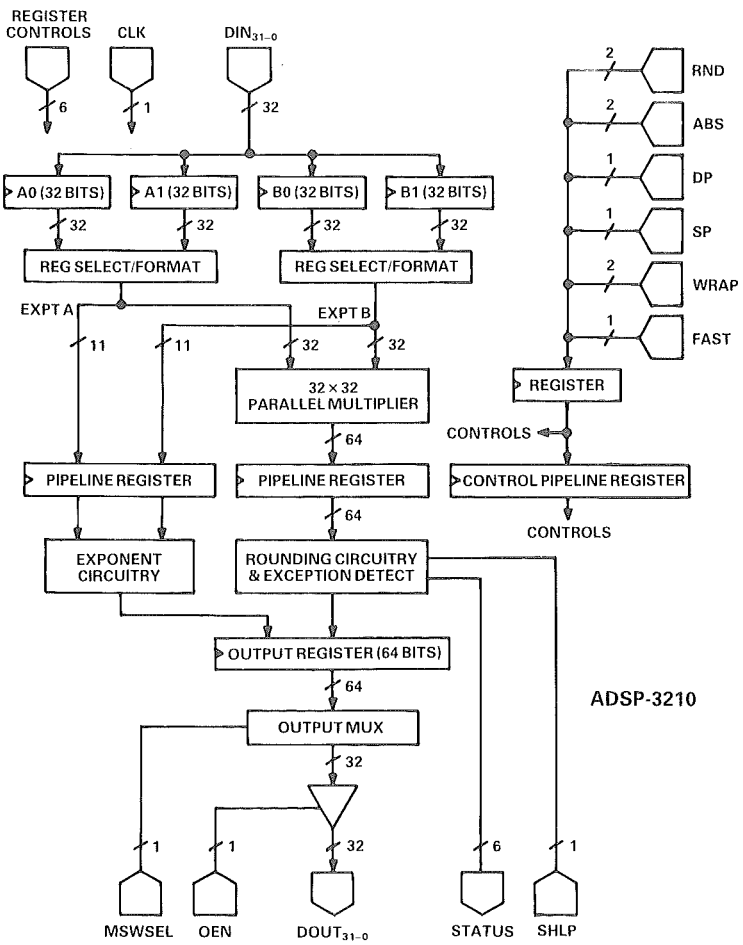

a. Floating-point multiplier.

Figure 21.19. Floating-point DSP units.

gorithm requires highly efficient program-sequencing, address-generating, and number-crunching capabilities. All of Figure 21.20's architectures will execute a 1,024-point complex FFT swiftly. For instance, (a)'s EMAC-based architecture (see Figure 21.18) performs this algorithm in 3 milliseconds. The floating-point architecture of (b) delivers a 1,024-point FFT in 4 milliseconds. Finally, the architecture of (c), featuring the 16×16-bit ADSP-1110A single-port MAC, a program sequencer (PS), and an address generator (AG), can execute a 1K complex FFT in 10 milliseconds.

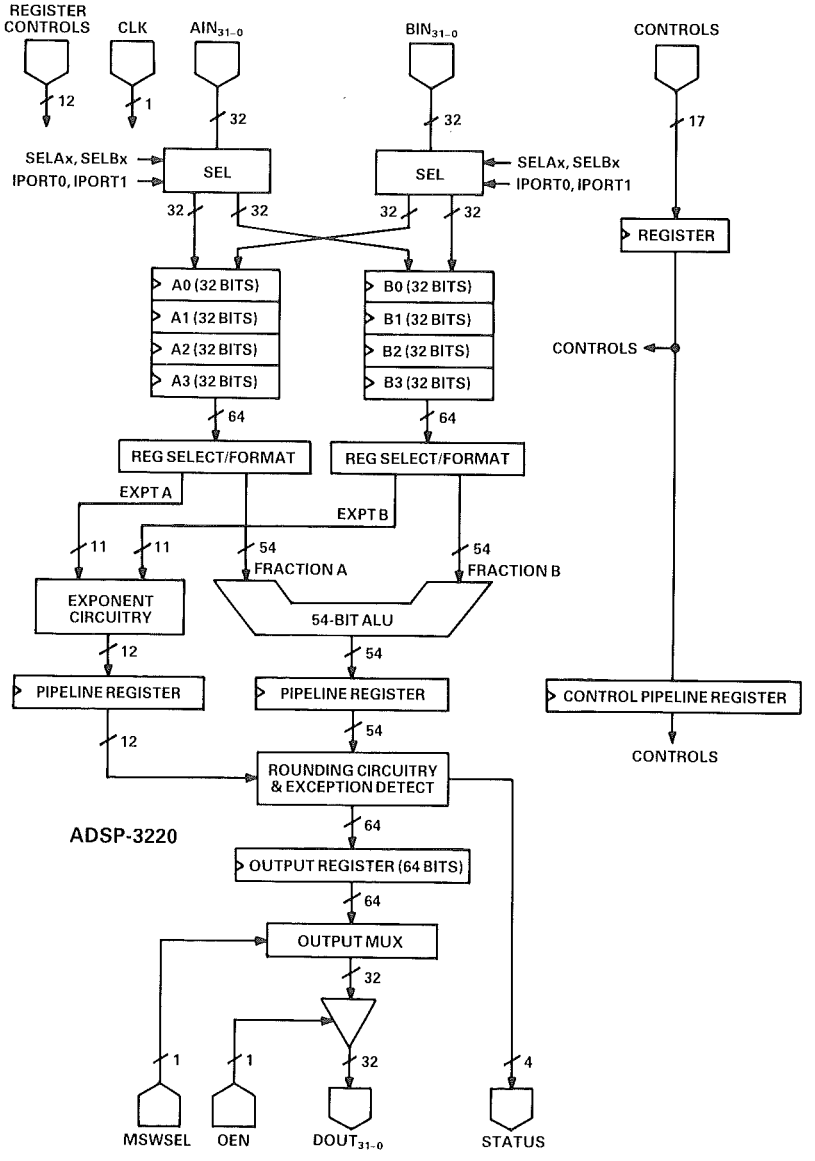

b. Floating-point ALU.

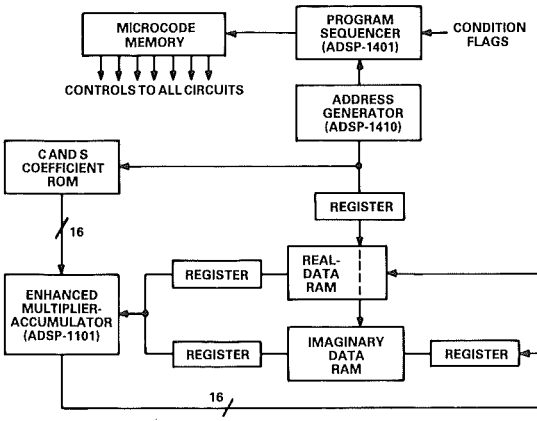

a. Enhanced-multiplier/accumulator (EMAC) based system.

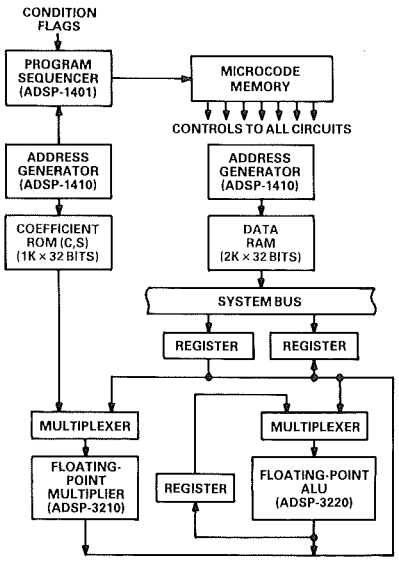

b. Floating-point processor.

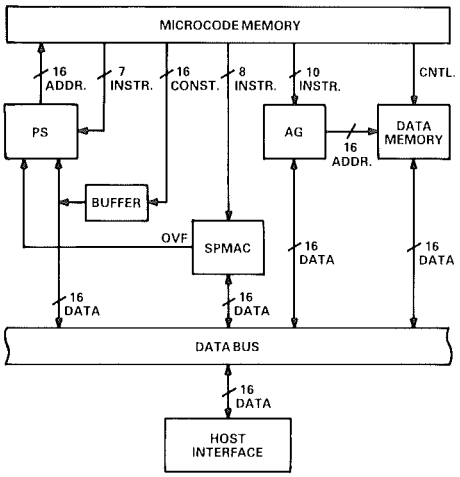

c. 16-bit fixed-point FFT processor.

Figure 21.20. Processor architectures.

21.3 DSP APPLICATIONS

DSP began as a specialized technology used in military applications. With government funding, a high-speed integrated-circuit array multiplier was developed and first offered commercially in 1976. This multiplier formed the heart of high-performance radar, sonar, and missile-control systems. Over time, however, the use of DSP has spread from specialized military niches into a broader set of industrial and commercial markets, as the table below confirms. A few important uses—modems, studio recording, ultrasound imaging, and vibration analyzers—are described in detail below, primarily to illustrate how DSP's advantages benefit the application.

PRINCIPAL DSP MARKETS

Instrumentation:	Spectrum analyzers, vibration analyzers, mass spectroscopy, chromatography
Audio:	Studio recording, music synthesis, speech recognition
Communications:	Modems, transmultiplexers, vocoders, satellite transmission, repeaters, voice storage and forwarding systems
Computers & Computer Peripherals:	Arithmetic acceleration, servo controls for disk head positioning, array processors, engineering workstations
Imaging:	Medical, satellite, seismic, bandwidth compression, digital television, machine vision
Graphics:	CAD/CAM, computer animation and special effects, solids modelling, video games, flight simulators
Defense Electronics:	Radar, sonar, missile/torpedo control, secure communications
Control:	Robotics, servo links, skid-eliminator adaptive control, engine control.

21.3.1 MODEMS

A tremendous amount of information is transmitted today over analog communication links, such as telephone lines. With the growing role of computer-based systems, this information is increasingly digital in nature (for example, digital data and digitized voice transmission). The challenge of transmitting digital data over analog links at high speeds, and reconstructing the received data with high noise immunity, thereby reducing communication costs, is met by a modulator-demodulator (modem).

In transmitting digital data over analog communication lines, a digital bit pattern is represented by modulating the phase, frequency, and/or amplitude of an analog signal. Figure 21.21 shows a simple scheme, which involves changing the frequency of the signal to denote a "0" or "1"; this frequency-shift keying (FSK) method can send 2,000 bits/second over a telephone line. A more sophisticated encoding method, quadrature phase-shift keying (QPSK), modifies the phase of the signal and is capable of transmitting data at four times the rate of simpler methods. When a modem is sending information, it encodes the digital data into the corresponding analog waveform; in receiving mode, it decodes the waveform and determines the bit pattern that was transmitted. This latter mode is the more difficult to implement.

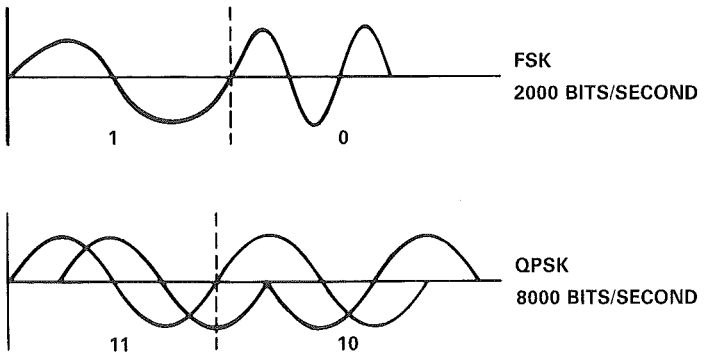

Figure 21.21. Frequency-shift keying (FSK) and quadrature phase-shift keying (QPSK).

If the transmission medium were noiseless, a modem's tasks would be limited to simple encoding and decoding—a relatively straightforward exercise. However, a phone line is a noisy transmission medium, corrupting the analog waveform. The more sophisticated the encoding scheme, the more disastrous the effects of noise and channel distortion. Therefore, a modem must effectively compensate for, or equalize, this channel distortion. To this end, high-speed modems (4,800 bits/second and above) turn to DSP for high-performance data recovery, using digital FIR filters.

Once the system channel distortion is filtered out, the modem must decide what bit pattern was originally sent. The compensation, or filtering, will rarely be perfect; that is, the amplitude, phase, and/or frequency of the filtered waveform will not generally correspond precisely to what was originally encoded. A least mean-squares (LMS) criterion (requiring fast multiplication capability) can be used to best estimate what the transmitted bit pattern is. A modem architecture for deriving the digital bit pattern from the received analog signal (and predicting the correct bit pattern in the presence of noise) can be seen in Figure 21.22. A modem using this relatively simple estimation method has a non-zero probability of bit error; more sophisticated DSP in a modem decreases the likelihood of such errors.

An additional complexity of telephone-line transmission is that its distortion properties change over time. Therefore, a modem's digital filter must be able to adapt to changes in the environment. This need to respond to a changing

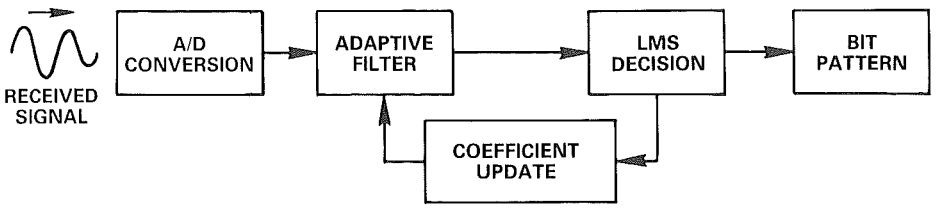

Figure 21.22. Least-mean squares modem processing architecture.

environment underscores another advantage of DSP—a digital filter's characteristics can be modified simply by changing its coefficients. Coefficient updating in a modem is determined by the observed drift of a property of the distortion in the system.

Aided by DSP, then, a modem can make it possible for high-speed data transmission to be implemented effectively. As Figure 21.22 illustrates, the DSP is the heart of a high-speed modem. The processing required generally can be handled by one digital multiplier, surrounded by the appropriate support devices (a program sequencer, an address generator). Alternatively, depending on the requirements of the modem, a single-chip processor may adequately handle all DSP requirements. The a/d conversion is an essential adjunct to the modem's signal processing. This structure allows the advantages of DSP to create a communications function that wouldn't have been possible with traditional analog signal-processing techniques.

21.3.2 STUDIO RECORDING

One of the most interesting applications of DSP is emerging in the audio processing performed in recording studios. This processing starts after the initial recording of voices and instruments in the studio; after a large number of steps, it ends with the recorded version that reaches the home stereo. Increasingly, DSP is being used to handle all intermediate steps.

The flow of activities in studio recording is complex and varied. Generally, multiple channels are used, with each track dedicated to one or more sources (instruments/voices). All channels need not be recorded at the same time. Each channel is subjected to extensive processing, including gain control, filtering, non-linear compression or expansion, reverberation adjustments, spectral equalization, and special-effects enhancements. The contributing channels are then mixed together to obtain a final arrangement with the desired overall effect.

Traditionally, channel processing and mixing were implemented entirely in the analog realm—with numerous disadvantages. Each channel's information—stored as an analog signal on magnetic tape—degrades as the cutting, splicing, and re-recording process progresses, undermining the benefits of the processing. The limited performance range available with analog processing sets a ceiling on the signal enhancement that can be attained. Also, analog circuitry can only handle one channel at a time; multi-channel mixers are expensive and difficult to control. Finally, if analog processing hardware is used, overall mixing flexibility can be achieved only through hardware modifications. In practice, this means that the mixing process loses its ability to creatively explore special effects.

Increasingly, audio processing is relying on digital techniques to improve audio quality. The first step in this transition was digital recording, which be-

came prevalent about five years ago. Audio signals are first converted to digital form before being stored on magnetic tape. Digital recording eliminates several sources of degradation that hamper analog recordings, including the effects of non-linearities and additive noise in the magnetic materials used for recording, and wow and flutter in the tape playback mechanism.

In studio mixing applications, however, digital recording does not eliminate all complications. In the mixing and enhancement process, information is passed from one tape to another—requiring D/A and A/D conversion processes, a source of noise. These conversions are no longer necessary if all processing and mixing are handled with DSP techniques.

In DSP-based studio recording systems (see Figure 21.23), signals are converted to digital as early as possible. In fact, some implementations place a remotely controlled amplifier/converter at the recording microphone. After conversion, the audio processing is handled digitally, with high performance and flexibility. Gain factors are handled with digital multiplication. Filtering and equalization can be handled with an IIR filter that replicates the performance of standard analog filters. Alternatively, digital FIR filters can provide high-performance linear-phase filters or complex comb filters. Dynamic-range control is easily included in the system by using a multiplier for non-linear compression/expansion computations.

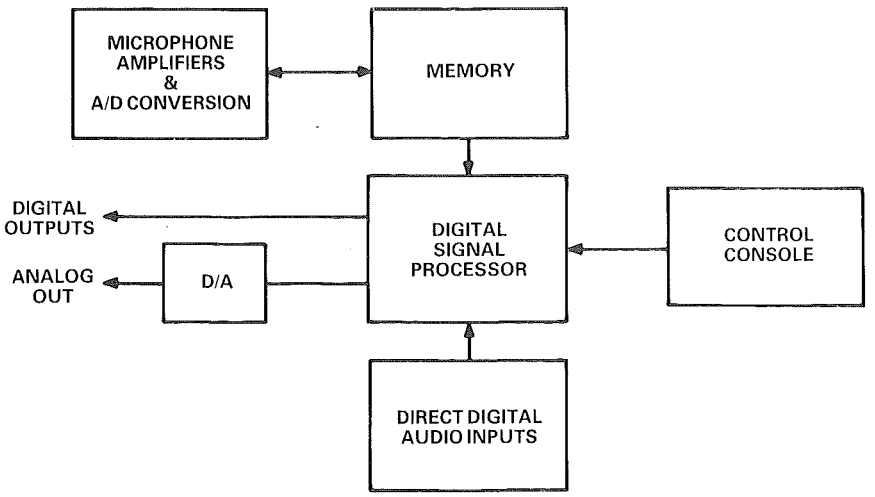

Figure 21.23. Block diagram for processing in studio recording.

The traditional mixing process is also easily implemented in a DSP-based system. Digital channels to be mixed are simply added together. Relative time delay lags can be easily introduced into the channel flows, allowing phase coherence to be explored without adding expensive delay lines to the system. An additional advantage is that the channel interconnections—which have to be hardwired in an analog processor—can be easily reconfigured in a DSP system.

In addition to improving on traditional operations, a DSP studio recording system opens up numerous new options. Unusual special effects are readily included in the system. Reverberation effects can be modeled, simulated, and integrated into the final recording. An FFT routine's spectral analysis of the signal forms the basis for frequency-domain filters that provide optimal equalization. Overall system flexibility allows the entire mixing system to be dynamically configured—processing steps can be re-ordered, mix groups and subgroups re-specified, and effects such as fading, equalization, and compression/expansion included at any juncture.

In practice, to perform the many required operations in timely fashion without introducing noise or distortion, studio recording systems face demanding processing tasks. Consequently, such designs rely on microcoded systems with numerous array multipliers. Such multipliers carry out the scaling, filtering, and FFTs required in the DSP block of Figure 21.23's block diagram.

Studio recording, then, follows the pattern of other applications using DSP. DSP techniques offers increased precision for processing steps traditionally performed with analog circuits. Of equal importance, DSP's flexibility paves the way for many new and creative processing steps. As in other areas, the DSP is shifting the role of converters; accurate ADCs and DACs are used in the system, but as close to the real-world interface as possible. The signal processing is conducted in the digital realm.

21.3.3 ULTRASOUND IMAGING

An important non-destructive imaging technology uses acoustic, or ultrasound, waves to investigate the interior of the human body. Like other medical imaging technologies, ultrasound draws heavily on the advantages of DSP for constructing and displaying images of internal organs.

In an ultrasound imaging system (see Figure 21.24), acoustic waves of a certain frequency are sent into the body. A phased array detector digitizes the reflected waves. After intensive number-crunching of the relative amplitudes and phase delays of returned signals, the shape of an internal organ can be inferred.

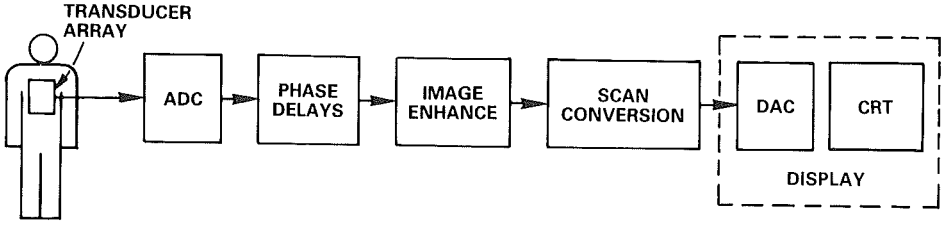

Figure 21.24. Ultrasound medical imaging.

High-speed arithmetic hardware is also used in an ultrasound imaging system to prepare the inferred image for display. The imaging data is originally expressed in polar coordinates—namely, distance from the imaging detector

and angle relative to the detector. In order to display the image on a screen, a transformation from polar to rectangular coordinates must be made. This transformation—called scan conversion—is handled by fast DSP multipliers.

Since the human body is a poor transmission medium for acoustic waves, the image is degraded substantially by noise. To improve the image's clarity, an ultrasound imager performs three-dimensional FIR filtering. A one-dimensional temporal filter improves the quality of a displayed image; each pixel's displayed value is determined by a weighted average over the previous few values in time. Further image enhancement is obtained by using a two-dimensional neighborhood filter; each pixel's display intensity is based on a weighted average over its nearest neighbors in the frame. One result of a neighborhood filter is to soften edges, reducing the "staircase" effect when a diagonal line is presented on a raster-type display.

As in other applications using DSP, the presence of high-performance number-crunching hardware allows the system to perform certain tasks that would otherwise be unachievable. In ultrasound imagers, for example, it is possible to draw certain inferences about the structure of arteries. Acoustic waves reflected by blood cells will experience a Doppler shift proportional to the velocity of these cells. If an FFT is performed on the returned wave, this Doppler shift (and therefore velocity) can be computed. Then, by analyzing the velocity of blood cells over the cross-section of an artery, information is gained about the artery's resistance characteristic—specifically, whether clotting is present.

21.3.4 VIBRATION ANALYZERS

Vibration analyzers make particularly extensive use of DSP. An example of the use of a vibration analyzer is in monitoring an expensive and hard to replace turbine in an electric power-generating station. As the turbine rotates, a bearing emits a characteristic spectrum of high- and low-pitched sounds. This frequency distribution contains information about the structure and health of the turbine. The frequency spectrum will shift markedly as the bearing deteriorates; if detected in time, this information allows cost-effective preventive maintenance to be performed and catastrophic failures to be averted.

The block diagram of Figure 21.25 represents the basic functions of a DSP-based vibration analyzer. First, the system under investigation must be stimulated with a waveform. To obtain a precise waveform, DSP techniques are implemented by fast digital logic. For instance, a sine wave can be generated by using a lookup table and multiplication-intensive interpolation schemes. Then, the sine wave can be shifted to any other location in the frequency spectrum by modulating it with the output of a digital oscillator.

Stimulated with a waveform, the system will respond in a manner analyzed by the vibration analyzer. The waveform, or signal, reflecting the system's

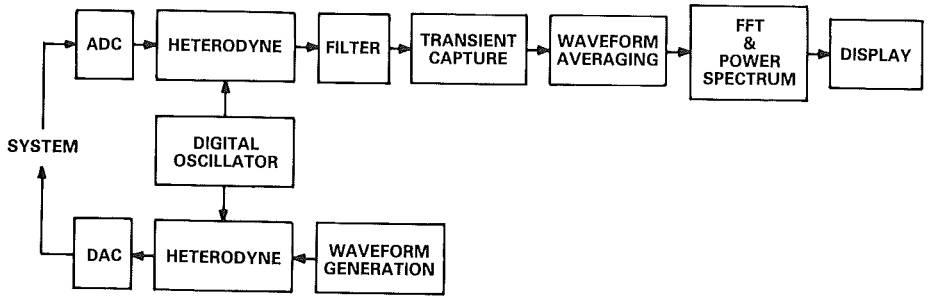

Figure 21.25. Vibration-analyzer architecture.

response is passed through a fast, high-resolution analog-to-digital converter and stored in a buffer memory. The vibration analyzer can then perform a range of processing tasks on this digitized signal in order to draw inferences about the structure of the system under test.

A key task is to filter the incoming signal. Fast DSP hardware allows all noise at undesired frequencies to be suppressed by a high-performance digital filter. Then, an FFT is performed, providing information about the spectral content of the received signal. This spectral information allows the transfer function to be estimated—information related to be basic structure of the system being monitored.

In some instances, it may be desirable to "zoom in" on a small portion of the frequency range; one technique that facilitates such zooming is to modulate the signal down to baseband, using the same digital oscillator used in stimulating the system. Then, an FFT, performed over a narrow bandwidth, provides very precise spectral resolution over a narrow frequency range.

Waveforms can be averaged either prior to or after performing the FFT; the choice may help toward improving the system's signal-to-noise ratios. With fast digital hardware in the system, sophisticated averaging schemes can be used, further enhancing system performance. Also, in displaying spectral information, certain calculations that are required can be performed, for example, power-spectrum computations. Polynomial-expansion approximations, which require fast multiply capability, can also be used by systems equipped with DSP hardware.

Once the signal has been digitized, it is a straightforward matter for the system to carry out other important computations. For example, auto- and cross-correlations can be readily calculated. Such information is useful, among others, in determining how closely the observed response matches one stored in memory. The memory-based response might correspond to a healthy system; if the correlation were to fall below a certain threshold, a failure might be likely and should be flagged.

Perhaps the biggest advantage accruing to a DSP-based vibration analyzer is its flexibility. Not only can a particular signal-processing scheme be carried

out with great accuracy, but simple changes in software allow the same hardware to carry out a very different set of tasks. In this sense, vibration analyzers are representative of many types of equipment that now use DSP, not only to supersede the performance of more-traditional implementations, but to cost-effectively add new, otherwise unrealizable, capabilities to the system.

21.4 GLOSSARY OF DSP TERMS

Accumulator—an arithmetic element that adds together, or accumulates, a sequence of inputs. A DSP multiplier with an accumulator on-chip is called a multiplier/accumulator (MAC).

Algorithm—A DSP algorithm, such as the fast Fourier transform, or a finite impulse-response filter, is a structured set of instructions, and/or operations, tailored to accomplish a signal-processing task. Each algorithm has a well-defined structure; however, variations in algorithm parameters, such as the number of input points or taps, allow the same basic algorithm to perform different functions.

ALU—An arithmetic and logic unit, which performs additions, subtractions, or logical operations (e.g., AND, OR, XOR) on operand pairs.

Attenuation—The damping-out, or suppression, of signal content. Filters will attenuate the frequency content of a signal that lies in the filter's stopband.

Barrel Shifter—A device that accepts a digital number as its input and—as a function of the controls—shifts the number up or down, or rotates the word as though it were placed on a barrel. A barrel shifter is used in a system for many tasks, including scaling and normalization.

Biquad—A particularly simple recursive, or infinite impulse-response (IIR), digital filter form, often used as a building block for constructing more complicated recursive filters. A biquadratic, or biquad, section uses the three most recent input points and the two most recent output, or feedback, values to compute each output point.

Block Floating Point—A compromise between fixed-point and floating-point arithmetic. Data grouped in "blocks" is assumed to be normalized with a common exponent (but, not being attached to the data words, the exponent need not be explicitly processed with the data). In essence, the process is carried out in fixed point, with its inherent speed advantage.

Convolution—In discrete computations, a mathematical operation, defined as the summation, or integral, of a product of two functions over a range of differences in the independent variable. In the time domain, one function is the impulse response, as a set of coefficients, h(i), over N time intervals; the other is the input, f(n − i), as a function of the differences between the time at the instant at which the function is being evaluated, n, and the input at earlier

instants, determined by the variable delay, i, from 0 to N. See equation 21.4b. In DSP, the convolution of an input signal, x, with the coefficients, h, results in the filtering of the input signal.

Correlation—a mathematical operation that indicates the degree to which two signals overlap. A high positive correlation reflects two signals that closely track each other. A negative correlation indicates that the two signals are closely related, but out of phase by roughly 180 degrees. If the correlation is close to zero, the two signals are unrelated.

Digital Signal Processing—DSP is a technology for high-performance signal processing that combines algorithms and fast number-crunching digital hardware.

Discrete Fourier Transform—the discrete Fourier transform (DFT) is a DSP algorithm used to determine the Fourier coefficient corresponding to a particular frequency.

FFT—An n-point fast Fourier transform (FFT) is computationally equivalent to performing n DFTs but, by taking advantage of computational symmetries and redundancies, can reduce the computational burden by several orders of magnitude.

FIR Filter—A finite impulse-response (FIR) filter is a commonly used type of digital filter. Digitized samples of the signal serve as inputs; each filtered output is computed from a weighted average of a finite number of previous inputs.

Fixed-Point Arithmetic—Each number is represented in a fixed arithmetic field of n bits, allowing integers in the range, 0 to $2^n - 1$, to be represented.

Floating-Point Arithmetic—Each number consists of a mantissa and an exponent, allowing wide dynamic range to be accommodated in the numbering system.

IIR Filter—An infinite impulse-response (IIR) filter is a commonly used type of digital filter. This recursive structure accepts as inputs digitized samples of the signal; each output point is computed on the basis of a weighted average of past output—or feedback—terms as well as past input values. An IIR filter is more efficient than its FIR counterpart, but poses more challenging design issues.

MAC—Multiplier/accumulator; see Accumulator.

Microcode—a set of instruction control signals stored in a program memory that govern the cycle-by-cycle operation of the various devices in a building-block architecture.

Passband—the frequency range over which a filter passes, to within some tolerance, the incoming signal content.

Pipeline—An architectural structure that allows two or more operations to be

carried out simultaneously, like the stages of an assembly line. While each basic operation requires several cycles to complete, a later stage of one operation is simultaneously with an earlier stage of another operation. This structure allows the effective throughput rate for each operation to be substantially increased.

Rolloff—a measure of filter performance defined as the rate-of-change of the filter's amplitude response with respect to frequency over a transition band.

Stopband—the frequency range over which a filter attenuates, to within some tolerance, the incoming signal content.