

PART III

ANALOG-DIGITAL CONVERTERS FOR SPECIAL APPLICATIONS

Chapter Thirteen

Video Converters

Video-speed a/d and d/a converters, with conversion rates in excess of one million words per second, present a unique challenge to the designer and user. Converter architecture, performance specifications, application rules, and testing procedures all differ significantly from those for lower-speed devices.

In the case of d/a converters, for instance, the accuracies at low speeds depend upon the design of the resistor network, the stability of the reference voltage, and nonlinearities within the converter's op amp. But as speed and resolution increase, full-scale settling time and the magnitude and duration of "glitches" take on overriding importance.

In the case of a/d converters, as conversion rates pass through the 1 MHz mark, numerous application-dependent performance parameters begin to deteriorate. Spurious in-band harmonic products—created by missing codes, differential nonlinearities, and other nonlinear frequency-dependent effects—may cause signal-to-noise levels and total harmonic distortion to degrade rapidly in improperly designed applications.

This chapter is intended to provide appreciation for and guidance in the design, specification, testing, and application of high-speed data converters.

13.1 APPLICATIONS OVERVIEW

High-speed a/d and d/a converters find use in a wide range of video, digital signal-processing, radar, transient-detection, and communications applications—where the analog bandwidths require digital word rates in excess of 1 million words per second. For these applications, converters must be characterized, not only by specifications common with lower-speed converters (such

as linearity, temperature coefficients, etc.), but further, in terms that are heavily application oriented. Table 13.1 lists typical examples of specifications, that in general have not been emphasized elsewhere in this book, and the applications where they are of interest.

We will briefly illustrate the range of high-speed data-converter applications, then examine converter specifications in greater depth.

SPECIFICATION	APPLICATION
Signal-to-noise ratio (S/N)	Radar, communications spectrum analysis
AC linearity	Radar, spectrum analysis
Noise power ratio (NPR)	Communications
Two-tone intermodulation distortion	Radar spectrum analysis
Transient response	Transient analysis, radar
Overvoltage recovery	Radar
Aperture uncertainty	All
Differential phase	Television
Differential gain	Television

Table 13.1. Application-specific specifications.

13.1.1 DISPLAY SYSTEMS

Driven to ever-higher scanning speeds to obtain increased resolution without flicker, new-generation raster display systems require extremely fast information transfer. This topic has been discussed at length in Chapter 6. It will suffice to reiterate here why high-speed digital-to-analog converters are needed.

The most popular forms of display systems include raster-scan alphanumeric, raster-scan graphics, vector-scan graphics, and storage tubes. Although there are differences in the way they operate, as well as certain characteristics that might suggest one type or another for various applications, display systems for digital information generally require some form of high-speed digital-to-analog converter for beam positioning or intensity modulation.

The raster-scan system, shown in Figure 13.1, is a good example for the purpose of illustrating the role of d/a conversion. Figure 13.1 depicts a typical computer-controlled raster-scan graphic display, using a single d/a converter to control beam intensity. In a color display, three such converters are required (and are available as a triad in a single package), one for each color gun of the CRT (red, green, blue). The system consists of a MOS random-access memory buffer for storing display data in digital form, one or more memory controllers for managing the updating of the display and controlling the refresh cycle of the CRT, and a programmable microprocessor for generating display graphics and manipulating the image. The entire system operates as

an intelligent peripheral to a host computer; most of the processing associated with image and graphic display is down-loaded to the graphic subsystem.

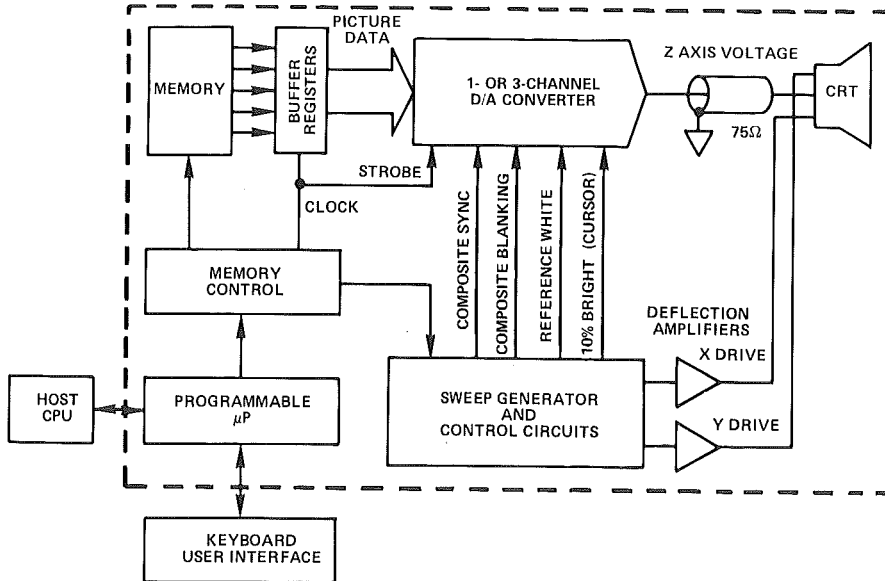


Figure 13.1. Digital display system.

The high-speed data conversion challenge becomes apparent when data rates for such displays are examined. If the picture resolution is specified as $1,024 \times 1,024$, there are 1,024 horizontal lines; each line has 1,024 independent dots, and each dot has its own programmable intensity level. Thus, there are 1,048,576 picture elements (pixels). If the picture on the CRT is to be refreshed 60 times per second, then a new pixel must appear on the screen at least every 15.9 nanoseconds, even without considering "overhead time," associated with vertical and horizontal blanking during the sweep-retrace portion of the cycle).

The d/a converter controls the Z axis of the CRT, to modulate the brightness of the raster-scan beam. For the example given above, the d/a converter must be capable of being updated somewhat faster than a pixel rate corresponding to 15.9 nanoseconds; it should be capable of settling to a new value in less than 10 nanoseconds. If the display plots a white dot on a black background, the d/a converter output must make a full scale transition between adjacent pixels. For sharp, clean raster-type displays, it is evidently of vital importance that the output d/a converter have fast settling times with imperceptible transient "glitches."

13.1.2 DIGITAL SIGNAL PROCESSING

Digital signal processing (DSP) is a broad term which encompasses a wide variety of digital techniques for the solution, manipulation, enhancement, and presentation of information in instrumentation and systems applications.

One of the more dramatic uses of DSP, which probably has the most familiarity to engineers not otherwise engaged in it, is *image processing*. Although many were unaware of it, the startling degree of detail to be found in images of the moon and Mars, displayed on home TV sets during the 1970's, and thereafter, was the result of digital techniques designed to extract a maximum amount of information for each image.

Both optical and digital techniques are used in image processing; optical techniques are simpler, but digital techniques are more flexible and have benefited by the increasing speed and availability of devices for digital data-handling technology—including analog-to-digital converters—and their decreasing cost. Because of the large amount of information typically handled in digital signal-processing, converters having both wide bandwidth and high resolution are generally required for such systems.

DSP, and especially in its applications to image processing, has been of major benefit to many fields; perhaps one of the most important is medicine. Noninvasive tools have become available that permit physicians to observe interior detail and physical functioning of body organs; these tools are helpful in both diagnosis and treatment. The growing list includes computerized axial tomography (CAT scanners), digital fluorography, and phased-array ultrasound. CAT scanners, which employ high-resolution a/d converters, have been discussed in Chapter 6. Here, we will touch on the last two techniques.

Digital Fluorography

Digital fluorography, often called digital radiography, is a computerized x-ray technology that is replacing conventional radiography using photographic film. The technique is relatively inexpensive; it can be incorporated with existing equipment; and it lowers the risks, discomfort, and costs for patients. It also provides more image information because film is limited in its ability to capture the dynamic range of x-ray signals.

X-ray signals consist of high-frequency radiation with the ability to pass through human tissue. The amount of penetration, and the resulting image character, depend on the densities of bone and tissue structure under examination. While a conventional x-ray machine captures these signals directly on a film that is sensitive to x-rays, a digital radiography machine captures x-ray signals with a TV camera, using a high-speed information-processing system to analyze and display the diagnostic image.

The major components of a digital radiography system are shown in Figure 13.2. Basically, the camera and film used in conventional x-ray examinations are replaced by the TV camera, shown in the illustration, along with the associated digital processing elements and display system.

The patient is positioned between the x-ray source and a photomultiplier fluoroscope with an image intensifier, which converts the x-rays into photons

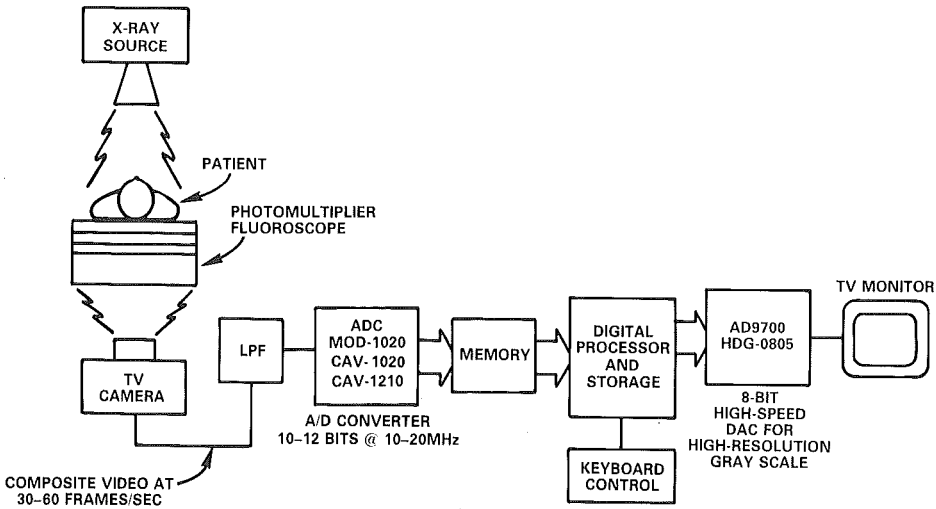


Figure 13.2. Digital fluorography.

(light). The low-level light signals are scanned by a high-quality TV camera, that changes the light into analog electrical signals. The resulting composite-video signal is updated at from 30 to 60 frames per second to eliminate flicker, low-pass filtered, and applied to an a/d converter. The digital data from the converter are stored in memory; and the results of the subsequent image processing can then be stored and/or displayed (essentially in real time, as perceived by the physician) for examination and analysis.

The need for good resolution and wide dynamic range imposes stringent requirements on the system a/d converter, to permit as much information as possible to be captured for use in processing the image. Typically, the ADC must have 10- to 12-bit resolution and be capable of 10-20 MHz word rates.

At the other end of the system, the d/a converter used in the raster-type display monitor must have high gray-scale resolution (for example, 8 bits), so that subtle gradations in tissue density can be observed; if the processor produces an image in color, a triple DAC with good resolution will be needed to display the red-green-blue color pixels.

Compared to conventional x-ray techniques, digital radiography offers improved diagnostic images with less radiation exposure, permitting a safer, less invasive diagnostic procedure. Also, the system's memory and sophisticated information-processing capabilities provide faster image generation and increased flexibility in image display and analysis, enabling the physician to arrive at a diagnosis, and appropriate course of treatment, more quickly.

Digital radiography can be subdivided into categories based on system characteristics, such as types of x-ray beams and detectors, number of pictures taken per second, the type of digital processing used, and—of course—the ultimate clinical application.

A typical example of a major use of digital radiography is *angiography*, an x-ray technique for examining blood vessels and the major organs they supply, especially the heart, head, and neck. Typically, a contrast dye, such as radioisotope iodine, is injected into the patient, and an x-ray scan is taken. The resulting image is then compared with earlier images (or *masks*) of the same area, obtained prior to the injection.

The comparison takes place in the processor by means of a digital subtraction algorithm, which isolates the iodine-infiltrated vessels and organs by eliminating all unaffected tissues from the image. Besides subtraction, the processor also uses contrast enhancement, averaging, edge enhancement, magnification, and other image-processing techniques to help present maximum information to the medical diagnostician.

Digital subtraction angiography (DSA) has emerged as a relatively safe, cost-effective way to examine blood vessels. Patient safety is enhanced, because the technique reduces the need for invasive arterial catheterization. Lower doses of the contrast dye are required for DSA, because image-processing subtraction and enhancement can amplify the image data; this, in turn, improves patient comfort and reduces motion during the procedure. The cost savings are also an important consideration. Elimination of the need for catheterization in many cases can replace a two-day hospitalization by an out-patient procedure. Since the data can be saved on inexpensive magnetic disks or tape, there is also a considerable saving of the cost of film and its storage.

Phased-Array Ultrasound

Phased-array ultrasound is most aptly described as a very-high-frequency sonar system. A typical arrangement is shown in Figure 13.3. An array of from 16 to 32 piezoelectric (crystal) transducers is arranged in a hand-held pickup head, which is placed against the patient's body in the area to be examined. On command from the operator, the transducers are excited with an electrical signal which causes each transducer to transmit a high-frequency (supersonic) pulse.




Figure 13.3. Phased array ultrasound.

The excitation is applied to the transducers in some sequence, rather than simultaneously, resulting in a spatially swept acoustic signal; the pattern can be varied by computer control for optimal results in various kinds of examinations.

As the sonic signals impinge on bone and tissue within the patient's body, echo signals will be returned to receiver elements in the array, whose focus is swept in synchronization with the range of returning echoes. These sonic echoes are transformed into electrical signals by the transducers, individually amplified, and applied to each individual channel's a/d converter. The ADC may have a resolution between 6 and 8 bits, and must be capable of word rates between 15 and 20 MHz (and more), for digitizing return signals with bandwidths of from 5 to 10 MHz.

After processing, the signal can be stored and/or displayed on a TV monitor, using a high-speed wide-dynamic-range (8-bit gray scale) d/a converter. Good image reproduction quality is required, because ultrasound systems are real-time diagnostic tools.

The high degree of mobility of the hand-held pickup head used in phased-array systems has made this technique particularly attractive for the study of the heart, because ultrasound is non-intrusive and lessens the physical dangers for the patient. Large strides have been made in recent years in examination techniques and interpretation of the results in this segment of medicine, called echocardiography.

Phased array is just one of several sub-categories in ultrasonic medicine. Others include linear array and mechanical-sector scanning. All are highly demanding on converter performance.

13.1.3 RADAR

In theory, all radars operate in essentially the same way, i.e., a pulse is transmitted from an antenna, and its echo is received from a target. The measurement of the echo's return-time (*range*) and other target information contained in the return pulse are extracted via signal processing.

In *monopulse radars*, the direction from the radar to the target is determined by the direction in which the antenna is pointed. In *tracking radars*, the antenna is continually re-positioned to follow the flight of the target. *Acquisition radars*, designed to show the presence of targets in the zone of interest, are not designed to track the target; instead, they scan periodically to find and indicate that one or more targets are in the vicinity. Performance is limited by the speed with which the familiar parabolic *dish* antenna can be moved.

For most modern applications, the monopulse radar has been supplanted by the *phased-array radar* system, in which the beam is steered electronically. A phased-array antenna comprises a multitude of transmit/receive antenna elements mounted on a flat plane, which is usually fixed in place. The number

of elements ranges from 20, or less, for airborne radar, to hundreds for large ground-based radars, such as might be used in missile-tracking systems.

Phased-array systems use the same antenna as both a search radar and a tracking radar by varying the phase- and time delays among the many transmit/receive elements via computers, which are an integral part of the system. In the search mode, the beam is focussed broadly and transmitted at right angles to the plane of the arrays; in tracking, one or more beams are focussed more narrowly and electronically steered as necessary to pinpoint targets and provide accurate tracking data on their locations.

Switching among the many search and track conditions is accomplished at very high rates; this multiplexing achieves the effect of using a single antenna as some number of antennas. Large phased-array systems can outperform parabolic antenna systems of equal size, and they can operate at lower frequencies, of the order of 500 MHz (vs. frequencies in the GHz range).

Processing a radar return signal, or echo, to obtain the maximum amount of information, is a complex procedure, but the technique is not markedly affected by the type of antenna system. Figure 13.4 shows the technique for processing signals in "I and Q" (*in-phase* and *quadrature*) radars. In these systems, the return pulse is separated into two channels of information with a 90° phase difference.

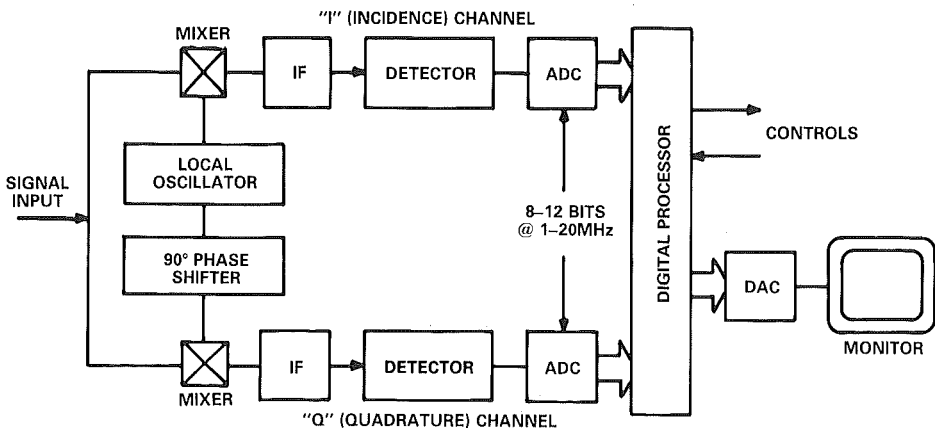


Figure 13.4. I & Q radar processing.

Direction and ranging information is rather straightforward; it is just a matter of knowing the direction in which the antenna is pointed and measuring the time interval between the transmission of the pulse and the receipt of its echo. However, the I and Q technique makes it possible to obtain more information about the target more quickly. Processing a set of returns with I and Q information allows the system to determine the size, speed, direction of travel, target type, and other parameters of the target.

The reason for this is that targets have distinctive radar “signatures,” which cause echo signals to differ slightly from one another, depending on the type of target. This kind of data can be extracted more easily with signal-processing algorithms when the return pulse has been split into its in-phase and quadrature components. Most modern radars use this kind of processing, regardless of how the antenna is positioned.

Radar systems call for fast, moderate-to-high-resolution a/d converters, with word rates from 1 to 20 MHz or more, and 8 to 12-bit resolutions, depending on the function of the radar, the location of the converter within the processing chain, and its expected role in the signal-processing function. The higher the speed and resolution, the faster and smaller the targets that can be followed; this is especially important where high signal-to-noise ratios are required.

High-speed, high-resolution converters (10 MHz/10 bits or better), in association with fast processing systems, make it possible, in special cases, to process the signals at intermediate frequencies (IF). This eliminates the need for physical detection (demodulation) circuitry, since any desired detection function can be achieved with high accuracy and repeatability—and low noise—in the course of processing.

For return to the analog world for display and/or recording, fast deglitched or low-glitch d/a converters are used in both vector and raster-type displays, typically with 12-bit and 8-bit resolutions.

13.1.4 TRANSIENT DETECTION

Transient recorders (Figure 13.5) are instruments which perform the function described by their name—they capture, for later analysis, unique “one-shot” events, such as those encountered in nuclear or seismic applications.




Figure 13.5. Transient recorder.

The input signal-conditioning circuits, the a/d converter, and the memory are all under the command of clock and timing circuits, which control the system's front-end gain, operate the memory's write circuits, and dictate the application of encode commands to the converter.

In one mode of operation, new conversions are triggered at a constant rate, and data flows continuously through a first in-first out buffer memory until an input threshold is exceeded, indicating the presence of an event; then all

data within a predetermined time before and after the occurrence of the event may be preserved.

Once stored in buffer memory, the user's program determines whether the data is to be read out into longer-term memory, displayed on a CRT, or processed.

A/D converters used in transient detectors range from low resolution and very high speed (e.g., 6 bits at 50 to 100 MHz word rates) to high resolution at lower speeds (e.g., 12 bits at 20 MHz).

13.1.5 COMMUNICATIONS

Baseband Characteristics

The *baseband* is the band of frequencies occupied by the signal before it modulates a carrier or subcarrier frequency to form the transmitted signal. An important objective of communications receivers in dealing with a received signal is to find out the baseband, if unknown, discard the carrier, and process the signal to determine what information it contains. This is easily accomplished with "friendly" signals; it is much more difficult with intercepted signals—a great deal of processing, analysis, decrypting, etc., may be required. Converters having high resolutions and the highest speeds are essential in order to minimize loss of data.

Fibre Optics and Satellite Communication

An increasingly common communications technique in recent years has been the use of a/d and d/a converters in digitizing, transmitting, and recovering baseband information via fibre optics and satellites. Fiberoptic links are less costly than coaxial cable, easier to maintain, and more secure from noise and interception. In these applications, the baseband frequencies are generally dc to 5 MHz, which is sufficient bandwidth for television or frequency-division multiplex (FDM) voice signals (Figure 13.6).




Figure 13.6. Fibre optic/satellite communications.

In typical systems, the analog input is filtered to prevent aliasing, then converted to parallel digital data with medium resolution (8 to 10 bits). To maintain the integrity of signals containing frequencies as high as 5 MHz, the word rates for sampling and conversion must be at least 10 MHz, to satisfy the Nyquist criterion, but are greater in practice to avoid degradation of data.

The parallel data are changed to a serial stream of data at rates of 200 megabaud or more, using a high-speed clock. In a fiber-optic link, the serial data stream modulates a laser diode, which emits pulses of light; this pulse train is transmitted via the link to a photodiode, which is sensitive to light at the same (light) frequency as that emitted by the laser diode.

The electrical pulse-train output of the photodiode is returned to parallel by the deserializer, and then converted back to analog by a d/a converter of appropriate speed and resolution.

Typically, a fiberoptic link can carry signals at 300-MHz data rates, which are adequate for up to three multiplexed digital television signals, or some combination of TV and FDM voice signals.

Satellite communication systems are conceptually similar to the system of Figure 13.6, but the laser diodes are replaced by microwave transmitters and receivers operating at gigahertz (GHz) frequencies.

FDM/TDM Transmultiplexers

In telephony, there are two standard formats for multiplexing voice signals. Frequency division multiplex (FDM) has been used throughout the world for many years to transmit long-distance telephone calls; time-division multiplex (TDM) principles have been known for some time but have been widely applied only recently.

FDM stacks voiceband signals in adjacent 4-kHz channels into 12-channel groups and 60-channel supergroups, using single-sideband amplitude modulation; in TDM, each voice signal is digitized, using pulse-code modulation (PCM), at a sampling rate of 8 kHz; The pulse streams which result are then interleaved in time and transmitted.

Inevitably, systems employing FDM and those employing TDM must communicate. To make this possible, digital signal processing provides a convenient means of translation, the FDM/TDM transmultiplexer. However, FDM signals must be digitized before they can be translated to TDM; this calls for an a/d converter. On the other hand, after TDM signals are translated digitally, they must be converted to analog. Figure 13.7 shows a translation process from FDM to TDM for handling signals from the outside within a given office—and ultimately back to FDM again when communicating with the outside system.

The special-purpose matched ADC and DAC used for this purpose, strictly speaking, are not video converters. Nevertheless, their most important




Figure 13.7. Digital FDM/TDM translation.

specifications are based on their ability to deal with high-frequency ac signals, a feature shared with video converters when used for signal-processing applications. In these applications, such characteristics as noise-power ratio (NPR) and signal-to-noise ratio (SNR) are important. The devices labeled in the block diagram are high-performance 14-bit types designed for high resolution and good mid-scale linearity; they operate at a word rate of 112 kHz.

13.2 ABOUT VIDEO CONVERTERS

13.2.1 D/A CONVERTERS

The accuracies of d/a converters at low speeds are primarily a function of the design of the resistor network and switches, the stability of the reference voltage, and the nonlinearities of the op amp (if used). At low speed and resolution, most modern d/a converters exhibit good linearity and monotonicity. However, at video speeds, full-scale settling time and the magnitude and duration of “glitches” become crucial.

DAC Settling Time

For a d/a converter, full-scale settling time is the worst-case time elapsed from the application of a digital input, representing a full-scale change in the analog variable, until the time when the output has entered (and remains within) a specified error band around its final value (usually given as a percentage of full scale or $\pm 1/2$ LSB). This figure takes into account all internal factors affecting settling time, i.e. turning the switches on and off, current changes within the resistor network, and time required by op-amp or buffer outputs to settle to within their error bands.

DAC Glitches

Glitches, or output spikes caused by *skew* (differences in turn-off and turn-on times of switches or logic), represent another error source in very high speed d/a converters. All DACs produce output glitches that occur during digital input transitions. The worst spikes occur at the major transitions, at half full scale (when the MSB changes state), and at $1/4$ and $3/4$ scale, when the next-most-significant bit change state. Figure 13.8 illustrates the glitches at these transitions (and there are glitches of lesser magnitude at other transitions, not shown) for increasing input—in the case where the lower-order bits turn off faster than a higher-order bit can turn on.




Figure 13.8. Glitches in d/a converter outputs.

DAC glitches cause significant distortion and create real problems in high-speed automatic test equipment and digitally controlled CRT displays. In CRTs the glitches actually appear on the screen as either shaded stripes, discontinuities or varying intensities in straight lines. In other applications, glitches can produce in-band harmonics, and generally higher noise levels.

Glitches can be characterized by measuring the "glitch impulse" (sometimes called glitch energy, or glitch charge), as Figure 13.9 illustrates for a glitch having a doublet shape. The glitch impulse is the net area under the voltage-versus-time curve and is expressed in picovolt-seconds, or millivolt-nanoseconds. It can be estimated by approximating the waveforms by triangles, computing the areas, and subtracting the negative area from the positive area, as illustrated in Figure 13.9. In a well designed DAC, the net glitch impulse can be made much less than 100 picovolt-seconds.




Figure 13.9. Minimal-area (glitch-impulse) doublet glitch.

100

Glitches in high speed d/a converters result from input bit timing differences, asymmetries in on-to-off and off-to-on times of d/a current switches, and circuit layout. Input bit timing differences, for instance, can cause severe glitches. Time skew in the parallel digital input will increase the width of the glitch by an amount equal to the skew. This time skew can be greatly reduced by adding a set of input registers either in the d/a or very near to it.

TTL signals, by nature, contribute significantly to the time skew. TTL logic utilizes saturated voltage-switching devices and has significant differences in delays and rise times between the positive-going and negative-going changes in logic levels. These differences contribute directly to the time skew at the front end of the DAC. For these reasons, a TTL-compatible d/a converter usually will have significant glitches and in most cases needs to be "de-glitched" as will be illustrated later.

Emitter-coupled logic (ECL), on the other hand, utilizes non-saturated current switching, and the delay times for positive and negative-going transitions are almost identical. Acceptable 8-bit performance can usually be obtained from an ECL d/a converter by adding small delay-equalizing deskewing capacitors on the first few most significant bits that drive the DAC, as shown in Figure 13.10. Careful board layout and selection of input registers help to decrease glitches.

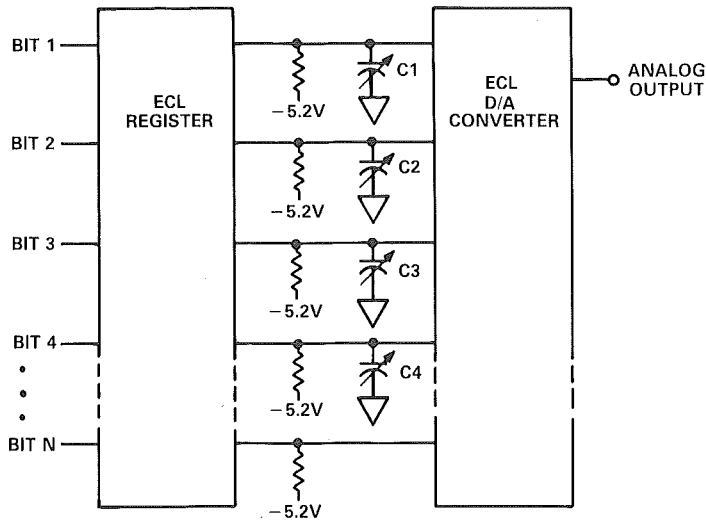


Figure 13.10. Trim Capacitors to neutralize skew.

Since the glitch in d/a converters is a nonlinear and code-dependent phenomenon, it cannot be simply filtered. One of the most effective methods to reduce the effects of glitches associated with video DACs is to use a track-and-hold amplifier (T/H). In Figure 13.11, a track-and-hold is connected to operate on the DAC's analog output. The T/H tracks the DAC output until just before the glitch occurs; then the T/H switches into the "hold" mode and remains there until the glitch has settled to an acceptable level, after which it returns to the track mode.

As Figure 13.11 illustrates, the Hold command begins just prior to and ends just after the worst-case glitch. This effectively masks the widely differing glitches associated with the D/A and instead introduces nearly identical T/H-related pulses at a frequency equal to the update rate. The track-and-hold may also introduce a pedestal (dynamic offset) error, but both the T/H pulse and pedestal errors can be significantly reduced through careful component selection and circuit design. Since the track-and-hold related pulses are consistent in magnitude and frequency, they can be readily low-pass filtered. The T/H technique loses effectiveness for update rates greater than about 20 MHz.

A viable alternative in very fast converters is to minimize the overall magnitude of the glitch by careful balancing of both the logic and the switching skews. If effectively performed, this can further reduce the glitch's net area by forcing it to be a fast doublet, which then becomes more susceptible to filtering, because it is at a doubled fundamental frequency and has no large one-sided excursion requiring a long settling time in linear filters. Though this is difficult to accomplish in TTL circuits because of the inherent asymmetry of saturated logic, nonsaturating ECL is essentially independent of the direction of the logic transition and can be deskewed (as noted above) by the addition of small capacitors.

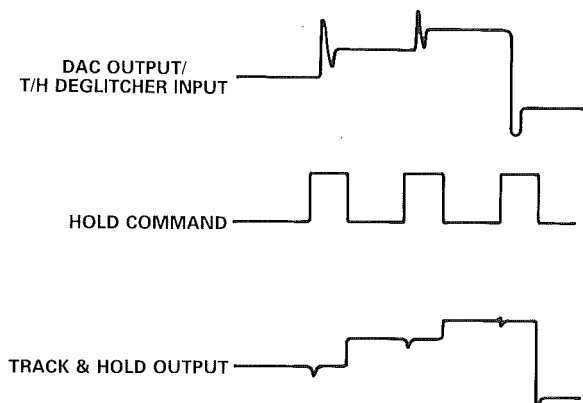
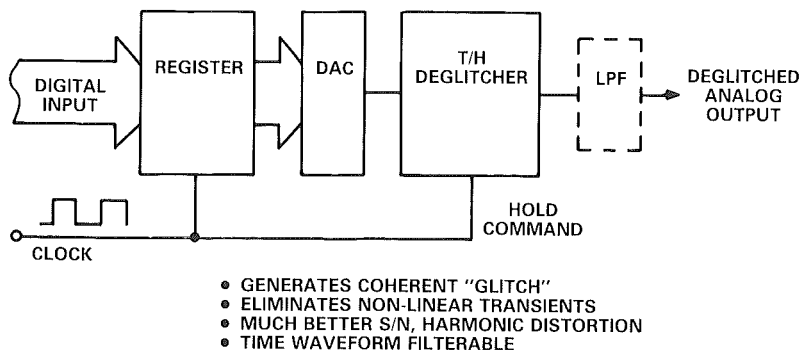


Figure 13.11. Deglitching a d/a converter.

12-Bit Deglitched DAC

Figure 13.12 is the block diagram of a high speed 12-bit DAC, which embodies—within a single thick-film 32-pin hybrid package—the deglitching techniques discussed above. The package includes the 12-bit current-output DAC network and switches, a current reference, an output amplifier with track-and-hold switching, timing circuits, and the associated electronics.

As a first step in minimizing the glitch, an input register removes the effects of differences in the arrival times of the individual bits. All bits are simultaneously latched and arrive at the input of the current-output DAC at the same time. Nevertheless, despite the best efforts at balanced design for the converter, there still exists a difference between on and off times for the switches, causing a glitch.

The block labeled “timing generator” includes a one-shot multivibrator and gating circuitry. This circuit provides a fixed hold period, starting at the time the strobe arrives and continuing beyond the expected glitch duration, irrespective of the strobe frequency. At the conclusion of the hold interval, the output amplifier is returned to the track mode, slews to the new value established by the digital input and settles quickly.

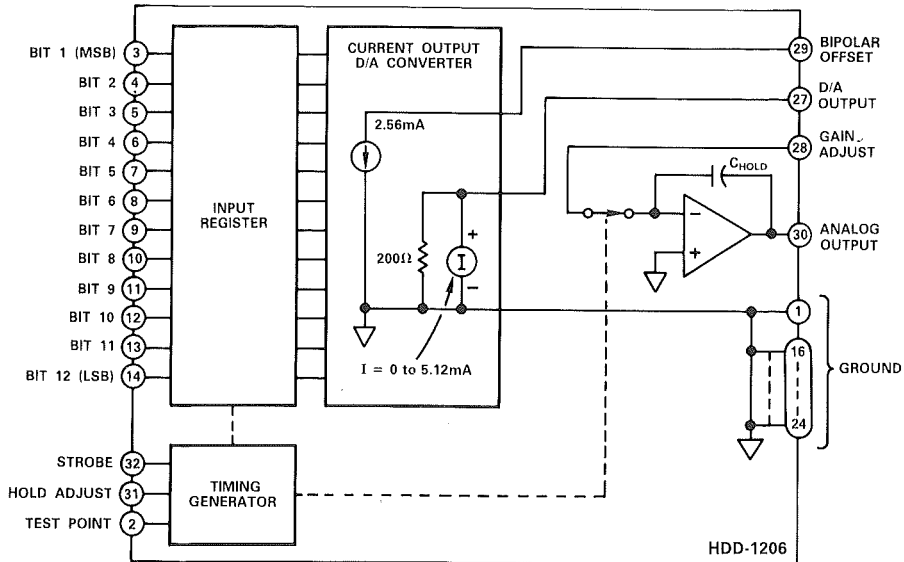


Figure 13.12. Functional block diagram of 12-bit deglitched DAC.

The role of the FET switches used in the track-and-hold is critical, because they affect the dissipation and complexity of the drive circuit, and the smoothness of the converter's output waveform. Chosen for low input capacitance, low gate cutoff voltage, and low drain-to-source "on" resistance, they make possible 6-MHz update rates with low power dissipation. Additional benefits are low charge transfer within the device—hence small residual spikes and simplified drive circuitry, resulting in efficient use of space and low cost.

DACs for Raster Scan

High speed DACs find growing use in raster-scan video-display applications. Devices especially designed for these applications not only meet the speed requirements, but also simplify the generation of composite video waveforms by providing digitally controlled auxiliary output currents of appropriate magnitude and polarity for blanking, synchronization, and 10% brightness levels.

Figure 13.13 shows the industry-standard video composite intensity waveform over $1\frac{1}{2}$ cycles of the horizontal sweep. The controlled range of the D/A converter's full scale (0 to -643mV) is from reference white (-71mV) to reference black (-714mV). In the illustration, the intensity is varying from full white to full black. At the beginning of the sync portion, the intensity signal drops to the blacker-than-black "front porch" (-785mV), and then to the extreme black level (-1071mV) during the horizontal retrace. As the next sweep starts, the intensity returns to the "back porch" (-785mV), and, as the first element of the picture is triggered, to the




Figure 13.13. Composite DAC output waveform.

controlled range of the D/A converter. During this scan, the 10% “brighter than white” level (0mV) has been activated briefly to display the cursor.

Typical DACs for this purpose, as exemplified by the 8-bit HDG-0805 DAC diagrammed in Figure 13.14, are specifically designed to meet the needs of raster-scan systems. Housed in 24-pin metal hybrid packages with 0.6" double-DIP spacing, they require only a single $-5.2V$ power supply. Such a DAC provides, in a single package, all the circuitry required for 256-level intensity modulation in raster-scan displays at video dot rates up to 100 MHz.

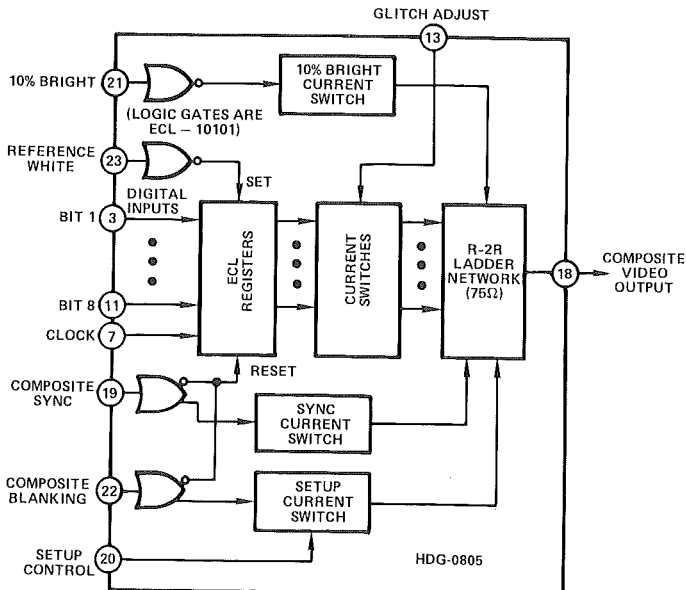


Figure 13.14. Block diagram of HDG-0805 D/A Converter.

For the HDG-0805, resolution is 8 bits, and full-scale settling to within 1 LSB is typically 7 nanoseconds. The output impedance is 75 ohms, thus simplifying impedance matching to video coaxial cable. Full-scale output current develops 1 V p-p video across a 75-ohm load. In order to minimize the glitch, a set of internal ECL registers provides minimum time skew between bits. Typical net glitch impulse is 50 picovolt-seconds. A Glitch-Adjust input lets the system designer optimize glitch performance.

As the block diagram shows, such devices provide self-contained digitally controlled sync and blanking capability, compatible with EIA Standards RS-170, RS-330, and RS-343A, to produce composite video waveforms like the one in Figure 13.13. In addition to the 256 levels of gray scale provided by the 8-bit digital input, three additional digitally controlled switches independently provide auxiliary output currents of appropriate magnitude and polarity for blanking, sync, and 10% bright levels.

Because of their small physical size, display DACs in the form of hybrids or ICs can be located quite close to the CRT's intensity input, eliminating degradation caused by coaxial-cable-related losses. The grounded metal package effectively screens out the effects of the high noise environment usually associated with CRT drives.

13.2.2 A/D CONVERTERS

The challenges of high speed a/d conversion are being increasingly met by a wide range of components from a variety of manufacturers. Unfortunately, in many cases, these devices do not achieve their full potential because of improper application.

Flash Converters

All-parallel, or so-called *flash* converters offer the fastest throughput of available ADC designs. Figure 13.15 illustrates a typical block diagram. The converter employs $2^n - 1$ latched analog comparators in parallel, where n is the number of bits. A resistive voltage divider provides the reference voltages for the comparators. The reference voltage for each comparator is one least-significant bit (LSB) higher than the reference voltage for the comparator immediately below it.

When an analog input signal is present at the input of the comparator bank, all comparators which have a reference voltage below the level of the input signal will assume a logic "1" output. The comparators which have their reference voltage above the input signal will have a logic "0" output.

Decoding logic then converts the comparators' outputs into a binary digital output. Because flash encoders having 8-bit resolution require 255 comparators, and comparable amounts of logic, it can be easily seen that these converters are relatively impractical to construct from discrete comparators and logic elements due to power, size, wiring, and cost considerations. Advances

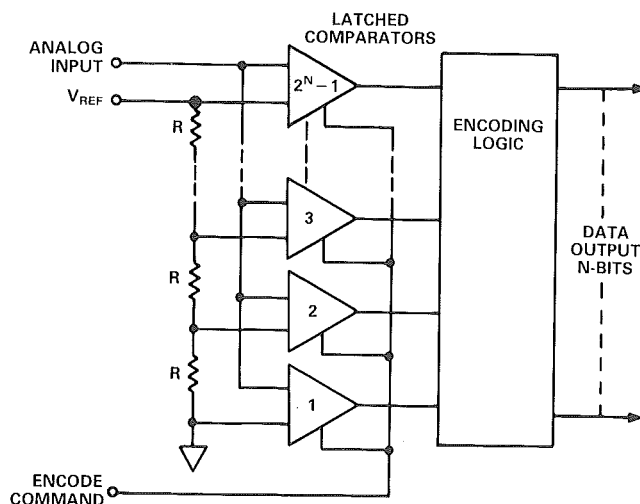


Figure 13.15. All-parallel a/d (flash) encoder.

in large scale integration techniques, however, have resulted in the development of commercially available monolithic devices ranging from 6 to 10 bits of resolution.

Even though the latches essentially perform a track/hold function, high-speed, high-resolution systems require an external track/hold for best performance. However, in low-resolution applications, if the relative aperture delay match between each comparator and the total system aperture uncertainty is better than, for example, about 100 picoseconds, no track-and-hold circuit is required for video signals having bandwidths up to 4 MHz.

Figure 13.16 is a functional block diagram of a 6-bit flash converter, the AD9000 a/d converter, packaged in a standard ceramic 16-pin DIP. It achieves 75 MHz word rates over a temperature range extending from -55°C to $+125^{\circ}\text{C}$, making it useful for a variety of applications in a wide diversity of environments. As the block diagram shows, 64 parallel comparators are employed to digitize fast-moving analog input signals. An overflow bit makes it possible easily to connect multiple units in a cascade arrangement to obtain up to eight bits of digital data at word rates comparable to those achieved by the devices operating independently. Wired-OR logic circuits within the device encode the comparator outputs into a binary format of six bits of parallel data, along with the overflow bit.

The outputs of the comparators are applied to latches controlled by the ENCODE input. When the encode command is low (digital "0"), the latches are transparent creating the "track" mode. When the ENCODE input changes to high (digital "1"), the latches go into a "hold" or latched condition, thus freezing the most recent digital outputs of the comparators and applying them to the encoding circuits.




Figure 13.16. AD9000 block diagram.

The signal held in the latches is converted to binary form by the encoders and applied to the output stages as a 6-bit digital representation of the analog signal which was present at the comparator inputs at the instant the ENCODE command made the change to the "hold" mode.

All-parallel flash a/d converters tend to have fairly random linearity errors. The overall linearity in monolithic flash converters is determined primarily by comparator offset-voltage matching and tolerance of the resistors comprising the voltage divider. Codes can be missed if adjacent comparators have offset voltages of opposite polarity and sufficient magnitude. Figure 13.17 shows a typical error characteristic (the difference between a linear unit slope and the typical ADC staircase of Figure 7.2) for a monolithic flash converter.




Figure 13.17. Error Characteristic (linear analog response - quantized staircase) for a/d converter. Only a few codes are shown.

Due to the high input capacitance associated with the large number of parallel comparators, the input to the encoder portion of such a converter must be driven with a video operational amplifier which is stable when driving a large capacitive load.

Using Track-Holds

Traditionally, designers have tended to use flash converters without track-and-holds, because the internal latches of the converter perform a track/hold function. However, there are distinct advantages to using a T/H.¹




Figure 13.18. Typical flash ADC-track/hold subsystem.

Figure 13.18 illustrates the use of a T/H with a flash ADC. The T/H “holds” the analog input constant while the comparators are settling. Then the comparators are latched and, while the data is being encoded to complete the conversion, the T/H returns to the “track” mode to acquire the new analog value for the next conversion. When the T/H is switched to hold, the comparators are unlatched and switch to the new state. Thus, the comparators are always converting a new “dc” value, rather one which is continually changing. This process provides a substantial improvement, for reasons discussed below.

As indicated earlier, all flash converters have inherent capacitance characteristics which affect their ability to operate on high-frequency analog signals. There are sets of problems that are specific to each of the inputs.

The analog input capacitance, C_{ai} , of each comparator (Figure 13.19) is determined primarily by the base-emitter junction capacitance of the transistor on the analog-input side. The value of this capacitance is affected by junction bias; forward-bias capacitance is higher than reverse-bias capacitance. Since the inputs of all comparators are in parallel, the total capacitance at the flash converter’s analog input is the sum of the individual comparators’ base capacitances.

¹See “Flash Converters Work Better with Track/Holds,” by Jerry Neal and Jim Surber, *Analog Dialogue* 18-2 (1984) 10-15.

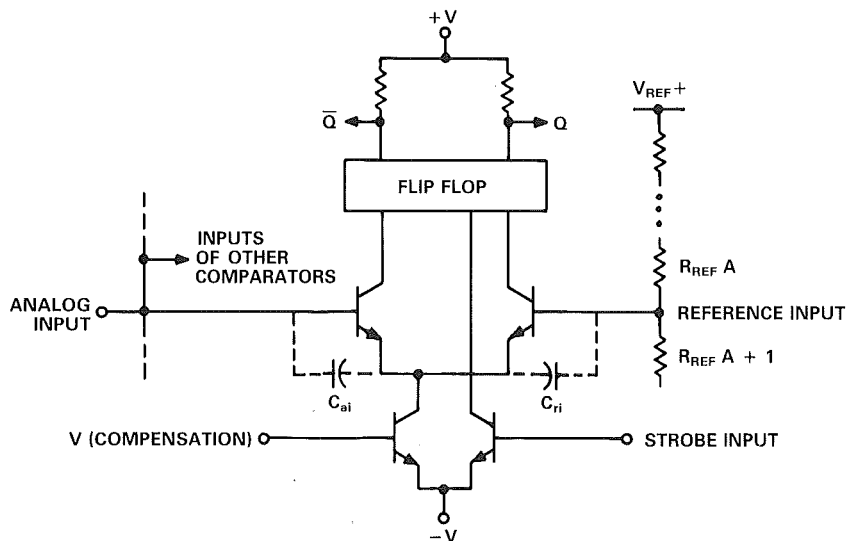


Figure 13.19. Typical flash ADC comparator cell.

Unfortunately, the total capacitance is essentially random, because it depends on the amplitude of the analog input signal; the signal amplitude and the on-off state of each individual comparator dictate its base-emitter voltage, hence its contribution to the total capacitance. The total capacitance can vary from 30 to 120 pF. Since the comparator can be driven by an analog buffer, this is not a serious problem (but it *must* be dealt with). However, the buffer doesn't solve capacitance problems associated with the reference input.

Since the comparator has a differential input, the reference input is similarly subject to capacitance variations introduced by its base-emitter capacitance, C_{ri} . But the problems they raise are not as easily dealt with; see Figure 13.20.




Figure 13.20. Equivalent circuit of the reference resistor network.

The charging path for each individual C_{ri} is through the reference ladder. This path constitutes an R-C time constant at the reference input; and the time required to charge the capacitances through the resistors tends to cause the reference voltages to lag fast changes in the current flowing through the reference ladder, dynamically distorting the voltage division of the reference ladder.

This capacitive reactance effect depends on the differential input voltage to the comparator, the speed at which it varies (frequency), and the comparator's position within the ladder. When signals having high slew rates are applied, the ac integral-linearity error of the comparator array will increase because of the reactance variations among the many comparators.

In addition, errors also occur due to noise introduced by parasitic coupling from the strobe circuit to the reference ladder via C_{ri} . Although the strobe is equally coupled to both sides, the lower source impedance at the analog side tends to reduce its amplitude, resulting in an imbalance between the comparator inputs.

One of the most effective ways to equalize the differential input impedances of the comparators is to use an external resistor in series with the analog input of the converter, as Figure 13.18 shows, to produce better common-mode rejection of unwanted noise voltages and improve the integral linearity of the comparator array. The resistance value should be about one-fourth of the total resistance of the reference ladder (i.e., the parallel resistance of the upper and lower halves of the ladder resistance).

Another potential source of dynamic conversion error in flash ADC designs is the variation of effective sample delays. Individual comparators within an array can be visualized as having variable delay lines in series with their latch inputs. The magnitude of delay for each comparator is determined by comparator inconsistencies, chip layout, and the strobe frequency.

For low-frequency analog inputs, these sample-delay variations among adjacent comparators are not a significant problem. But as the input frequency is increased, latch-time disparities among comparators can result in missing codes and excessive differential nonlinearities. The errors differ with slew-rate magnitude and direction. As a practical matter, it is realistic to expect sample-delay variations as large as 200 to 300 picoseconds, rather than to calculate them based on just the published aperture-jitter specifications.

Considerable improvement can be realized if we can replace the sample-delay variations of a flash converter by the much faster switching of a track/hold. For example, the track/hold shown in Figure 13.18 has a specified aperture uncertainty of only 5 picoseconds. Since the track-and-hold amplifier ahead of the flash converter freezes the input signal, it maintains a constant input while the comparators are latching, and it can also be timed to allow the comparator input capacitances to charge and settle before the conversion takes place.

It is important to remember that a track/hold amplifier used ahead of a flash converter has no effect on the conversion time; it simply allows the converter to digitize higher-frequency (faster slew rate) analog signals. The total time required for conversion, taking into account the delays of both the track/hold amplifier and the flash converter, can be calculated and compensated for by the system timing.

Multi-Stage Digitally Correcting Subranging A/D Converters

When higher resolutions than those obtainable with flash converters are needed, a multi-stage converter is required. Multi-stage converters, also referred to as subranging converters, provide high resolution with markedly fewer comparators and simpler logic than single-stage converters, but give up some of the inherently higher-speed capabilities of the single-stage types. Subranging converters are used instead of parallel-, or flash-type, converters whenever higher resolution and/or less complexity are required for the conversion function.

Figure 13.21 illustrates a 12-bit subranging a/d converter constructed with two encoders having resolutions that add up to 13 bits (6 and 7, in this case). The analog signal from the track-and-hold is applied through the two buffer amplifiers to a 6-bit flash encoder and a video delay line simultaneously. The 6-bit encoder converts the analog signal to binary, producing the six most significant bits. These bits are stored in a register and also applied to a 6-bit d/a converter having an accuracy of at least 12 bits. The output of the d/a converter is inverted and subtracted from the delayed track-and-hold output in a




Figure 13.21. Functional diagram of two-stage flash converter with digitally corrected subranging.

summation network, and the resulting “residue” is amplified to the 7-bit converter’s full-scale span.

The “residue” is then converted to digital form by a 7-bit flash encoder and represents the less-significant portion of information. The outputs of the 6-bit holding register and the 7-bit flash are then combined in an output register to yield a 12-bit parallel binary output.

Timing is of extreme importance in this type of a/d, as each element in the conversion process must settle to its optimum point before the required strobe signals are applied.

It would not be at all unusual for converters of this type to exhibit differential-linearity discontinuities around the bit-6 transition points, due to mismatch between the front-end and residue encoding circuits. However, in converters employing digitally corrected subranging (DCS), these discontinuities are eliminated by the use of digital correction logic, since the information to accurately characterize the bit-6 transition is already present in digital form because of the accurate d/a conversion and summing, and the 7th bit in the second conversion.

Here’s a simplified explanation of how it works: The result of the first conversion—in the holding register—is equal (digitally) to: (Input – Residue + 6-bit Error). It is converted to analog and subtracted from Input. Assuming a perfectly accurate DAC and subtraction, the result is (Residue – 6-bit Error). That is,

$$I - (I - R + E_6) = R - E_6 \quad (13.1)$$

This difference is scaled up and converted, giving a digital quantity equal to (Residue – 6-bit Error + 12-bit Error). Finally, the two digital words are added, giving: Input + 12-bit Error. That is,

$$(I - R + E_6) + (R - E_6 + E_{12}) = I + E_{12} \quad (13.2)$$

The extra bit is necessary because $(R - E_6)$ can exceed 1 LSB of 6 bits.

Although the explanation is simple, the execution is not. Fast 6-bit DACs with better than 12-bit accuracy are not easy to make, nor are fast subtracting amplifiers with adequate dynamic linearity. Remember that both the DAC and the output amplifier must settle before the second conversion can be completed.

However, this technique, incorporating fast emitter-coupled logic, has resulted in 10-bit converters capable of 40-MHz sampling rates and 12-bit converters that can handle 20-MHz sampling, with similar architecture to the one described by Figure 13.21 and depicted in Figure 13.22. The technique seems promising for even faster and higher-resolution converters.

¹See “Flash Converters Work Better with Track/Holds,” by Jerry Neal and Jim Surber, *Analog Dialogue* 18-2 (1984) 10-15.




Figure 13.22. Physical construction of the converter of Figure 13.21.

13.3 PRACTICAL DESIGN INSIGHTS

Much of the challenge in using high-speed converters lies, not in the electronic design, but rather, in the physical design and layout of the printed circuit boards. It is the old problem, reiterated many times in these pages, of the need for high-resolution analog signals to maintain their integrity in the presence of fast, pervasive digital edges. Grounding problems, cross-coupling, and parasitic effects associated with the physical circuit layout can degrade the performance of even the best designs.

13.3.1 BOARD-LEVEL GROUNDS

In general, a very low impedance ground is a must. Every portion of a printed circuit board which doesn't contain circuits or conducting runs should be ground plane. Sometimes circuit density must be purposely reduced in order to create more room for ground plane. All breadboards should be constructed on double sided copper clad boards. Avoid using general digital (purely insulating) breadboards and thin hookup wires.

Another basic rule for working with high-speed (high-frequency) printed-circuit (PC) board designs is to connect analog ground and digital ground together within the PC board. Board designs that employ separate digital and analog grounds often end up with ground loop problems that are very hard to solve. But there is a dilemma: although connecting the two grounds together at the board enhances the performance of converters at the board or subsystem level, it can create other problems at the system level. This problem will be discussed below.

As another practical rule in maintaining low ground impedance, use as many of the PC board's connector pins as possible (wired in parallel) for ground connections to the system. This lowers the contact resistance, thus avoiding IR-drop related noise. For perspective, consider that, for a 10-volt input range on a 12-bit a/d converter, the least significant bit (LSB) of the converter will have a value of only 2.5 millivolts (high-speed systems generally use even lower voltage ranges). Assume that a single pin of the PC connector, to be used for ground, has a resistance of 0.05 ohm—and that the PC card draws a total of 1.5 amperes (not unusual for cards carrying high-speed logic circuitry).

The subsequent voltage drop at the ground pin could be as much as 75 millivolts. If only digital logic were used, this voltage drop would be minuscule and hardly worth considering. However, if the circuit involves both analog and digital signals, and if a fraction of that voltage drop is coupled into the analog circuit, the effects could be devastating for converter accuracy.

Consider, for instance, the case of TTL logic. Since TTL is a saturated logic, ground currents vary widely, and varying current flowing through the ground often produces noise signals which modulate the ground plane. This noise, created by digital switching, can couple into the analog portion of the circuit. If only 10% of the 75 millivolt IR drop cited above couples into the analog signal, that would represent 3 LSBs of lost converter resolution. Connecting multiple pins in parallel to reduce effective contact resistance helps solve the problem.

13.3.2 LAYOUT

Component layout is just as important a consideration as grounding. A massive low-impedance ground will simply not help a poor basic layout. Signal cross-coupling generally represents the biggest problem. Digital signals should not couple to analog signals, and the analog channels must remain isolated from each other. Coupling between analog channels often leads to oscillation. Any subsystem or circuit layout operating at high speeds with both analog and digital signals needs to have those signals physically separated as much as possible to prevent crosstalk.

Digital signals leaving or entering the layout should use runs that have minimum length. The shorter the digital runs, the less the likelihood of coupling to the analog circuits.

Analog signals should be routed as far from digital signals as size constraints allow; and the two ideally should never closely parallel one another's paths. If they must cross, they should do so at right angles to minimize interference. Coaxial cables may be necessary for analog inputs or outputs—a demanding condition mechanically, but sometimes the only solution electrically.

If analog and digital signals must run parallel, design a ground path run between them. For signals entering and leaving a board, interpose ground pins between signal pins for added isolation.

When combining track-and-hold and a/d converter hybrids or modules on the same board, they should be mounted as near one another as practicable. All grounds need to be connected to the single, low-impedance ground plane; and the connections should be made right at the components themselves.

13.3.3 POWER SUPPLIES

Improper selection and application of power supplies can also impede high speed converter performance. Every power-supply line leading into a high speed PC card or data acquisition circuit must be carefully bypassed to its ground return to reduce noise entering the card. Ceramic capacitors, ranging in value from 0.01 to 0.1 microfarads should be used generously in the layout, mounted as closely as possible to the device or circuit being bypassed. In addition, at least one high-quality tantalum capacitor of 3 to 20 microfarads should be assigned to each power supply voltage, mounted as near as possible to the incoming power pins to keep potentially high levels of low frequency ripple off the card.

Low-noise, low-ripple, temperature-stable linear regulated power supplies are the preferred choices for high speed circuits. Switching power supplies often seem to meet those criteria. But because their ripple specifications are often expressed in terms of rms levels, the actual spikes generated in switchers (as an unavoidable by-product of the technique) may often produce hard-to-filter, uncontrollable noise peaks with amplitudes of several hundred millivolts. Their high frequency components may be extremely hard to keep out of the ground system.

If switchers cannot be avoided for high speed designs, then they should be carefully shielded and located as far away from sensitive circuits as possible. And of course, their outputs should be heavily filtered.

13.3.4 SYSTEM CONSIDERATIONS

Although analog and digital grounds should be tied together at the board, this can create problems for the system-level designer. At the system level, data converters should be considered as analog, not digital, components. The system design must be assigned to capable analog engineers who are experienced at defending millivolt signals against interference and wideband signals against degradation.

ADCs and DACs should be placed near other parts of the analog section in order to avoid transmission-line effects associated with long runs carrying high speed signals. Signal reflections and dispersion on long runs, for instance, can significantly reduce bandwidth and amplitude. In addition, if not adequately isolated from sensitive analog boards, digital subsystems can couple noise into analog circuits via the ground plane or power supplies, or via direct radiation to nearby analog components.

Each card in the system should be returned directly to the power supply common, using heavy-gauge wire. Where system considerations impose separate analog and digital board-level grounds, each should be returned to the power supply individually.

13.4 TESTING VIDEO CONVERTERS

In addition to being characterized by specifications common to general-purpose converters, such as linearity, temperature coefficients, etc., video converters require further characterization in terms that are heavily application oriented. Since the devices must be used at very high bandwidths in widely disparate applications, a diverse range of application-oriented specifications and testing methods have evolved.

Many tests of high speed converter circuits fail, due to improper use of test equipment at high speeds. To work with an oscilloscope at video speeds, for instance, requires special care, particularly when measuring settling time, troubleshooting a noise problem, or making certain other critical measurements. Oscilloscope ground clip-leads can lead to erroneous readings due to capacitive/inductive pick-up. Bayonet-type adapters, or BNC connectors should be used on or instead of probe tips. Some measurements may even require a special test fixture that plugs in directly to the scope without the use of cables. Component layout and ground considerations should also be carefully examined, as discussed in the previous section.

13.4.1 ADC TESTING

At low speeds and resolutions, a/d converters can be tested by operating in cascade ("back-to-back") with a high-speed, high-resolution d/a converter and using the reconstructed analog signal for the measurement. For testing high-resolution converters at high speeds, however, this introduces significant d/a-converter errors into the test loop. Most dynamic high speed a/d converter testing is today done using computer techniques. In typical tests, an analog signal is applied to the input of the a/d converter, while the converter's digital output is stored in a buffer memory. The digital samples are then time-weighted and a Fast Fourier Transform (FFT) is computed. From this basic data, various performance parameters can be calculated.

ADC Signal-to-Noise Ratio

As the name implies, this parameter measures the ratio of signal to noise at the output of the a/d converter. An a/d does generate some internal noise. This noise floor can create difficulties, particularly in high-resolution ADCs, where it can reduce resolution and dynamic range.

Figure 13.23 illustrates a typical test for signal-to-noise. A single frequency full-scale sinewave is applied to the a/d converter. The results are loaded into buffer memory, the samples are time weighted, an FFT is computed, and the

sinewave power is computed. Then a band-stop filter at the sinewave frequency is switched in, removing the fundamental signal; and the total remaining power is computed. The ratio is then computed and converted to dB; the result is the S/N ratio. In an ideal a/d converter with a full-scale sinewave input, the theoretical rms signal-to-noise ratio is $(6n + 1.8)$ dB, where n is the number of bits.




Figure 13.23. Signal-to-noise ratio test.

AC Linearity

In an ideal a/d converter, a pure sine wave on the analog input appears at the digital output as a pure (sampled) sine wave. In real-world ADCs, however, spurious signals due to nonlinear distortion within the ADC appear in the output. These spurious signals generally are combinations of the harmonics of the fundamental—and intermodulation products. The intermodulation products are produced when the fundamental and its harmonics beat with the sampling frequency. Generally, only spurious signals that fall within the video bandwidth ($\frac{1}{2}$ the sampling rate) are considered important. For example, in a system with sample rate of 10.74 MHz (video bandwidth less than 5.37 MHz), the third harmonic of a 3 MHz sine wave appears at 9 MHz. Spurious signals are generated at 1.74 MHz (sampling rate minus third harmonic) and at 19.74 MHz (sampling rate plus third harmonic). Note that, even though the harmonics may be out-of-band, the intermodulation products, such as the 1.74 MHz signal here, may fall within the band. It is desirable for such intermodulation products to be so small as to be negligible.

AC linearity can most easily be measured by applying a single frequency sine wave to the ADC input and observing the reconstructed d/a converter output with a spectrum analyzer. This of course may introduce additional errors due

to the DAC—which is why it should be chosen to have high speed, resolution, and linearity. AC linearity specifications generally include all in-band spurious signals and are specified in dB below full scale.

Two-Tone Intermodulation Distortion

Two-tone testing is a carryover from r-f testing techniques. With inputs consisting of sine waves at two frequencies, f_a and f_b , any active device with nonlinearities will create distortion products, of order $(m + n)$, at sum and difference frequencies of $m f_a \pm n f_b$, where $m, n = 0, 1, 2, \dots$. Intermodulation terms are those for which m or n is not equal to zero. For example, the second-order terms include $(f_a + f_b)$ and $(f_a - f_b)$, and the third-order terms include $(2 f_a + f_b)$, $(2 f_a - f_b)$, $(f_a + 2 f_b)$, and $(f_a - 2 f_b)$.

Traditionally, in r-f work, the two third-order difference terms are the ones principally considered, since their frequencies are the closest to the input frequencies (hence, in-band components); the others are filtered out (see Figure 13.24). However, in baseband applications, the second-order terms may well be more significant.




Figure 13.24. Two-tone test spectrum.





Figure 13.25. Intermodulation-distortion test.

Figure 13.25 illustrates the test procedure. Two half-scale sine waves of slightly different frequencies are summed, and the resulting signal is applied to the ADC under test. The resulting data loads into memory and is time weighted. An FFT is computed and the total fundamental power is calculated. Then the power of the third order intermodulation products, and their ratios compared to the total fundamental signal power, are computed.


Noise Power

Noise power (NPR) is the measure of the spectral power of all contributed errors, such as intermodulation and harmonic distortion, in a narrow frequency-slot within the baseband of the composite signal being processed. An example of such a communication system, including the test set-up that may be used, is shown in Figure 13.26.


In such a system, the baseband signal ranges from dc to about 8 MHz, and the a/d and d/a converters typically encode at rates of about 20 megawords/second. The NPR test consists of encoding a limited band of white noise, created by a noise generator, and examining this signal at the output of the DAC, using a noise receiver. The noise generator is equipped with band-stop filters, which eliminate very narrow slots (typically 3kHz) from the transmitted frequency spectrum (13.26b). At the receiving end, the noise receiver is equipped with complementary filters to allow the receiver to examine power spectral density of the noise contributed by the transmission medium (including the ADC and DAC) within these ideally noiseless slots (13.26c).



a. Typical test setup.



b. Transmitted spectrum, with slot.



c. Received Spectrum, slot partly filled with garbage.

Figure 13.26. Noise power ratio test setup.

This noise, the total cumulative effect of all transmission and encoding errors, such as intermodulation and harmonic-distortion products, aperture errors, and the like, is displayed as a weighted ratio of the output noise found in the slot to the power of the total transmitted noise spectrum. This number, expressed in dB, is called noise-power ratio—the larger its magnitude, the better. Noise-power ratio may also be computed in digital form, using fast Fourier transforms and/or down-conversion. Figure 13.27 diagrams a typical noise power ratio test procedure using FFTs.

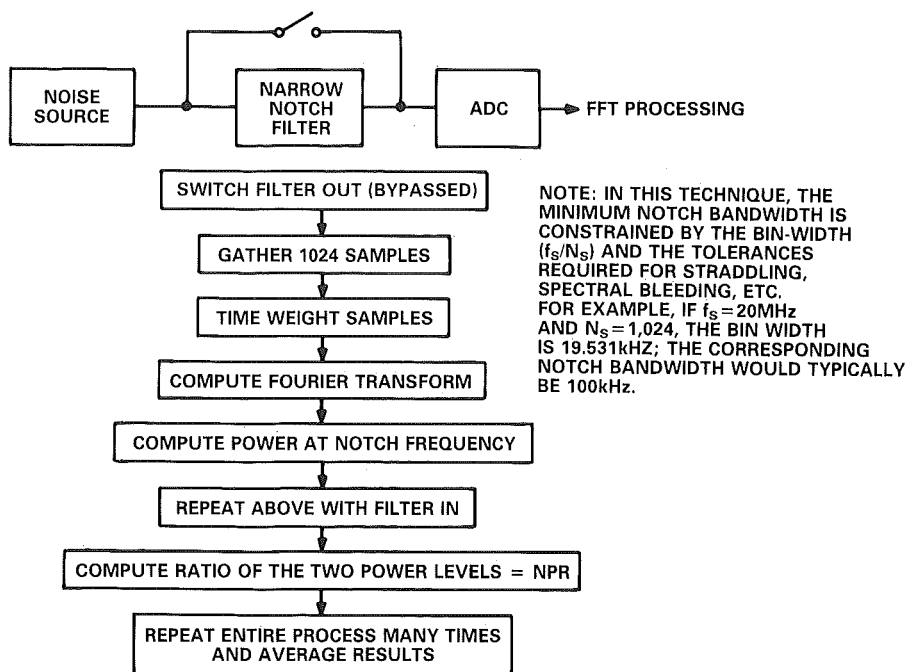


Figure 13.27. Noise-power ratio test.

Transient Response

As in the case of low-frequency converters, transient response is defined in terms of the time necessary for an ADC to achieve its rated accuracy after a full scale step function is applied to its input. Computerized testing does not readily lend itself to this test, so it is usually done manually, as illustrated by Figure 13.28.

In Figure 13.28, an encode generator is synchronized to a flat pulse generator. The flat pulse generator is adjusted for an output of just under full scale. A good high-frequency dual-trace scope is connected as shown and synchronized to the generators. The encode pulse is moved in time in relation to the flat pulse. At the setting for which the bit display indicates that the output is within 1 LSB of the final value, the amount of time from transition of the step function to the leading edge of the encode command is the transient-response settling time.




Figure 13.28. Transient-response test setup.

Differential Gain and Phase

This specification is performed specifically for video and television applications. Differential gain is defined as the percentage difference between the output amplitudes of a small high-frequency sine wave at two stated levels of a low frequency signal on which it is superimposed. Differential phase is the difference in the output phase of a small high frequency sine wave for two stated levels of a low frequency signal on which it is superimposed.

Figure 13.29 shows a typical test set-up. Specialized test equipment is used for these tests and a color video monitor is used for a final visual check of the performance of the ADC.




Figure 13.29. Differential-gain and differential-phase test.

13.4.2 DAC TESTING

Many of the testing methods that apply to low-speed DACs also apply to high-speed components. Measuring settling time of high-speed DACs, however, presents an especial challenge.

The high vertical gain required to determine accurately the *full-scale settling time* of fast d/a converters by observations during during the final settling

period tends to overload the oscilloscope during the initial phase of the response, driving it into a state that will give erroneous readings. This results in serious degradation of measurement accuracy. Therefore the preferred, although less direct, approach is to measure rise time and then calculate the settling time, or to determine settling time by observing the transition time between two voltage-comparator windows. In both cases, the oscilloscope is teamed with a digital driver (to drive the inputs of the d/a converter) and a flat-pulse generator. A window comparator circuit may also be necessary.

Calculating Settling Time

Because of the high vertical gain required to measure the full-scale settling time of an 8-nanosecond 8-bit d/a converter to within 0.4% of full scale accuracy, overloading develops and the results obtained from the oscilloscope are incorrect. However, the gain can be reduced, and the settling time can then (in many cases) be calculated from a rise-time measurement.

The rise time (10-90% points) of the all-ZEROs to all-ONES transition, t_r , is measured by means of a real-time or sampling oscilloscope that has a bandwidth of at least 500 MHz. In a single-pole system, the associated time constant, T , is related to t_r by $t_r = 2.2 T$. If an exponential function is assumed, then the d/a converter output, V , is calculated from the equation,

$$V = (1 - e^{-t/T}) V_o \quad (13.3)$$

where V_o is the final, settled output.

Table 13-2 lists the settling time required to reach a percentage of the final output, V_o , for various rise times, t_r .

Settling times to reach $\pm 1/2\text{LSB}$ of output	
Resolution ($\pm 1/2\text{LSB}$)	Settling time (t_s)
12 bits	$3.8 t_r^*$
11 bits	$3.5 t_r$
10 bits	$3.1 t_r$
9 bits	$2.8 t_r$
8 bits	$2.5 t_r$
7 bits	$2.2 t_r$
6 bits	$1.9 t_r$
5 bits	$1.5 t_r$
4 bits	$1.2 t_r$

* t_r = rise time for input step
from all-ZEROS to all-ONES
(10% to 90%)

Table 13.2. Settling time vs. resolution for single time-constant exponentials.

Accurate *midscale bit-transition settling time* (as opposed to full-scale) is relatively easy to measure using direct oscilloscopic techniques. Since an interval of only 1 LSB is being monitored, the response time is generally quite a bit

shorter, and the glitch at the midscale carry point (011 . . 11 to 100 . . 00) usually has an amplitude of less than 100 mV—maintained very briefly—little danger exists of excessively overdriving the input amplifier of the oscilloscope.

Window Measurement of DAC Settling Time

An alternative way to measure full-scale settling time (from the start of the transition) is to employ fast window comparators to establish T_1 , when the response has started—i.e., exceeds 1 LSB—and when it is within 1 LSB of the final value (T_2). The oscilloscope is used only to observe the binary switching of logic levels rather than a graded analog response. Figure 13.30 shows a typical measurement setup.

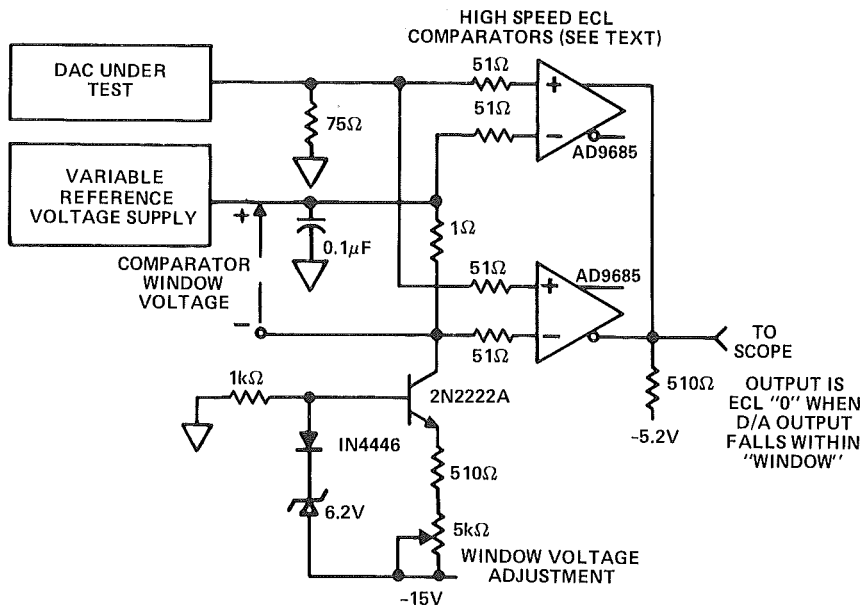


Figure 13.30. Settling-time test setup employing comparators.

To measure full scale settling time, T_1 and T_2 must be measured. The window is set for 1 LSB (for instance, 2.5 mV for the HDG-0805 8-bit display DAC). For T_1 , the reference is set at the most positive value that will still leave the comparator output at logic 0 during the time immediately preceding the output transition of the DAC. T_1 is defined as the time the fast comparator output reaches 50% of its transition following the start of the DAC output transition (Figure 13.31).

For T_2 , the comparator reference is shifted positively until the comparator output returns to logic 0 during the time following the DAC's output transition, and the 50% point of the 1 to 0 transition is the shortest possible time after T_1 . $T_2 - T_1$ represents the settling time, as defined above, independently of the fixed delay through the DAC input registers.

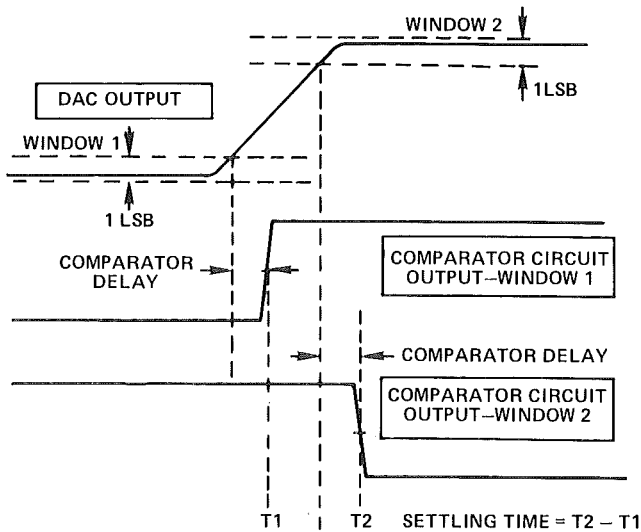


Figure 13.31. Response plots in settling-time test.

Since T_1 is established with both comparators in a state of large overdrive, and T_2 with the comparators in small overdrive, it is desirable to consult the specifications for the specific comparators being used to see if the difference in comparator response is sufficiently significant to require a correction, and to establish the size of the correction. Figure 13.32 shows the comparator waveforms for T_1 and T_2 superimposed in a double exposure. The settling time measures about 7.5 nanoseconds.




Figure 13.32. Waveforms in settling-time test.

Chapter Fourteen

Converters for Resolvers and Related Devices


The conventional approach to signal conditioning and data acquisition (Chapters 2 and 4) is useful in the many cases where a variety of transducers, chosen for a variety of reasons, must have their outputs interpreted or their inputs furnished digitally. However, there are a great many applications in position measurement where a simple subsystem, consisting of a specialized electromagnetic sensor having standardized analog data format and a special-purpose microprocessor-compatible data converter, provides accurate, reliable, cost-effective measurement and control. This chapter is about some subsystems of this kind.

The boom in robotics, computer-aided manufacturing, and factory automation has generated increasing demands for linear and rotational position-measuring transducers that are rugged, reliable, and easily automated. Resolvers, synchros, Inductosyns^{TM*} (flat-transformer-like position transducers), and LVDTs (linear variable differential transformers) have been available to meet these needs for years. Until recently, however, they have found only limited use in digital systems, primarily because the interfacing was somewhat complicated and expensive.

The situation has now changed considerably. Integrated-circuit manufacturers, such as Analog Devices, offer compact digital interface circuits, such as resolver-to-digital and synchro-to-digital converters (Figure 14.1a), in the form of hybrid integrated circuits. These devices make it significantly easier to use resolvers, synchros, Inductosyns, and—more recently—LVDTs in automated position-measurement applications. Data-acquisition systems that

*TM—Inductosyn is a registered trademark of Farrand Controls, Inc.

team these converters up with the respective transducers achieve superior reliability, noise performance, accuracy, and survivability in harsh environments when compared with other commonly used position-transducer alternatives, such as optical encoders and potentiometers.



(Courtesy of Moore-Reed, Ltd.)

Figure 14.1. Resolvers and resolver-to-digital converters.

14.1 POSITION MEASUREMENT IN PERSPECTIVE

Modern machine-tool and robot systems increasingly rely on precision position transducers as the primary feedback element in their closed-loop control systems. The robots of Figure 14.2, for instance, have nineteen angular movements, each of which must be measured and controlled. Not only do resolvers, in combination with resolver-to-digital converters (RDCs), provide absolute digital representations of angular position; they also offer the option of an analog voltage output proportional to angular velocity. Designers of robot and machine-tool controllers find this signal indispensable for control-loop stabilization, or speed-profile control.

Resolvers, synchros, Inductosyns, and their associated digital converters have already found use in a wide range of applications, and the list is growing. A brief listing might include:

Avionics:

- airborne radar scanners
- aircraft control surface feedback (flap control, for instance)
- acquisition of pitch, roll, heading data for aircraft navigational equipment
- sonars

Industrial measurement and control, including:

- robots
- machine-tool controllers
- speed controllers
- control of printing machinery




Figure 14.2. Multi-axis robot arms.

Military:


missile gyros and launchers
 naval fire control systems, data retransmission units, and compass
 repeaters
 radars
 torpedos and torpedo launchers
 electronic countermeasures equipment
 tank gun sights and gun controls
 ground-based missile launchers

Of available position sensors, most—such as resolvers, synchros, rotary inductosyns, potentiometers, and optical encoders—are inherently *rotary* measuring devices. Through the use of lead screws, however, such rotary transducers are readily converted into *linear-position* transducers. In contrast, linear Inductosyns, linear variable differential transformers (LVDTs), and straight wire potentiometers measure linear positions directly. Position sensors can be further categorized as providing either *absolute* or *incremental* measurements.


14.1.1 OPTICAL ENCODERS

Among the most-popular position measuring transducers, optical encoders find use mainly in relatively low-reliability and low-resolution applications. An *incremental* optical encoder (Figure 14.3a) is a disc divided into sectors that are alternately transparent and opaque. A light source is positioned at one side of the disc, and a light detector at the other side. As the disc rotates, the output from the detector switches alternately on and off depending on whether the sector appearing between the light source and the detector is opaque or trans-

parent. Thus, the encoder produces a stream of square wave pulses which, when counted, indicate the angular position of the shaft. Available encoder resolutions (the number of opaque and transparent sectors per disc) range from 100 to 65,000, with absolute accuracies approaching 30 arc seconds ($1/43,200$ per rotation).



a. Incremental encoder.



b. Absolute encoder.

Figure 14.3. Optical encoders.

Most incremental encoders feature a second light source and detector, at an angle to the main source and detector, to indicate direction of rotation. Many encoders also feature a third light source and detector to sense a once-per-revolution marker. Without some form of revolution marker, absolute angles are difficult to determine.

As a potentially serious disadvantage, incremental encoders require external counters to determine absolute angles within a given rotation. If the power is momentarily shut off, or if the encoder misses a pulse due to induced noise or a dirty encoder disk, the resulting absolute angular indications will be in error. For robotics applications, in particular, reliable determination of absolute angular position is increasingly required. The absolute position of a robot's arm is critical in most applications, for instance, welding or assembly.

Absolute optical encoders overcome these disadvantages but typically cost more than twice as much as their incremental counterparts. An absolute optical encoder's disc is divided up into 2^N sectors (Figure 14.3b). But each sector is further divided radially along its length into opaque and transparent sections, forming a unique N-bit digital word with a maximum count of 2^N . The digital word formed radially by each sector increments in value from one sector to the next, usually employing Gray code (see Table 7.6 and Figure 7.5).

A set of N light detectors, arranged radially on one side of the disc, responds to N light sources positioned on the other side. The detectors' output forms an N-bit digital word corresponding to the disc's absolute angular position. Industrial absolute optical encoders achieve up to 16-bit resolutions, with absolute accuracies that approach the resolution (20 arc seconds).

Both absolute and incremental optical encoders, however, experience difficulties, particularly in harsh factory environments. When subjected to vibration or sudden shocks, they may fall out of optical alignment. Light-source lifetimes can be seriously temperature limited, and incandescent light sources can experience damage through vibration. The gallium-arsenide light sources used in many optical encoders may suffer degraded performance due to aging and extremes of temperature. For this reason, designers often hesitate to mount optical encoders where they are needed most, such as directly on drive motors that operate in harsh environments. For low-resolution (10-bits or less) applications not requiring the utmost in reliability, however, optical encoders are easy to use and low in cost.

14.1.2 POTENTIOMETERS, CONTACT ENCODERS, RVDTs, LVDTs, AND GLASS SCALES

As common position-measuring devices, *potentiometers* exhibit exceptionally low cost but suffer from significantly poorer reliability when compared to other transducers. Since typical pots employ a brush or wiper contacting a resistive track, they often experience wear-related electrical and mechanical noise problems. Their performance depends upon the stability of their reference voltage supply, and the accuracy of the signal conditioning circuitry and associated analog-to-digital converter. Designers occasionally use potentiometers for coarse position measurement, in conjunction with other transducers. Available potentiometers approach 12-bit resolutions, with absolute accuracies approaching 7 arc-minutes (about $1/3000$ per rotation).

So-called *brush* or *contact encoders* work on a principle similar to that of optical encoders. However, instead of using optics to generate the angular signal, contact encoders employ conducting and insulating surfaces in conjunction with brush contacts. This type of encoder cannot sustain high rotational speeds and remains limited to applications requiring small and infrequent shaft movements. Contact encoders with resolutions of up to 10 bits and accuracies of 26 arc-minutes (about $\frac{1}{830}$ per rotation) are available.

Rotary (or linear) variable differential transformers (RVDTs, LVDTs) are in common use. LVDTs, for instance, find use in metrology, gauging, and even cash dispensers. They feature contactless operation, small size, and high resolution and repeatability. Some devices do suffer from drift as high as 0.08%/°C. RVDTs find use extensively in avionic applications, such as vane-position indication in turbines, aircraft angle-of-attack indicators, and control-surface feedback. RVDTs suffer from a limited rotational range (85 degrees) and limited accuracies (1 to 2%).

Direct interfacing of LVDTs and RVDTs to digital systems has recently become practical through the use of LVDT-to-digital and RVDT-to-digital converters, such as those in the Analog Devices 2S56 series. They have many features in common with resolver-to-digital converters.

Glass-scale or Moire-fringe transducers represent still another transducer technology. Glass scales are widely used for high-accuracy linear and rotational measurement. A linear scale for example will give a basic resolution of 10 microns, which can often be improved by a 4- or 5-bit interpolation method. Glass scales are significantly less expensive than some alternative transducer technologies, but lack ruggedness.

14.1.3 RESOLVERS AND SYNCHROS

Because of the serious disadvantages of competing transducer technologies, machine-tool and robotics manufacturers increasingly turn to resolvers and Inductosyns as the position transducers of choice. These components particularly excel in demanding factory applications requiring small size, long-term reliability, absolute position measurement, high accuracy, and low-noise operation.

Figure 14.4 is a diagram of a typical brushless resolver. In appearance, synchros and resolvers resemble small cylindrical AC motors. They vary in diameter from 0.8 inches to 3.7 inches. One end of their bodies has an insulated terminal block, and the other end a mounting flange. Rotor shafts are normally threaded and splined.

In their simplest forms both synchros and resolvers employ single-winding rotors that revolve inside fixed stators. In the case of a simple *synchro*, the stator has three windings oriented 120 degrees apart and electrically connected in a Y-configuration. Resolvers differ from synchros in that their




Figure 14.4. Synchro and resolver windings. Brushless resolver construction.

stators have only two windings oriented at 90 degrees.

Because synchros have three stator coils in a 120° orientation, they are more difficult than resolvers to manufacture and are therefore more costly. Today, synchros find decreasing use, except in certain military and avionic retrofit applications.

Modern *resolvers*, in contrast, are available in a brushless form that employs a rotating transformer on the rotor to couple the rotor signals. Because brushless resolvers have no slip rings or brushes, and because there are only two stator coils wound at right angles, they are cheaper and easier to make compared to synchros. They are also more rugged than synchros because there are no brushes to break or dislodge. In fact, the life of a brushless resolver is limited only by its bearings. Furthermore, most brushless resolvers are specified to work over 2-40 volts and at frequencies from 400 to 10,000 Hz. This makes them easier to tailor to the user's needs.

Typical resolvers achieve the following accuracies:

Standard Size 11:	± 5 arc-minutes
Selected Size 11:	± 3 arc-minutes
High Accuracy Slab Type:	± 1 arc-minutes
Two-speed electrically geared type:	0.5 arc-minutes

In operation, synchros and resolvers resemble rotating transformers. The rotor winding is excited by an AC reference voltage, at frequencies up to a few kilohertz. The magnitude of the voltage induced in any stator winding is proportional to the cosine of the angle, θ , between the rotor coil axis and the stator coil axis. In the case of a synchro (Figure 14.4), the voltage induced across any pair of stator terminals will be the vector sum of the voltages across the two connected coils.

For example, if the rotor of a synchro is excited with a reference voltage, $V \sin \omega t$, across its terminals, R1 and R2, then the stator's terminals will see voltages of the form

$$\begin{aligned} V(S1 \text{ to } S3) &= V \sin \omega t \sin \theta \\ V(S3 \text{ to } S2) &= V \sin \omega t \sin (\theta + 120^\circ) \\ V(S2 \text{ to } S1) &= V \sin \omega t \sin (\theta + 240^\circ) \end{aligned} \quad (14.1)$$

where θ is the shaft angle.

In the case of a resolver, with a rotor a-c reference voltage of $V \sin \omega t$, the stator's terminal voltages will be

$$\begin{aligned} V(S1 \text{ to } S3) &= V \sin \omega t \sin \theta \\ V(S4 \text{ to } S2) &= V \sin \omega t \sin (\theta + 90^\circ) \\ &= V \sin \omega t \cos \theta \end{aligned} \quad (14.2)$$

Resolver-to-digital converters (RDCs) transform these voltage relationships into digital representations of the actual angle, θ .

When teamed up with such converters, synchros and resolvers easily achieve 12-bit resolutions, corresponding to an angular resolution of 5.3 arc-minutes. Absolute accuracies can approach 20 arc seconds, or better, for models having higher resolution. If the resolver, connected to a 0.1-inch-pitch leadscrew, is used with a 12-bit converter, linear resolution becomes 25 microinches.

Resolvers and synchros offer the advantages of proven long-term reliability, absolute position measurement, high accuracy and resolution, small size, and moderate cost. In fact, a standard resolver in combination with an RDC costs less than many optical encoders, especially in the case of optical encoders specifically designed for use in harsh industrial environments. In addition, users report that typical resolvers and synchros have a mean time before failure (MTBF) four or more times longer than that for optical encoders.

In the past, resolvers and synchros were handicapped by the need for complicated, expensive interface circuitry. But with the availability of easy-to-implement RDCs and SDCs in hybrid form, this roadblock no longer exists.

14.1.4 INDUCTOSYNS

Synchros and resolvers inherently measure rotary position, but they can make linear position measurements when used with lead screws. An alternative, the Inductosyn* (Figure 14.5) measures linear position directly. In addition, Inductosyns are accurate and rugged, well-suited to severe industrial environments, and do not require ohmic contact.




Figure 14.5. Inductosyn, showing detail of slider pattern.

The linear Inductosyn consists of two magnetically coupled parts; it resembles a multipole resolver in its operation. One part, the scale, is fixed (e.g. with epoxy) to one axis, such as a machine tool bed. The other part, the slider, moves along the scale in conjunction with the device to be positioned (for example, the machine tool carrier).

The scale is constructed of a base material such as steel, stainless steel, aluminum, or a tape of spring steel, covered by an insulating layer. Bonded to this is a printed-circuit track, in the form of a continuous rectangular waveform pattern. The pattern typically has a cyclic pitch of 0.1 inch, 0.2 inch, or 2 millimeters. The slider, about 4 inches long, has two separate but identical printed circuit tracks bonded to the surface that faces the scale. These two tracks have a waveform pattern with exactly the same cyclic pitch as the waveform on the scale, but one track is shifted one-quarter of a cycle relative to the other. The slider and scale remain separated by a small air gap of about 0.007 inch.

Inductosyn operation resembles that of a resolver. When the scale is energized with a sine wave, this voltage couples to the two slider windings, inducing voltages proportional to the sine and cosine of the slider's spacing within a cyclic pitch. If S is the distance between pitches, X is the slider displacement within a pitch, and the scale is energized by the voltage, $V \sin \omega t$, then the slider windings will see terminal voltages of

*Inductosyns are manufactured by Farrand Controls, Inc., Valhalla, NY, and their licensees. Inductosyn® is a registered trademark of Farrand Controls, Inc.

$$V(\text{sine output}) = V \sin \omega t \sin \frac{2\pi X}{S}$$

$$V(\text{cosine output}) = V \sin \omega t \cos \frac{2\pi X}{S} \quad (14.3)$$

As the slider moves a distance equivalent to one pitch, the voltages produced by the two slider windings are similar to those produced by a resolver rotating through 360 degrees. However, in contrast to a resolver's highly efficient transformation ratio of 1:1 or 2:1, typical Inductosyns operate with transformation ratios of 100:1. This results in a pair of sinusoidal output signals in the millivolt range which generally require amplification. Dual preamplifiers are available to amplify both signals in identical fashion (see Figure 14.14).

Since the slider output signals are derived from an average of several spatial cycles, small residual errors in conductor spacing have minimal effects. This is an important reason for the Inductosyn's very high accuracy. In combination with 12-bit Inductosyn-to-digital converters (IDCs), 0.1-inch-pitch linear Inductosyns readily achieve 25 microinch resolutions.

Rotary Inductosyns can be created by printing the scale on a circular rotor and the slider's track pattern on a circular stator. Such rotary devices can achieve very high resolutions. For instance, a typical rotary Inductosyn may have 360 cyclic pitches per rotation, and might use a 12-bit Inductosyn-to-digital converter. The converter effectively divides each pitch into 2^{12} , or 4,096, sectors. Multiplying by 360 pitches, the rotary Inductosyn divides the circle into a total of 1,474,560 sectors. This corresponds to an angular resolution of less than 0.9 arc seconds. Higher-resolution models can achieve resolutions to 0.05 arc seconds (about 25-bit resolution), with absolute accuracies approaching 1.5 arc-seconds.

In the above example, the absolute orientation of the Inductosyn is determined by counting successive pitches in either direction from an established starting point. Although the basic Inductosyn is a single-track incremental device, absolute distance measurements can also be made by combining two tracks. The most common device is often called an "N and N - 1" Inductosyn. The device employs two tracks with, say, 255 and 256 pitches. A separate converter on each track and a small amount of digital logic results in absolute position measurements at all times *including turn-on*. Schemes such as this provide the most cost-effective high-resolution system available, typically achieving resolutions comparable to those of absolute optical encoders (for instance) at less than one fourth the cost.

14.1.5 COMPARING ALTERNATIVES

Figure 14.6 compares in a general way the absolute accuracies and costs of the competing position-measuring transducers. In addition to these parameters, resolution is an important consideration. The resolution necessary for a given application, however, largely depends upon the smoothness required

for the finished workpiece and/or the number of bits of digital resolution specified for input to the computing system. Today, 10-bit resolutions are most common (1,024 increments per revolution), but 12-bit (4,096 increments) and 14-bit (16,384 increments) resolutions are becoming more popular, particularly in advanced machine-tool controllers and drive systems. The designer must carefully match the system resolution to the absolute accuracy of the transducer used. It makes no sense to use a 1%-accuracy potentiometer with a 14-bit control system.

Repeatability often has greater significance than absolute accuracy, particularly in robot or tracer applications. Repeatability should at least match the resolution of the control system. Repeatability is particularly important in "teaching-type" robots where it is crucially important that the robot return to the same instructed positions each time.




Figure 14.6. Relative cost vs. accuracy of implementing common position transducers digitally.

14.2 RESOLVER-TO-DIGITAL CONVERTERS

14.2.1 TRACKING CONVERTERS

The attractiveness of resolvers, synchros and Inductosyns is significantly enhanced by the availability of easy-to-implement interface circuits, for example, tracking resolver-to-digital converters (RDCs), such as those pictured in Figure 14.1.

Figure 14.7 is a functional diagram of a typical tracking resolver to digital converter. To begin with, it is worth noting that RDCs and Inductosyn-to-digital

being measured. The resolver's stator output voltages, from equation 14.2, are written as

$$\begin{aligned} V_1 &= V \sin \omega t \sin \theta \\ V_2 &= V \sin \omega t \cos \theta \end{aligned} \quad (14.4)$$

where θ is the angle of the resolver's rotor. The digital angle ϕ is applied to the cosine multiplier, and its cosine is multiplied by V_1 to produce the term,

$$V \sin \omega t \sin \theta \cos \phi \quad (14.5)$$

Digital angle ϕ is also applied to the sine multiplier and multiplied by V_2 to produce

$$V \sin \omega t \cos \theta \sin \phi \quad (14.6)$$

These two signals are subtracted by the error amplifier to yield an AC error signal of the form

$$V \sin \omega t (\sin \theta \cos \phi - \cos \theta \sin \phi) \quad (14.7)$$

This, in turn, through a simple trigonometric identity, reduces to

$$V \sin \omega t \sin (\theta - \phi) \quad (14.8)$$

The phase-sensitive detector demodulates this AC error signal, using the resolver's rotor voltage as a reference. This results in a DC error signal proportional to $\sin (\theta - \phi)$.

The DC error signal feeds an integrator, the output of which drives a voltage controlled oscillator (VCO). The VCO, in turn, causes the up-down counter to count in the proper direction to cause

$$\sin (\theta - \phi) \rightarrow 0 \quad (14.9)$$

When this is achieved

$$\theta - \phi \rightarrow 0 \quad (14.10)$$

Therefore,

$$\phi = \theta \quad (14.11)$$

to within one count. Hence the counter's digital output, ϕ , represents the angle θ . The latches enable this data to be transferred without interrupting the loop's tracking.

The circuit of Figure 14.7 is equivalent to a so-called type-2 servo loop, because it has, in effect, two integrators. One is the counter, which accumulates pulses; the other is the frequency-shaping filter at the output of the phase-sensitive detector. In a type-2 servo loop, with a constant rotational velocity input, the output digital word continuously follows, or tracks the input, without needing externally derived convert commands, and with no steady state phase lag between the digital output word and actual shaft angle. An error signal appears only during periods of acceleration or deceleration.

In contrast, since a type-1 servo loop has only one stage of integration, its error signal must remain non-zero for constant rotational velocities. In a type-1 loop, the digital angle output therefore lags the resolver's actual angular position by an amount proportional to the required error signal. As a consequence, type-1 loops are hardly ever used.

Here is a heuristic explanation of how this works: In Figure 14.7, if the resolver shaft is rotating at a constant rate, then the output digital word must follow at a constant rate, the counter must be counting pulses at a constant rate, and they can be produced at a constant rate only if the voltage at the input to the VCO (voltage-controlled oscillator) is constant and nonzero.

Therefore, for a constant rotational velocity, the second integrator (the frequency-shaping network) must maintain a constant nonzero dc output voltage to drive the VCO. This can only occur when its input, the loop error signal, remains at zero volts (since if it were not zero, the output of the integrator would have to be changing). Consequently, a type-2 servo loop enforces zero steady-state error signal (i.e., $\phi = \theta$) when the shaft is rotating at constant velocity. In other words, the digital output signal follows the resolver's actual shaft angle.

As an added bonus, the tracking RDC of Figure 14.7 provides an analog dc output voltage directly proportional to velocity. This is a useful feature if velocity is to be measured or used as a stabilization term in a servo system—it makes tachometers unnecessary.

Since the operation of these converters depends only on the ratio between input signal amplitudes, attenuation in the lines connecting them to resolvers doesn't substantially affect performance. For similar reasons, these converters are not greatly susceptible to waveform distortion. In fact, they can operate with as much as 10% harmonic distortion on the input signals; some applications actually use *square-wave* references with little additional error.

Because the tracking converter doubly integrates its error signal, the device offers a high degree of noise immunity (40 dB-per-decade rolloff). The net area under any given noise spike produces an error. But typical inductively coupled noise spikes have equal positive and negative going waveforms. When integrated, this results in a zero net error signal. The resulting noise immunity, combined with the converter's insensitivity to voltage drops, lets the user locate the converter at a considerable distance from the resolver. Noise rejection is further enhanced by the phase detector's rejection of any signal not at the reference frequency, such as wideband noise.

The family of converters exemplified in Figure 14.7 will handle resolver or Inductosyn signals at input frequencies in the range 400-10,000 Hz. The converter's inputs are designed to accept voltages of 2 volts rms, which is the lowest reference voltage likely to be encountered. Higher voltages can be scaled down as needed. Consequently the converters can interface to just about any

resolver or Inductosyn - and can produce digital output words with resolutions (depending on which family member) of 10, 12, 14, or 16 bits.

14.2.2 PHASE-ANALOG APPROACH

Some users of resolvers, synchros and Inductosyns choose not to use proprietary converters, such as the one illustrated in Figure 14.7. Instead, they seek economies by designing their own interface circuitry. The so-called phase-analog method represents the most common approach.

In contrast with the tracking converter method, the phase-analog converter excites the resolver's stator (instead of rotor) with reference signals $V \sin \omega t$ and $V \cos \omega t$. The rotor winding then sees an induced voltage of the form $V \cos(\omega t + \theta)$, where θ is the shaft angle. Controlled by zero-crossing detectors, a pulse generator and counter together time the phase delay between rotor and stator signals. This delay is, of course, directly proportional to the rotor shaft angle θ .

While this approach is less costly for large numbers of channels, it is noise-sensitive and not capable, in practice, of the accuracy with which theory blesses it. Here are some points to consider, when comparing the phase analog approach with commercially available tracking-type converters. Tracking converters have:

- High tracking speed and no velocity error or lag.
- Higher accuracy. Whereas phase analog converters are capable of high resolution, depending upon clock rate, the technique has difficulty achieving accuracies beyond 10 or 11 bits in practice.
- Better immunity to variations in resolver signal and noise. The tracking converter concept virtually eliminates problems with signal level variations and noise. Phase-analog converters, on the other hand, produce large errors when the reference source is unstable, and require filters on the resolver rotor when using non-sine wave references.
- Cost. Here, the phase analog approach may have an advantage today in multichannel systems (more than about 5 channels) if accuracy and noise considerations are of less importance. Even so, new low-cost hybrid converters (and the monolithics that are likely to appear in the future) will extend their cost effectiveness to increasing numbers of channels.
- *Perhaps most important*, the accuracy of the phase analog approach is highly dependent upon the ratio of the sine and cosine drive-signal amplitudes, as well as the actual phase difference between those signals. Figure 14.8 plots the number of bits of achievable accuracy as a function of the drive signal amplitude and phase relationships. If, for instance, the drive signals have a 1% amplitude mismatch (an amplitude ratio of 1.01) and deviate from a 90° phase relationship by only 0.5° (corresponding to the $\beta = 89.5^\circ$ curve in Figure 14.8), then maximum achievable accuracy is less than 10 bits.




Figure 14.8. Achievable accuracy level vs. amplitude ratio and drive-signal phase error ($= |\beta - 90^\circ|$) in a phase-analog conversion system.

This underscores the need for analog design expertise in implementing the phase-analog approach.

14.2.3 SUCCESSIVE-APPROXIMATION APPROACH

The successive-approximation conversion technique has been discussed in Chapter 7 and elsewhere. Although widely used for general-purpose a/d conversion and once popular for synchro-to-digital conversion, this technique is not often used in commercial SDC designs. Nonetheless, it warrants discussion, since elements of the technique are often adopted for in-house designs, especially in multi-channel applications.

In one approach, the successive approximation converter accepts sine and cosine inputs (via Scott-T transformers if the information is in synchro form), converts them to digital, and finds the angle via a tangent lookup or an algorithm. Thus, the loop does not act in a continuous tracking mode; each conversion is independent and requires a convert command signal to initiate it. Noise represents a serious problem for this approach. Any transient appearing on the sine or cosine lines as they are being converted will cause large errors. For this reason, many designs employing this technique take several readings and produce an average.

The inputs to the converter are two dc voltages representing $V \sin \theta$ and $V \cos \theta$, where θ is the input angle. The converter itself cannot accept ac voltages in the synchro or resolver format; the angular information requires signal conditioning to translate it to dc.

What then is the advantage of the successive-approximation converter? It is sometimes cheaper, when a large number of channels of synchro or resolver information have to be converted into digital form, along with other analog data, such as temperature, pressure, etc.

Suppose for example there are 16 channels of slowly changing data to be processed, including eight channels of synchro data. Clearly, the best angle-measuring performance could be obtained by using 8 tracking converters, from which fresh angular data from any channel could be taken at any time by digitally multiplexing the outputs on a bus.

However, the cost of 8 tracking converters (in recent years, but not necessarily at the time you read this or in the future) might appear to make this method undesirable. As an alternative lower-cost approach, a system of Scott-T-connected transformers, peak detectors, and sample/hold amplifiers, together with a multiplexed successive approximation converter, could serve as a solution to this problem. Figure 14.9 illustrates such an approach, showing three channels.

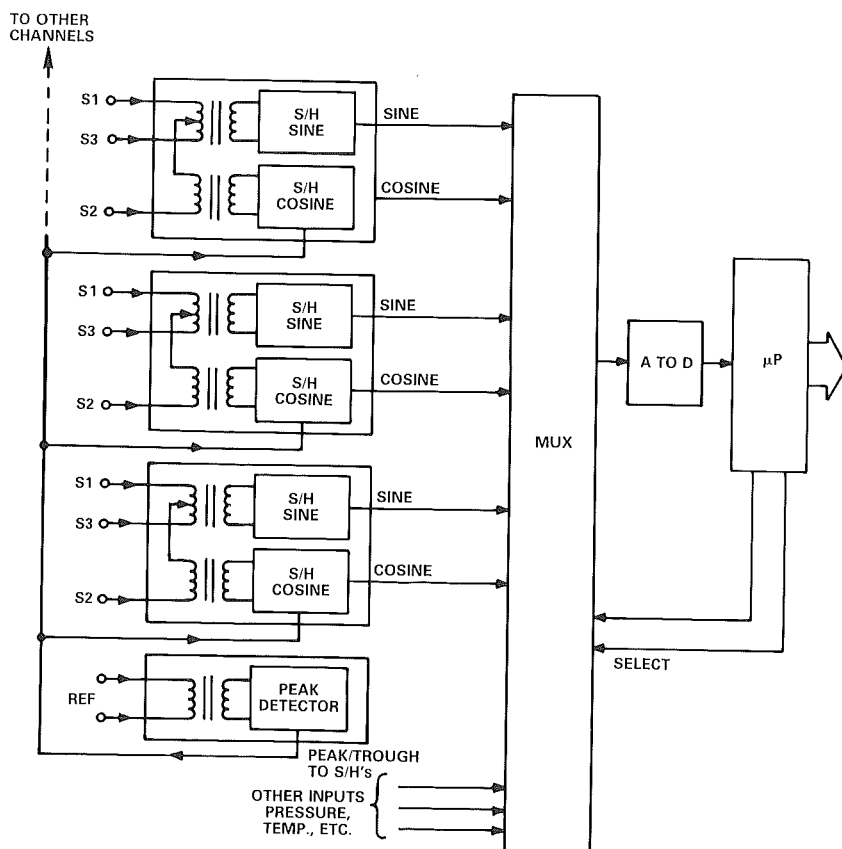


Figure 14.9. Multichannel conversion system using successive-approximation conversion.

The transformers convert from synchro to resolver format, if necessary; the peak detectors determine the amplitudes of the sine and cosine components of each angle in synchronism with the peaks of the reference waveform, and

the sample/holds store the values until they are read out via the multiplexer and converted.

In an ideal system, where no phase shift exists between the reference and signal waveforms, the resolver-format signals would be sampled at the peaks of their waveforms. In practice, small phase shifts which may be present have no effect on the accuracy of the system. When a channel is to be converted, the sine and cosine outputs of its sample/hold are applied in sequence at the inputs of the a/d converter, and conversions are initiated.

Note that sampling converters, such as this one, are more sensitive to noise on the input than are tracking converters. This is because the input is usually obtained from a single sample of the voltage, while the tracking converter's frequency-shaping circuits effectively average the waveform over many periods of the reference. For applications where the inputs are varying slowly, it is possible to use precision phase-sensitive rectification and smoothing instead of sample/hold modules, in which case the smoothing will give some immunity to noise.

14.2.4 TACHOMETER OUTPUTS

Most servo-loop designs, particularly in systems where position measurement is critical, employ some form of tachometer signal for stabilization or speed control. For this reason, it is important to note that many modern converters offer the user a tachometer output with specifications comparable to what can be achieved with top-grade electromechanical tachometers.

Therefore, if a resolver is used as a feedback element, and the RDC provides a tachometer output, the designer can eliminate an electromechanical tachometer from the system design.

14.3 DIGITAL-TO-RESOLVER CONVERTERS

Digital-to-synchro and digital-to-resolver converters (DSCs and DRCs) interface digital systems to synchro or resolver angular control systems, and to other forms of application, such as PPI radar displays. The converters take binary digital inputs, representing angles, and produce outputs in either synchro or resolver format at standard frequencies and voltages. Practical converter outputs normally include a 2 volt-ampere amplification stage, which lets the units drive many electromechanical loads directly.

The most common loads include:

- Control Transformers (CT)
- Torque Receivers (TR)
- Control Differential Transmitters (CDX)
- Torque Differential Transmitters (TDX)

Synchro or resolver control transformers, used as feedback elements in electromechanical servo control loops, represent the most common loads. Figure

14.10 shows such a control loop being driven from a digital-to-synchro converter.

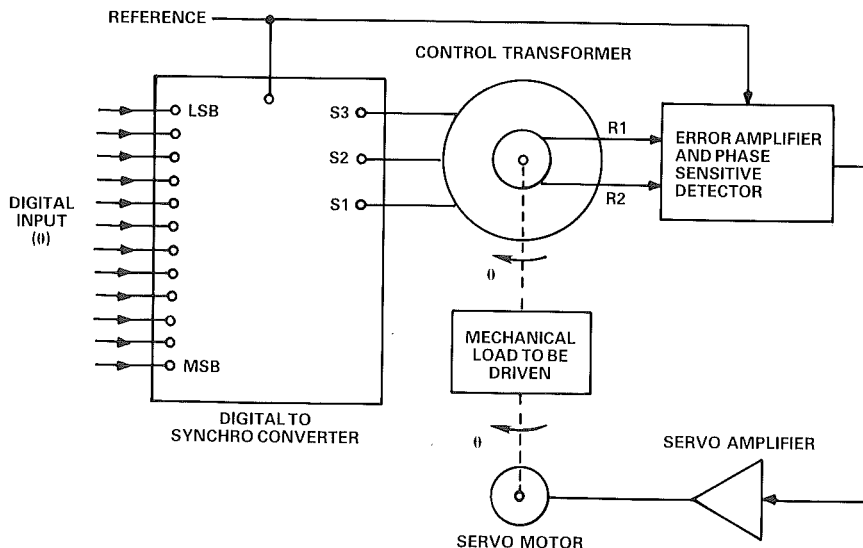


Figure 14.10. Digital-to-synchro converter in control-loop setpoint application.

Figure 14.11a illustrates the design of a digital-to-synchro converter and Figure 14.11b a digital-to-resolver converter. Differences in the secondary connections of the output transformers signify the only difference between the two designs. In the DSC, the output transformers are connected in a Scott configuration, with the outputs 120 degrees apart; the DRC's transformers are connected to provide outputs 90 degrees apart.

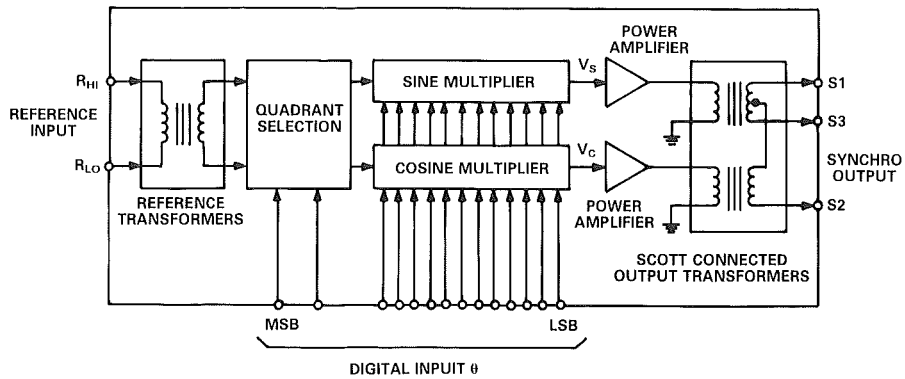
Like the tracking-type SDCs and RDCs, the converters of Figures 14.11 use multiplying digital to analog converters incorporating sine and cosine laws. The converters multiply the reference input voltage, $V \sin \omega t$ by $\sin \theta$ and $\cos \theta$ to produce the resolver-form voltages:

$$\begin{aligned} V(\text{sine}) &= E \sin \omega t \sin \theta \\ V(\text{cosine}) &= E \sin \omega t \cos \theta \end{aligned} \quad (14.12)$$


where θ is the desired shaft angle applied in digital form to the sine and cosine multipliers.

For digital-to-resolver converters, these voltages are fed directly into the power amplifiers and isolation transformers. In the case of DSCs, the voltages feed into the power amplifiers and the Scott-T-connected transformers to produce the required synchro-format voltages:

$$\begin{aligned} V_{S1-S3} &= E \sin \omega t \sin \theta \\ V_{S3-S2} &= E \sin \omega t \sin (\theta + 120^\circ) \\ V_{S2-S1} &= E \sin \omega t \sin (\theta + 240^\circ) \end{aligned} \quad (14.13)$$



a. Digital to synchro.



b. Digital to resolver.

Figure 14.11. Digital-to-synchro and digital-to-resolver converters.

In certain converter designs, the sine and cosine multipliers do not follow the sine and cosine laws exactly, but vary from them by up to $\pm 7\%$. Although the sine and cosine are not accurate *per se*, in such devices, their ratio forms a very close approximation to the tangent, and they are useful in establishing shaft angles accurately. However, the value of E , and hence the *magnitude* of the resulting vector in Equations 14.12 and 14.13, is a function of the angle θ ; the phenomenon is called *radius-vector variation*. If a digital-to-resolver converter of this type were used to drive the X and Y plates of an oscilloscope or radar display, the resulting trace—which is expected to be circular—would appear distorted, as shown in Figure 14.12, with accurate angles but erroneous amplitudes.

Modern converters (for example, the Analog Devices DRC1745) achieve low radius-vector variations, typically $\pm 0.1\%$, and consequently experience few problems of this sort, whether driving an electromechanical servo system, an oscilloscope, or a radar display. Some older converter designs, with significant radius-vector errors, can still be found driving certain electromechanical

loads that depend on the angle of the vector and are tolerant of magnitude errors. In those cases, it is still important that the *ratio* of sine and cosine, i.e., the tangent, retain sufficient accuracy.




Figure 14.12. Effect of radius-vector variation on oscilloscope circle plot.

14.4 DESIGNING WITH RESOLVERS, SYNCHROS, INDUCTOSYNS, AND CONVERTERS

Selecting the right resolver, synchro, Inductosyn, and converter for a particular application has been considerably simplified by tables, such as those digested to form Figure 14.13. Before employing such tables, however, the designer should begin by asking a few fundamental questions.

Which measurement technology fits the needs of the application most closely? For instance, does the application call for angular resolutions of 10 bits or more, with absolute errors of less than 30 arc-minutes? Does the application need absolute measurement of angles (or will incremental representation suffice)? Must the transducer operate within a mechanically rugged or stressful environment? If the designer answers yes to any or all of these questions, then the use of resolvers, Inductosyns, or perhaps synchros is clearly called for.

Next, the designer should determine the amount of space available for the transducer, the required range of motion, amount and type of power available, and any unusual system interfacing requirements.

With this data in hand, the designer can select an appropriate transducer and converter from the tables. Figure 14.13 is a sampling of data appearing in tables available from Analog Devices. For each listed transducer, the tables list the range of compatible converters and the subsequent combined system performance.

RESOLVER TYPE	BASIC SPECIFICATIONS		ANALOG DEVICES CONVERTER PART NUMBER				
			1732/510 12 Bits	1S20/510 12 Bits	1S40/510 14 Bits	1S60/510 16 Bits	1S61/510 16 Bits
Singer Kearfott CR01093103 (Size 11) (400Hz)	System Resolution	(arc min)	5.27	5.27	1.32	0.33	0.33
	System Resolution	(bits)	12	12	14	16	16
	Converter Accuracy	(arc min)	21.00	8.50	5.30	4.00	10.00
	Resolver Accuracy	(arc min)	3.00	3.00	3.00	3.00	3.00
	System Accuracy	(arc min)	21.21	9.01	6.09	5.00	10.44
	Worst-Case Accuracy	(arc min)	24.00	11.50	8.30	7.00	13.00
	Tracking Rate	(rev/sec)	50.00	50.00	12.50	3.00	3.00
Harowe 11BRCX-30-C (Size 11) (1000Hz) Intrinsically Safe	System Resolution	(arc min)	5.27	5.27	1.32	0.33	0.33
	System Resolution	(bits)	12	12	14	16	16
	Converter Accuracy	(arc min)	21.00	8.50	5.30	4.00	10.00
	Resolver Accuracy	(arc min)	7.00	7.00	7.00	7.00	7.00
	System Accuracy	(arc min)	22.14	11.01	8.78	8.06	12.20
	Worst-Case Accuracy	(arc min)	28.00	15.50	12.30	11.00	17.00
	Tracking Rate	(rev/sec)	100.00	50.00	12.50	3.00	3.00
Thomson CSF 12T 11 RX4a (Size 11) (10kHz)	System Resolution	(arc min)	5.27	5.27	1.32	0.33	0.33
	System Resolution	(bits)	12	12	14	16	16
	Converter Accuracy	(arc min)	21.00	8.50	5.30	4.00	10.00
	Resolver Accuracy	(arc min)	5.00	5.00	5.00	5.00	5.00
	System Accuracy	(arc min)	21.58	9.86	7.29	6.41	11.18
	Worst-Case Accuracy	(arc min)	26.00	13.50	10.30	9.00	15.00
	Tracking Rate	(rev/sec)	100.00	170.00	42.50	10.50	10.50
Moore Reed 11RS236 (Size 11) (2kHz)	System Resolution	(arc min)	5.27	5.27	1.30	0.33	0.33
	System Resolution	(bits)	12	12	14	16	16
	Converter Accuracy	(arc min)	21.00	8.50	5.30	4.00	10.00
	Resolver Accuracy	(arc min)	1.00	1.00	1.00	1.00	1.00
	System Accuracy	(arc min)	21.02	8.56	5.39	4.12	10.05
	Worst-Case Accuracy	(arc min)	22.00	9.50	6.30	5.00	11.00
	Tracking Rate	(rev/sec)	100.00	50.00	12.50	3.00	3.00
Clifton Precision HSBJ-20-C-1 Multispeed 1:1 and 32:1 Size 20 (1200Hz)	System Resolution	(arc min)	0.16	0.16	0.04	0.01	0.01
	System Resolution	(bits)	17	17	19	21	21
	Converter Accuracy	(arc min)*	0.66	0.27	0.17	0.13	0.31
	Resolver Accuracy	(arc min)*	0.50	0.50	0.50	0.50	0.50
	System Accuracy	(arc min)	0.82	0.57	0.53	0.52	0.59
	Worst-Case Accuracy	(arc min)	1.16	0.77	0.67	0.63	0.81
	Tracking Rate	(rev/sec)†	3.12	1.56	0.39	0.09	0.09

*Refers to 32:1 Channel

†Refers to 1:1 Channel (i.e., Resolver Shaft)

a. Resolver to digital.

Figure 14.13. Excerpts from Converter Selection Charts.

14.4.1 DESIGNING WITH RESOLVERS AND SYNCHROS

Transducer Accuracy and Size

The first column in Figure 14.13a lists a variety of resolvers, their accuracies, and their sizes. For instance, the first section illustrates data for a resolver offered by Singer Kearfott, along with combined characteristics when employing various matching resolver-to-digital converters from Analog Devices. A model CR01093103 resolver has a body diameter of 1.1-inches (denoted as "size 11"), and an absolute accuracy of 3 arc minutes. This means that the resolver's analog electrical output, if measured ideally, will represent the resolver shaft angle to within ± 3 arc minutes.

System Resolution

The numbers in each subsequent column measure overall system performance when the resolver in the first column is mated with the RDCs noted in

BASIC SPECIFICATIONS		ANALOG DEVICES CONVERTER PART NUMBER					
		1732/560 12 Bits	1S10/560 10 Bits	1S20/560 12 Bits	1S40/560 14 Bits	1S60/560 16 Bits	1S61/560 16 Bits
Inductosyn Size	4"						
Number of Pitches	256						
Inductosyn Frequency	10kHz						
System Resolution	(arc secs)	1.24	4.94	1.24	0.31	0.08	0.08
System Resolution	(bits)	20	18	20	22	24	24
Converter Accuracy	(arc secs)*	4.92	5.39	1.99	1.24	0.94	2.34
Inductosyn Accuracy	(arc secs)*	3.00	3.00	3.00	3.00	3.00	3.00
System Accuracy	(arc secs)	5.76	6.17	3.60	3.25	3.14	3.81
Worst-Case Accuracy	(arc secs)	7.92	8.39	4.99	4.24	3.94	5.34
Tracking Rate	(revs/sec)†	0.39	2.65	0.66	0.17	0.04	0.04

*Refers to 256 Pitch Channel

†Refers to Inductosyn Itself

b. Rotary Inductosyn to digital.

BASIC SPECIFICATIONS		ANALOG DEVICES CONVERTER PART NUMBER					
		1732/560 12 Bits	1S10/560 10 Bits	1S20/560 12 Bits	1S40/560 14 Bits	1S60/560 16 Bits	1S61/560 16 Bits
Inductosyn Size	0.1"						
Inductosyn Accuracy	0.0001						
Length	9.99"						
System Resolution	(microinches)	24.41	97.66	24.41	6.10	1.53	1.53
Converter Accuracy	(microinches)	97.22	106.48	39.35	24.54	18.52	46.30
System Accuracy	(microinches)	139.47	146.08	107.46	102.97	101.70	110.20
Worst-Case Accuracy	(microinches)	197.22	206.48	139.35	124.54	118.52	146.30
Velocity	(inches/sec)	10.00	68.00	17.00	4.25	1.05	1.05

c. Linear Inductosyn to digital.

the top row. For instance, when the CR01093103 resolver is connected with a 1S60/510 (16-bit) converter, *system resolution* becomes 0.33 arc minutes, or 16 bits. This is calculated from:

$$\frac{360^\circ \times 60 \text{ arc-minutes}/^\circ}{2^{16} \text{ increments}} = 0.33 \text{ arc-minutes/increment} \quad (14.14)$$

Converter Accuracy

The specification for *converter accuracy* measures the worst case error of the RDC itself, specified in arc-minutes. For instance, the 1S60/510 converter's digital output will always be within ± 4.00 arc minutes of its analog input.

System Accuracy

This entry measures the overall rms error of the resolver in combination with the converter. System rms accuracy, in arc minutes, is calculated as the square root of the sum of the squares of the resolver accuracy specification and the converter accuracy specification. For the case of the CR01093103 resolver and 1S60/510 converter shown in the Figure 14.13a, for instance, the system rms accuracy is calculated to be 5.00 arc minutes.

Worst-Case Accuracy

This entry, also measured in arc minutes, represents the algebraic sum of the resolver accuracy specification and converter accuracy specification. For the example considered here, the converter's digital output will always indicate

the resolver's shaft angle to within ± 7.00 arc minutes of the resolver's actual shaft angle.

Tracking Rate

Converter *tracking rate* is a measure of the fastest shaft rotational speed, in revolutions per second, that the converter can measure without losing lock. The figures shown in the tables of Figure 14.13 are guaranteed minimums.

Multiple-Speed Resolvers

Multi-speed resolvers have sets of both coarse and fine windings on the stator. When the shaft is turned through 360 degrees, the output from the coarse winding is equivalent to that obtained by turning the shaft of a single-speed resolver through 360 degrees. The fine output, however, completes a number of electrical cycles that depends on the quoted resolver ratio. Typically, the ratio will be a binary number, such as 8:1, 16:1, or 32:1, in order to simplify the combinational logic.

Multi-speed resolvers permit much higher accuracies and resolutions than would be possible with single-speed systems. The resolvers need two converters, one on the fine output and one on the coarse, although it is the performance of the fine converter which determines the overall accuracy and resolution. A "shift and add" circuit combines the outputs of the two converters to give a single digital word representing the shaft angle, which is absolute at all times including switch-on.

Consider, for example, the Clifton Precision Multi-speed resolver, type HSBJ-20-C-1, the last section of Figure 14.13a. The ratio in this case is 32:1. If we use a 14-bit converter on the fine winding output, the resolution will be:

$$14 + \log_2(32) = 19 \text{ bits.}$$

The worst-case accuracy (error) will be:

$$0.5 + 5.3/32 = 0.67 \text{ arc-minutes.}$$

In other words, the coarse winding does not contribute to the error and converter accuracy error is reduced to $1/32$ of its specified value. In the limit, the accuracy can never be better than 0.5 arc-minutes, no matter what converter is used, but the *resolution* could be increased by using a higher resolution converter on the fine output. For accuracies better than this, Inductosyns, which are really high ratio multi-speed resolvers, can be used very effectively.

Frequency

Finally, note that each example in Figure 14.13 specifies nominal frequency of operation. Frequency primarily influences allowable tracking rates. Higher frequencies lead to higher allowable tracking rates, as can be seen if the tracking rate specs in Figure 14.13a for 400 Hz and 10 kHz resolvers are compared.

14.4.2 DESIGNING WITH ROTARY INDUCTOSYNS

Figure 14.13b illustrates system performance of Farrand rotary Inductosyns in combination with Inductosyn-to-digital converters from Analog Devices. In analyzing system performance, it is useful to think of multi-pole rotary Inductosyns as multi-speed resolvers. In a rotary Inductosyn, there are 2 poles/pitch, and the number of pitches is equivalent to the number of speeds of an equivalent multi-speed resolver.

In the case of the 256-pitch rotary Inductosyn and IRDC1732/560 (12-bit) converter of Figure 14.13b, for instance, resolution is calculated as:

$$\frac{360^\circ \times 60 \text{ arc-min}/^\circ \times 60 \text{ arc-sec}/\text{arc-min}}{256 \text{ pitches} \times 2^{12} \text{ increments}/\text{pitch}} = 1.24 \text{ arc-sec}/\text{increment} \quad (14.15)$$

Converter Accuracy

In the case of the rotary Inductosyn, *converter accuracy* in the table of Figure 14.13b is calculated from the inherent converter error, 21.00 arc-minutes for the IRDC 1732/560, divided by the number of Inductosyn pitches, as follows:

$$\frac{21 \text{ arc-min} \times 60 \text{ arc-sec}/\text{arc-min}}{256 \text{ pitches}} = 4.92 \text{ arc-sec} \quad (14.16)$$

System RMS Accuracy and Worst-Case Accuracy

In the case of rotary Inductosyns, *system rms accuracy* and *worst-case accuracy* are calculated in the same way as for resolvers, as discussed above. In addition, Inductosyns will experience a further error, of the order of fractions of an arc second, due to gain-mismatch in the preamplifier.

Tracking Rate

At comparable operating frequencies, the *tracking rate* of a rotary Inductosyn-and-converter combination can be calculated from the basic tracking rate of the converter divided by the number of pitches. In the case of the IRDC1732/560, TRACKING RATE is calculated as:

$$\frac{100 \text{ revolutions/second} \times 360^\circ/\text{revolution}}{256 \text{ pitches}} = 140.6^\circ/\text{second} \quad (14.17)$$

where 100 pitches per second is the tracking speed of the IRDC 1732/560 at 10 kHz.

14.4.3 DESIGNING WITH LINEAR INDUCTOSYNS

Figure 14.13c lists the overall performance of a typical linear Inductosyn from Farrand Controls when used with Inductosyn-to-digital converters from Analog Devices.

The Inductosyn is characterized by the length of its pitch, its overall length, and its inherent accuracy. Overall system performance figures are listed in the right hand columns.

Resolution

For instance, in the case of the Inductosyn shown, with 0.1-inch pitch, operating with an IRDC1732/560 Inductosyn-to-digital converter, *system resolution* (in microinches) is calculated from the Inductosyn *pitch* (in microinches) divided by 2^N , where N is the converter's rated number of bits. In this case:

$$\begin{aligned}\text{Resolution} &= 100,000 \text{ microinches}/2^{12} \\ &= 24.41 \text{ microinches}\end{aligned}\quad (14.18)$$

Converter Accuracy

The *converter accuracy* of Figure 14.13c is calculated by multiplying the Inductosyn *pitch*, in microinches, by the *converter accuracy* number in arc-minutes, divided by a conversion factor. For example, consider again the case of the IRDC1732/560. *Converter accuracy* is calculated by:

$$\begin{aligned}\text{Converter accuracy} &= 21 \text{ arc-minutes} \times \frac{100,000 \text{ microinches/pitch}}{360^\circ \times 60 \text{ arc-min}/^\circ} \\ (\text{error}) &= 97.22 \text{ microinches}\end{aligned}\quad (14.19)$$

System Accuracy

As in previous cases, *system rms accuracy* is calculated from the square root of the sum of the squares of the *converter accuracy* and *Inductosyn accuracy*.

Worst-Case Accuracy

This term is calculated by simply adding the *Inductosyn accuracy* to the *converter accuracy*.

Velocity

The velocity specification represents the maximum Inductosyn velocity in inches per second for which the IDC doesn't lose lock. Velocity is calculated from the *tracking rate* of a comparable converter at comparable operating frequency, multiplied by inches per pitch. For instance, in the case of the IRDC1732/560 converter of Figure 14.13c, and an 0.1-inch Inductosyn, the converter's *tracking rate* is 100 revolutions per second. Since one revolution of a resolver is equivalent to one pitch of an Inductosyn, 100 revolutions per second is equivalent to 100 pitches per second. Maximum *velocity* for this case is then calculated by:

$$\begin{aligned}\text{Velocity} &= 100 \text{ rev/second} \times 1 \text{ pitch/rev} \times 0.1 \text{ inches/pitch} \\ &= 10 \text{ inches/second}\end{aligned}\quad (14.20)$$

14.4.4 DESIGNING WITH DIGITAL-TO-RESOLVER/SYNCHRO CONVERTERS

When designing with digital-to-resolver, or digital-to-synchro converters, the accuracy of the converter has an important influence on the overall perform-

ance of the system. Such systems often are in the form of a feedback control loop, which also includes a resolver and RDC for position sensing.

To deal with an important concern when designing with digital-to-resolver converters, the designer should be certain to provide sufficiently large amounts of power to ensure proper performance of the resolver itself.

14.5 APPLICATIONS OF RESOLVERS, SYNCHROS, INDUCTOSYNS, AND CONVERTERS

The increasing industry attention to factory automation, robotics, and computer-controlled machine tools is resulting in expanded use of resolvers and Inductosyns.

14.5.1 MACHINE-TOOL APPLICATIONS

Figure 14.14 illustrates a typical central computer-based position-control system where an Inductosyn and an Inductosyn-to-digital converter constitute the primary feedback elements. As seen in the tables of Figure 14.13, the IDC produces a 12-bit digital word proportional to the Inductosyn slider's position within one pitch of the scale. The IDC's output also includes a direction signal (DIR), indicating the direction of motion, and a ripple-clock signal (RC), which indicates when the scale moves from one cyclic pitch to the next. A similar circuit configuration can be used in applications where a resolver and converter are used with a lead screw.




Figure 14.14. Inductosyn and IDC in a typical application.

Designers of systems employing resolver/leadscrew or Inductosyns may prefer a serial output to a parallel output. This is often the case where a resolver is replacing an incremental optical encoder. Such a system is shown in Figure 14.15. The IRDC 1731 has a serial output of 4000 counts per pitch, as well as a zero-crossing pulse which indicates when the traveler has moved from one pitch to the adjacent one. External counters convert the serial output to a parallel binary or BCD word. In this example, it is a parallel BCD word, which could drive a display to facilitate readout of position by the machine operator. Also, when used with Inductosyns, the choice of 4000 counts per pitch (instead of 4096, or 12 bits) makes it easier to divide the 2 mm pitch or the 0.1 inch pitch of typical Inductosyns into more-easily processed increments.




Figure 14.15. RDC in linear measurement application, with leadscrew.

14.5.2 GUN CONTROL

Military fire-control systems, such as that outlined in Figure 14.16, often use synchros to measure angular position, employ a processing system which accepts signals in synchro format and produces outputs, usually in coarse/fine

synchro format, to point a gun or missile launcher. Tracking synchro-to-digital converters can be used to process the synchro inputs, and a digital-to-resolver converter, in combination with control transformers, provides the output control signal.

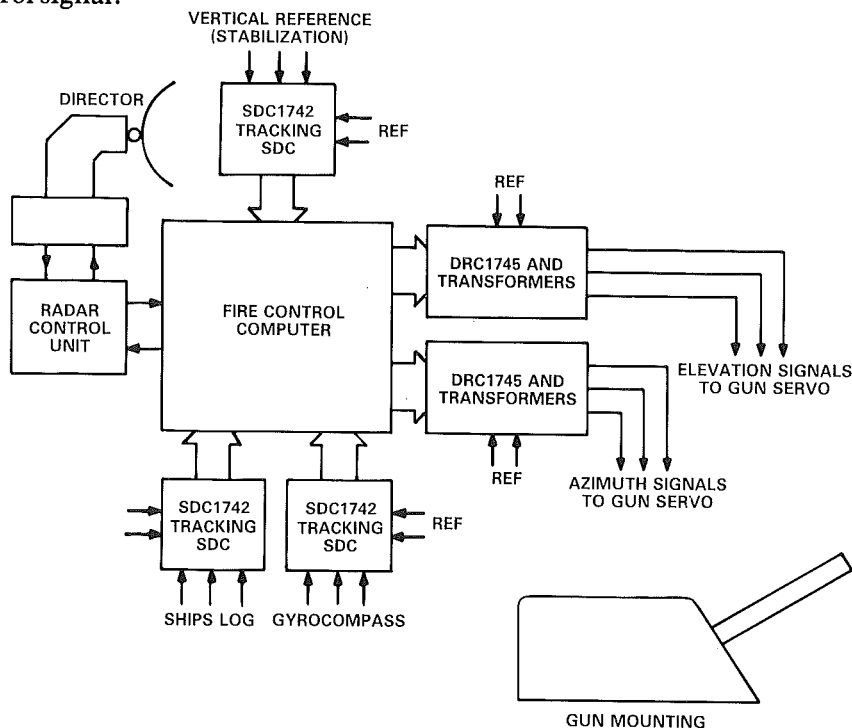


Figure 14.16. Fire-control application.

The fire-control computer of Figure 14.16 employs inputs from the ship's tracking radar in addition to various other ship sensors.

Figure 14.17 illustrates an alternative closed-loop mechanical positioning system that makes use of the velocity output voltage of the 1S24 RDC—a device having a high-grade tachometer output.

14.5.3 PLAN-POSITION INDICATOR

Figure 14.18 shows how a hybrid digital-to-resolver converter can be used in conjunction with a synchro-to-digital converter to generate the waveforms required for radar PPI displays.

The synchro signal from a control transformer representing the radar antenna angle is converted to a binary word by the synchro to digital converter. This digital angle is applied to the digital input of the DRC. A dc voltage is applied to the analog input to control the radius of the displayed raster. The outputs of the DRC, which are proportional to the sine and cosine of the antenna angle, provide the voltages required by the X and Y time base of the PPI display.

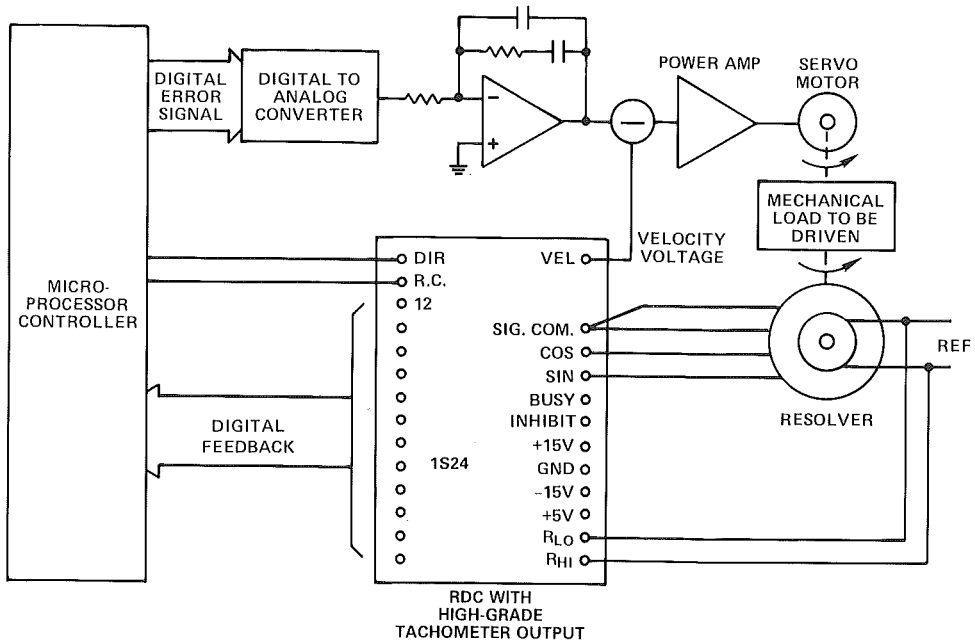


Figure 14.17. Resolver-to-digital converter in servo-loop application, showing use of velocity output for tachometric feedback.

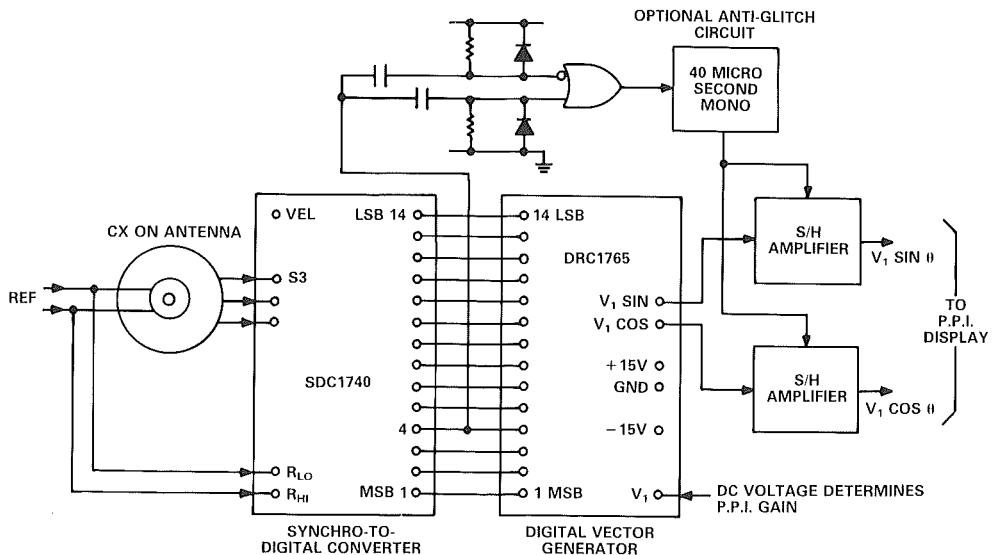


Figure 14.18. Converters in PPI application.

Radar antenna direction information is sometimes transmitted in serial digital form from the antenna to the processing equipment. This method is used, for example, in the case of microwave transmission of position data. At the receiving end, it may be required to convert these azimuth change pulses, or ACPs,

which usually number 4096 per revolution, into dc sine and cosine voltages to drive the PPI display. The ACPs are usually accompanied by an azimuth rotation pulse, or ARP, which occurs once per revolution of the antenna.

Figure 14.19 illustrates a circuit designed to accomplish the conversion. Note that the azimuth rotation pulse resets the counter to zero every revolution in case any spurious pulses have occurred during the the previous revolution.

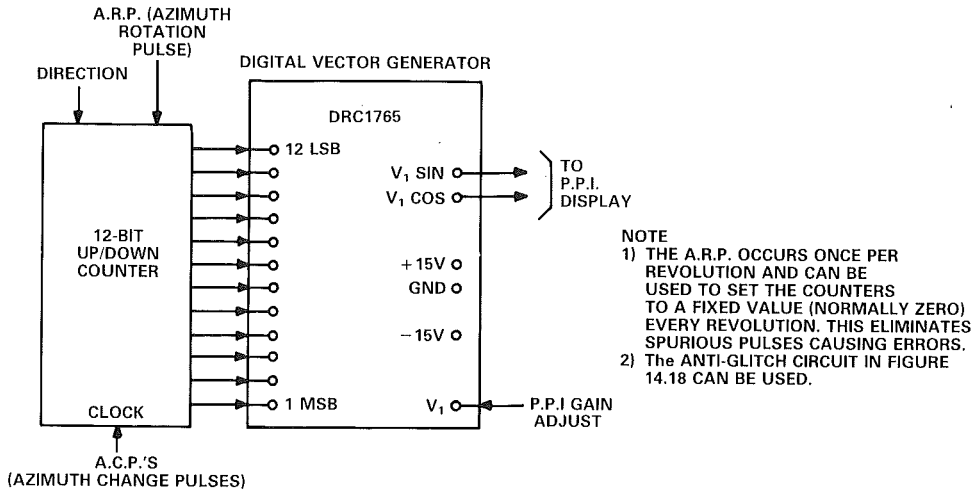


Figure 14.19. PPI application with serial input.

Chapter Fifteen

Voltage-to-Frequency Converters

15.1 INTRODUCTION

A voltage-to-frequency converter (VFC) accepts an analog voltage or current signal and provides at its output a train of pulses or square waves with periods inversely proportional to the average value of the input over each cycle of output. For a given input value, the number of pulses per second, or *frequency*, is proportional to that input value. Thus, the output digital word from a counter clocked by a VFC and read out at regular intervals will be proportional to the analog input. When used in an appropriate feedback loop, a VFC circuit will function as a *frequency-to-voltage* converter (FVC).

The VFC is a versatile building block that is quite useful in modern data-acquisition systems. In its most elementary application (measuring the frequency resulting from an applied voltage), an analog-to-digital data conversion is achieved in a very simple and economical manner. An integrating device, the VFC voltmeter has guaranteed monotonicity, very high resolution, and rejection of hum and noise. The time required to convert an analog voltage into a digital number is related to the full-scale frequency of the VFC and the required resolution of the measurement.

In general, a VFC as an a/d converter is slower than successive-approximation and "flash" types (See Chapter 7) but comparable in speed to other types of integrating converters, such as dual-slope types. A VFC voltmeter provides a kind of flexibility that none of the other data converters offers: conversion speed is easily traded for measurement resolution, when a VFC is used in conjunction with a computer-controlled frequency counter. In addition, a VFC voltmeter can provide measurement resolution far greater than can be achieved with any other integrated-circuit data converter. Resolution can, in

principle, be increased almost indefinitely by simply waiting long enough to accumulate a sufficient number of pulses to resolve the output frequency to the degree desired.

For example, an Analog Devices AD650 VFC, operating with a full scale frequency of 1 MHz, will achieve a resolution of 18 bits,—or four parts in a million—by allowing a counter to accumulate pulses for little more than a quarter of a second. The ultimate resolution of a VFC voltmeter is limited by noise and drift, and not by architectural constraints (i.e., length of a resistance ladder, number of comparators, or width of an internal counter). A VFC voltmeter offers a cost advantage when high resolution (rather than high speed) is of primary importance. And the cost savings can be quite dramatic if many signals are to be monitored simultaneously with high resolution.

A voltage-to-frequency converter can also perform the function of a digital-to-analog data converter (DAC), again with very high resolution. To achieve this, the VFC is configured as a frequency-to-voltage converter (FVC), and the digital information is presented as a train of pulses or square waves at a given frequency. In the very simplest DAC application—configured as a “tachometer”—using converters such as the ADVFC32 or AD650, each cycle of the input frequency causes a fixed quantity of charge to be dumped into a lossy integrator (this will be discussed in relation to the circuit of Figure 15.13). This kind of circuit is useful in low-speed applications, such as driving a panel meter or controlling power applied to a motor or a heater. At constant frequency, the output voltage is not purely dc; it has some ripple at the input frequency, but its average value is strictly proportional to frequency.

It is possible use filtering to reduce the amount of ripple on the voltage output, but only at the expense of settling time. An FVC with a much more favorable tradeoff between ripple and settling time can be achieved by using a VFC in a phase-locked loop (PLL). In the PLL (Sec. 15.4.3), the frequency of the input is compared with the frequency produced by a VFC, and the signal applied to this VFC is adjusted until the two frequencies are equal. This signal, which is proportional to frequency, becomes the voltage output. This circuit achieves very fast settling time of the voltage output with low ripple—at the expense of increased circuit complexity. When an AD650 v/f converter is used in a PLL application, the circuit can easily handle signal bandwidths of 70 kHz with a dynamic range (a measure of noise and ripple) of 80 dB.

Other uses of the VFC are simple variations and combinations of the two fundamental circuits described above. For example, an integrator can easily be built from the same components as a VFC voltmeter. A counter is used to accumulate pulses of the output frequency, and the total number of pulses over the gate time is interpreted as an integral (i.e., total charge delivered to the VFC), rather than as an average frequency. Very long integration times can be achieved using this method. For example, to evaluate the total energy in a signal, an analog multiplier continuously calculates the power delivered to

a load ($V \times I$), and the VFC can integrate the power signal to keep track of the total watt-hours used.

Another important application of a VFC is to transmit a high-accuracy analog signal through a noisy environment without interference. To accomplish this, the VFC converts the analog signal to a pulse train at a frequency proportional to the input, and that frequency signal is then transmitted. The frequency signal, consisting of large-amplitude digital pulses, is much less susceptible to interference than a high-resolution analog signal. At the receiving end, the signal is converted back into a voltage by one of the methods mentioned earlier.

The selection of the medium of transmission of the frequency signal is primarily determined by the nature of the task at hand. For example, galvanic isolation is achieved with a simple transformer or opto-isolator; extremely high-voltage isolation or transmission through environments having severe radio-frequency noise levels can be accomplished with a fiber-optic link; and telemetry can be accomplished with a radio link. In a sense, time can be considered as a noisy environment: a common need is to transmit information from the past into the future; this can be easily fulfilled by recording the frequency signal on magnetic tape and then replaying the tape when needed in the future.

15.2 DEFINITIONS AND SPECIFICATIONS

A voltage-to-frequency converter is a device that produces an output frequency in response to an input signal. The input signal may be presented as either a voltage or a current (in the discussions and examples that follow, voltage will be used as the input variable, but the reader should bear in mind that it could just as well be current). In the ideal case, the output frequency is directly proportional to the input. Thus,

$$\frac{f_{\text{out}}}{f_{\text{FS}}} = \frac{V_{\text{in}}}{V_{\text{FS}}} \quad (15.1a)$$

where f_{out} is the frequency of the output signal, V_{in} is the input voltage, and f_{FS} and V_{FS} are the respective full-scale quantities. Since they will be used as constants, they may be combined in a transfer constant,

$$f_{\text{out}} = G V_{\text{in}} \quad (15.1b)$$

where G is the gain of the VFC in events-per-second (Hz) per volt.

A real device will not have this perfectly linear transfer function—there will be offsets, gain errors, and some warping or *nonlinearity*. Nonlinearity is the primary specification of a VFC device; it warrants careful consideration since it determines the irreducible error of a system which uses the device (other errors can be reduced at a fixed temperature by trimming and calibration, either statically or dynamically).

In a VFC with *offset*, zero frequency does not occur at zero voltage. This means that, for example, the VFC ceases to oscillate when the input is at a small but non-zero voltage, or that it will not cease to oscillate with an input of zero volts. The offset of a VFC transfer function at constant temperature is typically not a problem, since it may be adjusted to be zero with very simple circuit trimming. If the offset should change, as a function of temperature, however, the resolution of the VFC will be compromised, because it is not possible to distinguish between a small change in the input and a change in the offset.

The *dynamic range* of the system will also be compromised by any offset drift. A dynamic range of a million to one means that a VFC which accepts a 10-volt full scale signal should also accept a 10-microvolt signal and produce an output frequency one millionth as great, with any difference due only to non-linearity. However, if the offset of the VFC transfer function should drift away from its trimmed value of zero, then a larger error in the ratio of frequencies produced by the 10-microvolt and 10-volt signals will result. If the offset drifts by $+5\mu\text{V}$, for example, it would effectively reduce the $+10\text{-}\mu\text{V}$ input to $+5\mu\text{V}$, and the frequency ratio would be two million. However, if the drift were $-5\mu\text{V}$, the effective minimum input would be $15\mu\text{V}$, and the ratio would be two thirds of a million. This example of distortion at the low end dramatically shows the need for an “autozero” circuit in any wide-dynamic-range system. An example of a circuit to provide compensation of offset errors will be found in the Applications section (15.4).

The linearity error of a VFC is usually specified by the *end-point* method. That is, the error is expressed in terms of the deviation from the ideal voltage-to-frequency transfer relation after calibrating the converter at full scale and “zero”. To be able to verify the linearity specification, one must have available a switchable voltage source (or a DAC) that has a linearity error less than 20ppm, and one must use very long measurement intervals to minimize counting uncertainty. Fortunately, reputable manufacturers submit every VFC to automatic testing for linearity, and it will not usually be necessary to perform this verification, which is both tedious and time consuming. However, if a nonlinearity test is required, either as part of an incoming quality screening or in final product evaluation, an automated “benchtop” tester would prove useful.¹

Figure 15.1 shows voltage-to-frequency transfer relationships, with the non-linearity exaggerated for clarity. The first step in determining nonlinearity is to connect the measured end points of the operating dynamic range (for example, 10mV and 10V) with a straight line; this straight line is the ideal relationship, which is desired from the circuit. The second step is to find the differ-

¹See “V-F Converters Demand Accurate Linearity Testing,” by L. DeVito, *Electronic Design*, March 4, 1982, for an example of implementation of such a system, based on the Analog Devices’ LTS-2010 benchtop tester.

ence between points on this line and the actual response of the circuit at a few frequencies between the end points—typically, ten intermediate points will suffice. The difference between the actual and the ideal response is a frequency error, measured in Hz. Finally, these frequency errors are normalized to the full-scale frequency, and expressed either as parts per million (ppm) or parts per hundred (percent) of full scale.

For example, if 100 kHz is full scale and the maximum frequency error is 5 Hz, the nonlinearity would be specified as 50 ppm, or 0.005%. The gain of the VFC is simply the slope of the “end-point” line which was used in the determination of the nonlinearity as described above.




Figure 15.1. Voltage-to-frequency converter nonlinearity.

The specification of *gain*, (or *full-scale error*) is usually quite loose—especially in monolithic devices. This specification is not very critical, since the gain can always be trimmed at a given temperature with simple circuit adjustments. Digital systems can be designed for periodic gain calibration to remove the effects of temperature drifts.

The offset of the voltage-to-frequency transfer relation is simply the intercept of the extrapolated “end-point” ideal straight line, through the measured end points, with the frequency axis, as shown in Figure 15.2. It is equally valid, in concept, to refer the offset to the input, i.e., express the offset as a voltage, by finding the intercept of the “end-point” line with the voltage axis. However, this would be inconvenient for computations of nonlinearity using the “end-point” line, since it would mix dependent and independent variables; also, offset may stem from more than one source.

The offset to be expected from a VFC must be calculated using the specifications of both input voltage offset and input bias current. The offset is primarily due to offset voltage in the operational amplifiers used as pre-amps, integrators, or voltage-to-current converters—see Section 15.3. However, in the

charge-balance type of VFC (e.g., the AD650), the input bias current of the integrator's operational amplifier also contributes an offset. And in multivibrator-type VFCs, such as the Analog Devices AD537, the input bias current of the op amp used as a voltage-to-current converter will also contribute an offset. For this reason it is more helpful to express the combined offset effects in terms of frequency.




Figure 15.2. Voltage-to-frequency converter offset.

Gain tempco, or full-scale error versus temperature, is simply a measure of the change in the slope of the “end-point” line with a change in temperature. This parameter of a VFC must be carefully considered, because the ultimate accuracy of a system may be limited by temperature changes and not by the linearity of the converter. For example, consider an Analog Devices Model 460 VFC, with a 1-MHz full-scale frequency: the linearity error is ± 150 ppm, and the gain tempco of the L grade is ± 15 ppm/ $^{\circ}$ C. With 10 volts full-scale, the 150-ppm nonlinearity will allow an absolute measurement accuracy of 1.5 millivolts; however, the gain TC will contribute an error of equal magnitude if the temperature should change by 10° C after calibration.

In general, the linearity specifications of VFC devices are so good that they cannot be fully exploited unless deliberate steps are taken to minimize offset and gain changes with temperature. It is usually a simple matter to install the device (especially if it is monolithic) in a constant-temperature oven, to stabilize its operating temperature against a changing ambient. Another approach to gain stabilization is to use an “auto-gain” circuit, which can adjust the scale factor of the VFC in order to produce a constant output in response to a known reference input. Such a system is described in the Applications section (15.4).

The *power-supply rejection ratio* (PSRR as a ratio, PSR in logarithmic form—dB) is a specification of the change in gain of the VFC as the power supply

voltage is changed. The PSRR is usually expressed in units of parts-per-million change of the gain per percent change of the power supply—ppm/%. For example, consider a VFC with a 10-volt input applied and an output frequency of exactly 100 kHz when the power supply potential is ± 15 volts. Changing the power supply to ± 12.5 volts is a 5 volt change out of 30 volts, or 16.7%. If the output frequency changes to 99.9 kHz, the gain has changed 0.1% or 1000ppm. The PSRR is thus 1000ppm divided by 16.7% which equals 60ppm/%.

Alternatively, the PSRR may be expressed in units of percent change in gain per volt change of the power supply—i.e., %/volt. For the example given, the equivalent specification would be $0.1\%/5\text{ V} = 0.02\%/ \text{volt}$. The user must beware when comparing specifications of various devices: units are not always commensurate. For example, PSRR of the AD537 is specified as $\pm 0.1\%$ per volt max. However, to convert this into ppm/% we must first assume a nominal operating power supply voltage. If we take the usual ± 15 volt supplies, then a one percent change of supplies is seen to be 0.3 volt, which yields a PSRR of $1000\text{ppm}/3.3\% = 300\text{ ppm}/\%$; if, on the other hand, the device is operated on a single 5-volt supply, the sensitivity will be 1/6 as much, i.e., 50 ppm/%.

15.3 SURVEY OF VFC DEVICES

Today's system designer has a choice among a variety of voltage-to-frequency converter types, including both monolithic circuits and modules; each has strengths and weaknesses which must be understood to make an appropriate choice for the task at hand. There are two basic voltage-to-frequency converter architectures; one uses a multivibrator, the other is a charge-balancing technique.

15.3.1 MULTIVIBRATOR TYPES

Figure 15.3 shows an example of a monolithic VFC design using a current-controlled multivibrator as the primary timing element. An operational amplifier converts the input voltage into a proportional current; this current determines the charging rate of external capacitor, C , which in turn determines the frequency of operation of the multivibrator. The device has an internal bandgap reference (Chapter 20), which provides both a constant 1-volt output and a thermometer output scaled to 1mV per kelvin.

The frequency output is delivered via an open-collector transistor, and only the emitter of this output device is returned to digital ground. The dedicated digital ground allows the complete separation of noisy digital spikes from the high-accuracy analog circuits.

Note that signal current is relayed to the multivibrator from the input amplifier through n-p-n transistors; this means that current of only one polarity can

be delivered to the timing capacitor. Thus, the net input voltage must always be positive.

In the case of a small input signal that has a large amount of noise, the instantaneous input voltage may be negative, which would call for a current of opposite polarity. However, this circuit cannot provide a bipolar current, so the input amplifier will saturate—and continuous integration of the input signal is lost. This means that positive peaks of noise will increase the charging current, but negative noise peaks cannot decrease the current below zero in order to compensate. The result is that the noise-rejection characteristics of the multivibrator VFC are not as good as those of the charge-balance architecture.




Figure 15.3. Multivibrator-type VFC architecture.

However, apart from this disadvantage, there are a number of advantages. The power consumption of a multivibrator VFC can be very low; for example, the AD537—a representative example of the type—can operate on a single supply voltage as low as +4.5 volts and consume a maximum of 2.5mA quiescent current. The device can operate with full-scale output frequencies up to 100 kHz and voltage-to-frequency nonlinearity of 0.1% (or 1000 ppm), referred to full scale. In addition, its output is a square wave, which allows it to be a-c coupled without introducing dc level shifts when the frequency is changed.

This combination of characteristics makes such a device an ideal choice for telemetry applications, where measurements at a remote location must be relayed to a central data collection area. The flexibility of single-supply opera-

tion and low power consumption allows battery-powered operation with long life, or convenient two-wire operation, in which a single wire pair will deliver power to the remote device and relay the frequency back to the central location. The two wire telemetry application is shown in Figure 15.4. The thermometer output of the device is especially convenient for remote thermometry.




Figure 15.4. VFC in two-wire single-supply operation. Both input power and output signal share the interconnecting line.

As suggested earlier, it is important that the signal appearing at the device inputs not be contaminated by hum or noise. If the signal is tainted, then the dynamic range of the system will be reduced. To preserve the continuous integration of the signal and thereby avoid the consequences of the device's susceptibility to noise, it is desirable to provide filtering if possible, and then to constrain the minimum d.c. level of the input voltage signal to be greater than the expected zero-to-peak value of the residual noise. This is often not a problem, especially when using the thermometer output. In remote measurements, care should be taken to locate the VFC as close as possible to the signal source in order to minimize noise pick up, since noise is usually introduced on long lengths of cable; this mode of operation is of course the forte of a VFC: the measurement circuitry is placed near the signal source, and only a digital frequency signal need traverse any length of cable.

15.3.2 CHARGE-BALANCE TYPES

If flexibility and low power consumption are the hallmarks of the multivibrator type of VFC (as exemplified by the AD537), then high speed, high linearity, and high noise rejection are the salient features of the *charge-balance* VFC architecture, employed by devices like the ADVFC32 and the high-performance AD650. The latter type is used as our example, because it provides

some features that are especially useful in conversion systems—not found on most general-purpose VFCs, e.g., separate digital ground, offset nulling for the integrator op amp, and a bipolar offset current for convenient half-scale offset when converting bipolar signals. Key specifications of such a device include typical nonlinearity of 50ppm for 100-kHz full-scale frequency and maximum nonlinearity specifications of 1000 ppm for 1-MHz full-scale frequency. Power requirements are conventional split supplies, at ± 9 to ± 18 volts, with maximum quiescent current of 8 mA.

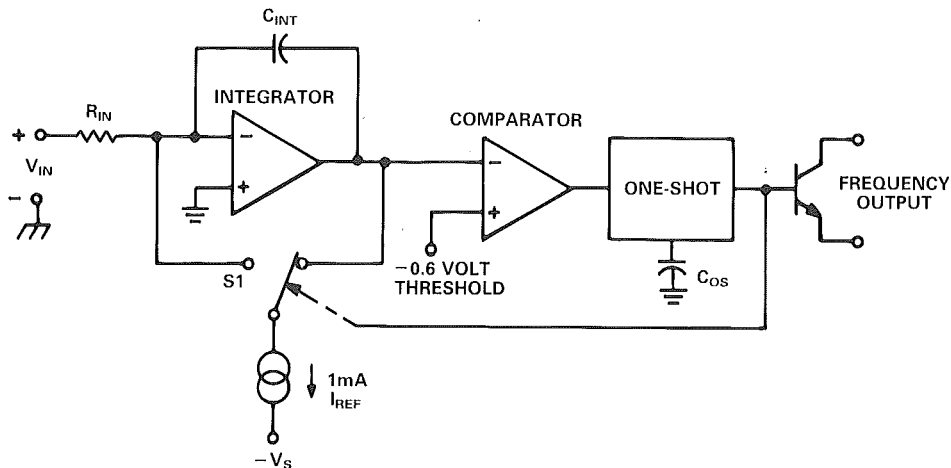


Figure 15.5. Charge-balance converter architecture.

Figure 15.5 shows a block diagram of such a device (AD650), connected to operate as a voltage-to-frequency converter. It comprises a summing integrator, a current source and steering switch, a comparator, a one-shot, and an output transistor. The integrator, in a feedback loop, forces an internal feedback current to exactly balance (over time) the input signal current, I_{IN} , which is either furnished directly to the op amp's summing point, as a current, or developed across R_{IN} by the input voltage. The current from the precision current source is applied as a feedback via S1 as short, accurately timed bursts of current—of opposite polarity to the input—in effect, as precisely defined packets of charge ($\Delta Q = I \Delta t$). The net charge is accumulated in the integrating capacitor, and the output of the inverting integrator is compared against a threshold; when the output falls below the threshold, indicating that the accumulated input charge is larger than the accumulated feedback, the comparator changes state and stimulates the one-shot to emit another pulse of current; at the same time an output pulse is emitted.

Thus, the number of charge packets required to keep the feedback loop balanced (each accompanied by one pulse of the output transistor), depends upon the magnitude of the input signal. Since the input signal is being balanced by a proportional number of discrete, equal charge packets, a linear

voltage-to-frequency transformation is accomplished. The frequency output is furnished via an open-collector transistor.

Here is a more detailed description and analysis of the circuit's behavior: When the output of the one-shot is low, the current steering switch, S1, diverts all the current to the output terminal of the op amp, where it has no effect on the integrator's rate of charge-accumulation; this is called the Integration Period (Figure 15.6a). When the one-shot has been triggered, and its output is high, the switch, S1, diverts the current to the summing junction of the op amp; this is called the Reset Period (Figure 15.6b). The figure shows the various branch currents and the integrator output voltage in both states. It should be noted that the output current from the op amp is the same for either state (in the integrating period, the current from the current source flows directly from the output; in the reset period, it flows through the capacitor), thus minimizing transients.




Figure 15.6. Charge-balance converter operation.

The positive input voltage develops a current ($I_{IN} = V_{IN}/R_{IN}$), which charges the integrating capacitor C_{INT} , the output voltage of the integrator ramps downward towards ground. When the integrator output voltage crosses the

comparator threshold (-0.6 V), initiating the Reset Period, the comparator triggers the one-shot; its time period, t_{os} , establishing the Reset Period, is determined by the one-shot capacitor, C_{os} (Figure 15.5). The integrator now ramps upward during the interval, t_{os} , by an amount:

$$\Delta V = t_{os} \frac{dV}{dt} = t_{os} \frac{1\text{ mA} - I_{IN}}{C_{INT}} \quad (15.2)$$

After the Reset Period has ended, the device starts another Integration Period, as shown in Figure 15.6, and starts ramping downward again. The amount of time, T_1 , required to reach the comparator threshold is:

$$\begin{aligned} T_1 &= \frac{\Delta V}{dV/dt} = t_{os} \frac{1\text{ mA} - I_{IN}}{C_{INT}} \frac{1}{I_{IN}/C_{INT}} \\ &= t_{os} \frac{1\text{ mA} - I_{IN}}{I_{IN}} \end{aligned} \quad (15.3)$$

Thus, the output frequency is:

$$\begin{aligned} f_{OUT} &= \frac{1}{t_{os} + T_1} = \frac{I_{IN}}{t_{os} \times 1\text{ mA}} \\ &= \frac{V_{IN}}{R_{IN}} \frac{1}{t_{os} \times 1\text{ mA}} \end{aligned} \quad (15.4)$$

Note that C_{INT} , the integration capacitor, has no effect on the transfer relation; its principal effect is to establish the amplitude of the sawtooth signal out of the integrator.

A key factor of the analysis is t_{os} , the one-shot time period. This time period can be divided into two time segments, approximately 300ns of propagation delay, and an interval which depends linearly on timing capacitor C_{os} . When the one-shot is triggered, a voltage switch, that clamps the capacitor at analog ground, is opened, allowing the capacitor voltage to change. An internal 0.5mA-current source connected to the capacitor terminal draws current through C_{os} , causing its voltage to decrease linearly. At approximately -3.4 V , the one-shot resets itself, ending the timed period and starting the V/F conversion cycle over again. The one-shot time period is:

$$t_{os} = \Delta V \frac{C_{os}}{I} + T_{delay} \quad (15.5a)$$

substituting the above values,

$$t_{os} = 3.4\text{ V} \frac{C_{os}}{0.5\text{ mA}} + 300\text{ ns} \quad (15.5b)$$

Therefore,

$$f_{OUT} = \frac{V_{IN}}{R_{IN}} \frac{1}{2(C_{OS} \times 3.4V + 300 \text{ ns} \times 0.5 \text{ mA})} \quad (15.6)$$

Component selection is not difficult. Since R_{IN} and C_{OS} are the only two parameters available to set the full-scale frequency and accommodate any given input voltage range, all of the necessary design information can be presented in graphical form. The selection guide of Figure 15.7 allows the user to quickly and easily pick a combination of R_{IN} and C_{OS} to suit almost any operating requirements.




Figure 15.7. VFC parameter selection guide.

The linearity is also affected by the choice of R_{IN} and C_{OS} ; this information has also been provided in Figure 15.7. In general, larger values of C_{OS} and lower full scale input currents (larger values of R_{IN}) result in better linearity. Although the selection guide is based on a 0-to-10-volt input signal range, the results are easily extended by simply scaling the input resistor in proportion to the desired input voltage span.

Selection of the integrating capacitor, C_{INT} , depends more on the particular application of the VFC than upon the input voltage span and output frequency range. If the input voltage is a steady signal with no interference in the form of hum or noise pulses (for example, in an application as a clock source, where the input signal is derived from a clean voltage reference), C_{INT} is most easily determined by the equation:

$$C_{INT} = \frac{100\mu\text{F/s}}{f_{MAX} \text{ Hz}} \text{ microfarads} \quad (15.7)$$

The minimum value of C_{INT} is 1000 pF.

A major advantage of integrating-type converters, such as the charge-balance VFC, is the ability to reject large amounts of additive noise by integrating both signal and noise. If the noise has a mean value of zero, then its effect on the measurement is reduced or even eliminated. This kind of integrating analog-to-digital converter uses a counter to accumulate pulses of the VFC output frequency for a fixed gate time. An integrating converter behaves like a filter having a transfer function of the form:

$$H(f) = \frac{\sin \pi f T}{\pi f T} \quad (15.8)$$

where f is the frequency of the input signal and T is the gate time of the counter. For very low frequencies the transfer function is unity; this means that the dc and low-frequency components are measured correctly by the integrating converter. However, for a frequency which corresponds to the gate period, and for its harmonics, the transfer function is zero. This means that additive noise is completely rejected at certain frequencies; the gate time is usually selected to provide rejection of power line hum.

Figure 15.8 is a logarithmic plot of the magnitude of the transfer function of Equation 15.8, normalized to 60 Hz, as a function of frequency. Note that, even for frequencies not integrally related to the reciprocal of the gate time, the envelope of the function provides some rejection of additive noise. Since the pulses of the output frequency of the VFC are totalized over a gate time, T , to give an average frequency over that time interval, samples of the input signal are generated at a rate of:

$$f_s = 1/T \quad (15.9)$$

where f_s is the sampling rate and T is the gate time of the counter. As mentioned elsewhere in this volume, the Sampling Theorem explains that it is not possible to reconstruct from the samples any component of the input signal having a frequency greater than half the sampling frequency. Furthermore, if there are components present (noise or hum) at frequencies higher than half the sampling rate, those components will be aliased into the baseband. The amplitude of these aliased signals is determined by Figure 15.8; however, the

apparent frequencies of the interfering components will all lie below half the sampling rate. Therefore, additional filtering may be needed if normal-mode noise components unrelated to line frequency are significant.




Figure 15.8. Magnitude response of sampling filter as a function of frequency.

These properties of noise rejection at certain frequencies depend upon continuous integration of the input signal; fortunately, the charge-balance architecture of the AD650 VFC depicted in Figure 15.5 performs continuous integration. Consider the Integration Period of the conversion cycle, as shown in Figure 15.6. When the input voltage is low, the downward ramp of the integrator output will be slow. If there is a noise pulse on the input signal which causes the instantaneous input voltage to be negative, then the integrator output will actually tend to move upward, away from the comparator threshold, but only during the duration of the noise impulse. Clearly, this is a transient situation; the instantaneous voltage does not remain negative for very long, and eventually the finite d.c. value of the input voltage causes the integrator output to ramp down and trigger the comparator. If a positive noise peak causes the integrator output to move down and trigger the comparator prematurely, no error is committed, since the period of the next VFC cycle will be longer, and the average frequency will still be correct.

The only situation which will allow an error to be made in the charge-balance process is if the continuous integration of the input signal is interrupted. For example, if the negative noise peak is too large in amplitude or too long in duration, the integrator output can drift all the way up to the plus rail and

saturate. When the input again becomes positive, and the integrator returns to normal operation, there will be a lapse in the memory of the signal, leading to errors.

This situation is most likely to arise with low power supply voltages and long, unshielded input leads, which are subject to hum and interference pick up. The amount of noise which would be required to interrupt the continuous integration depends upon the frequency of the interference and the values of the input resistance, R_{IN} , and the integrating capacitance, C_{INT} . The voltage waveform on the output of the integrator, in the presence of interfering noise, is simply the sum of the sawtooth—as derived in Figure 15.6—and the response of the integrator to the noise alone. For example, consider a value of 1000 pF for C_{INT} , corresponding to a full-scale frequency of 100 kHz. Let us use a 40.2-k Ω input resistor and a 330-pF one-shot capacitor. Suppose that the interference is hum at 60Hz; the output magnitude of the voltage from the integrator at this frequency is:

$$V_{OUT} = \frac{V_{IN}}{2\pi \times 60 \text{ Hz} \times 1000 \text{ pF} \times 40.2 \text{ k}\Omega} \quad (15.10)$$

Thus, there is a gain of 66 for the interfering hum. Noise at higher frequencies (for example, impulse noise from a spark plug) will be amplified less, and therefore constitute less of a problem than hum.

The highest voltage normally seen at the output of the integrator (see Figure 15.6) is the peak of the sawtooth signal at the end of the Reset Period. From Equation 15.2 we see that ΔV is most positive when the input signal (I_{IN}) is smallest. Using Equation 6 to calculate t_{os} (2.5 μ s), we find from Equation 15.2 that the maximum value of ΔV is 2.5 volts. Since the rising edge of the sawtooth starts from the -0.6 volt comparator threshold, the maximum value of the waveform is 1.9 volts.

Let us now assume that the device is operating on ± 9 -volt power supplies; with an allowance of 3 volts as headroom for the op amp output stage, the integrator output voltage cannot exceed 6 volts. Since the sawtooth signal would go to 1.9 volts in the absence of any interference, we now see that there are 4.1 volts of noise margin before the hum causes the op amp to saturate and sawtooth-plus-hum waveform to form a “flat top”. Since there is a gain of 66 for the 60-Hz hum (see Equation 15.10), the noise may have any amplitude up to 62mV.

To provide greater noise rejection, the value of the integrating capacitance can be increased. If C_{INT} were increased to 5,100 pF in this example, the allowable 60-Hz noise amplitude would be 470mV. This factor-of-7.6 improvement is due in part to a $5.1 \times$ reduction in gain seen by the hum, as calculated in Equation 15.10, and partly to a reduction in the ΔV of the sawtooth signal, as calculated using Equation 15.2. Another way to increase noise margin is to increase the input resistance. If R_{IN} of the example were 100 k Ω , then the value of C_{os}

would have to be 100 pF in order to maintain the full-scale frequency at 100 kHz. In this example, we could allow 210 mV of hum pick-up amplitude with a 1000-pF C_{INT} . Finally, combining approaches, if the integrating capacitance were 5,100pF, then 1.2-volt noise-pickup amplitude at 60 Hz could be tolerated at the op-amp input with no significant errors.

Synchronous VFCs

A variation of the charge-balance type of VFC is the *synchronous voltage-to-frequency converter* (SVFC). Devices of this kind require an external clock signal; the period of the clock signal is used as the length of the Reset period (Figure 15.6), making the scale factor proportional to the external clock frequency (In Equation 15.4, let $t_{os} = 1/f_{clock}$). In this way, the temperature drifts associated with the one-shot and its capacitor are eliminated, since the clock can be derived from a crystal oscillator for great temperature stability.

In an ordinary charge-balance VFC, both the dc reference (the 1-mA current source in Figure 15.6) and the one-shot can cause the gain to drift with temperature. Even if the on-chip portion of the one-shot were perfect, the temperature coefficients of the one-shot capacitor would still introduce errors. In the synchronous VFC, since the one-shot is completely eliminated, only the dc errors remain, and they can be made very small.

In general, the linearity and offset errors of SVFCs are comparable with those of other types of charge-balance devices, but the gain drifts are much better. In the SVFC, the pulses of the output frequency are generated in phase with the input clock; that is, the Reset period is initiated and terminated by a rising (or falling) edge of the input clock. This characteristic can be exploited in system applications (see section 15.4.5).

15.3.3 ASSEMBLED MODULAR DEVICES

Monolithic charge-balance voltage-to-frequency converters, such as the AD650 and ADVFC32, offer high-speed operation and outstanding linearity. However, for the most demanding applications, these devices cannot match the temperature stability of the best high-performance modular VFC devices. For example, the best grade of the Analog Devices Model 458 provides a temperature coefficient of gain equal to $\pm 5\text{ppm}/^\circ\text{C}$ at a 100-kHz full-scale, compared to the monolithic AD650's 150ppm/ $^\circ\text{C}$ gain-TC for the same frequency scale. However, both devices have approximately equal nonlinearity, in the several hundred ppm range. As a further example, the Model 460 modular VFC's gain-tempco is $\pm 15\text{ppm}/^\circ\text{C}$ at 1-MHz full-scale, and nonlinearity is 150 ppm at the same frequency scale; this can be compared with the AD650's gain tempco of $-400, +200\text{ppm}/^\circ\text{C}$ and nonlinearity of 1000ppm on the 1-MHz scale.

However, a price must be paid for the superior performance of the modular

devices, in size and power consumption. They are substantially more expensive, by a factor of about 10. Table 15.1 compares the salient performance specifications of the VFC devices discussed here.

	INTEGRATED CIRCUITS			DISCRETE MODULES	
	AD537	ADVFC32	AD650	458	460
Frequency (Max. Full Scale):	150 kHz	500 kHz	1 MHz	100 kHz(fixed)	1 MHz(fixed)
Linearity Error (PPM)					
10 kHz Full Scale	700	100	50	—	—
100 kHz Full Scale	1000	500	200	100	—
500 kHz Full Scale	—	2000	500	—	—
1 MHz Full Scale	—	—	1000	—	150
Gain Drift (PPM/°C)	50	75	75	5	15
Offset Voltage (mV)	5	4	4	10	10
Bias Current (nA)	100	100	20	—	—
Supply Voltage	4.5	± 9	± 9	± 13	± 13
Range (Volts)	to 36	to ± 18	to ± 18	to ± 18	to ± 18
Quiescent Current (mA)	2.5	8	8	+ 25, — 8	+ 25, — 8
Minimum Dissipation (mW)	12	144	144	429	429

TABLE 15.1 Specifications of Typical Devices

15.4 APPLICATIONS

V/F converters, as indicated earlier, can serve in a wide variety of applications. In this section, we will discuss in depth a few of the more popular and useful applications, including VFC voltmeters, signal isolation, and two possibilities for phase-locked-loop f/v conversion.

15.4.1 VFC VOLTMETER

A voltage-to-frequency converter can be used as the heart of a high-resolution voltmeter. In an instrument where the ultimate consumer of the data is human, or where extremely high resolution is necessary, a conversion rate of one or two per second is adequate. In this case, the VFC provides accuracy and linearity effectively, economically, and compactly. In a system where speed requirements are more aggressive, the VFC can still be used, but with somewhat reduced resolution; in these applications, the VFC provides the benefits of guaranteed monotonicity and no missing codes. This is a very important consideration in closed-loop control systems, where any non-monotonic behavior can lead to instability, with possibly dire consequences.

The conversion speed of a VFC voltmeter is related to full-scale frequency of the VFC and the required resolution of the measurement. For example, a measurement with 14 bits of resolution requires that changes as small as one part in 16,384 of full scale be detected. If we use a VFC with a 1-MHz full

scale, and accumulate pulses of the output frequency for $1/60$ of a second, then the full-scale reading of 16,666 would exceed 14 bits. Note that the gate time will provide infinite rejection of 60-Hz power line hum and its harmonics if continuous integration of the signal is preserved, as detailed in Section 15.3.

The accuracy of a VFC voltmeter is bounded by the linearity of the VFC. For example, consider a Model 460 modular VFC, with nonlinearity of ± 150 ppm and a full scale of 1 MHz. Since a 14-bit measurement implies resolution of 61 ppm of full scale, it is seen that the Model 460 will provide a relative accuracy to within 3 LSB. If resolution greater than 14 bits is required, the gate time can be increased to accumulate more pulses of the output frequency. In principle, there is no limit to the resolution obtainable with a VFC-based voltmeter, provided one is willing to wait. In practice, however, noise will limit useful resolution.

In order to determine the maximum usable resolution of a VFC-based voltmeter, it is necessary to measure a constant, highly stable voltage repeatedly and observe the variation in successive answers. As the gate period of the counter is increased, a point will be reached where the variation between successive measurements will equal the resolution of the measurement. In other words, there will be jitter in a supposedly constant frequency; there is no point in extending the resolution any further.

In order to get full accuracy from a VFC as a voltmeter, it is necessary either to adjust gain and the offset of the device by itself or to calibrate the measurement chain, starting with preamplifiers, signal conditioners, multiplexers, etc. These adjustments are easy to accomplish with trim potentiometers; however, this type of calibration only compensates for channel-to-channel or unit-to-unit variations. Since the errors caused by temperature drifts and component shifts after the initial calibration are not addressed by this technique, especially in system applications, it is often necessary and desirable to use an automatic gain and zero adjustment.

Such auto-cal adjustments are easy to do in software; simply measure the output frequency produced in response to two known input voltages—for example 10 volts and 1.2 volts. Then, by assuming a linear voltage-to-frequency transfer relation (always a safe bet with a VFC) the actual input voltage can be calculated for any output frequency. Also, if several intermediate voltage-reference points are provided, it is possible to make some first order corrections for nonlinearity.

These techniques are valuable if computational ability is already in place in a system (and what instrument these days does not have a microprocessor?) However, it may be more desirable to have a hardware auto-zero and auto-gain system if, for example, the necessary computations would overburden the computer, or if the results are required at a rate which is too fast to allow time for the arithmetic. Another system where a hardware gain and offset calibra-

tion would be necessary is in a frequency-to-voltage application, for example, as a tachometer, where a calibrated output voltage is produced in response to an input frequency.

As an example of the approach, a system will be described that has automatic adjustment of both gain and offset for an AD650 monolithic VFC. The *offset autozero loop* uses a sample-and-hold amplifier (SHA), as shown in Figure 15.9. The input voltage to the VFC is switched between the signal to be meas-




Figure 15.9. Autozero circuit for v/f converter.

ured and a ground reference, via an analog switch. In the *sample* mode, while the VFC's voltage input is grounded, the SHA measures and establishes the required correction for the VFC's offset; in the *hold* mode, the SHA maintains the zero adjustment, while the VFC is performing its measurement.

The sample-and-hold adjusts the offset of the AD650 by forcing a current at pin 13, one of the offset null pins. During the *sample* mode (VFC input grounded), the sample-and-hold and the integrating amplifier form a negative feedback loop; since the sample-and-hold's input is grounded, the signal returned to its feedback input must also be at zero; the sample-and-hold drives just enough current into the nulling input of the integrator to bring the output of the integrator to zero and keep it there. Since the output of the integrator is constant, its input current must be zero; thus the combined effects of amplifier offset voltage and bias current have been forced to zero. When the auto-zero cycle

to compare with a reference frequency the output of the VFC when connected to a 10-volt reference. The phase detector's voltage output drives a voltage comparator, which decides which frequency is higher; this output drives the UP/DOWN control of a TTL counter, which adjusts the gain of the DAC in the appropriate direction.

If, for example, the VFC's output frequency is higher than the reference frequency, the comparator's output will cause the counter to count down, reducing the gain of the VFC. The counter will continue to count down until the comparator changes sign, indicating that the signal frequency has been adjusted to be within 1 LSB of the reference frequency. When the auto-gain operation is completed, the data is latched into the DAC, which will hold the gain setting of the VFC until the next calibration cycle. The auto-gain loop requires the use of synchronous counters (TTL ...S169 are used in this circuit) to avoid the effects of ripple carry errors, which would cause the MDAC to assume erroneous values before settling to the proper result, leading to instability in the loop.

15.4.2 SIGNAL ISOLATION

A very important application of a VFC is analog signal isolation: transmission of an analog signal across a barrier with high accuracy. The analog signal is converted into a pulse train at a proportional frequency; this pulse train is relayed across a barrier and then converted back into an analog signal. The signals can thus be transmitted across a barrier accurately and without interference (assuming that the pulses are significantly higher in amplitude than any interfering signals), since they have been converted into a digital form.

The barrier may be physical (e.g., distance); it may be electrical, as in transmitting data across a large potential difference; or it may be a noisy environment such as a factory floor. The barrier may even be time; one may wish to save data by recording it on magnetic tape (for example) and reproduce it faithfully at some future time. The signal, once transformed into a pulse train with digital properties, may be transmitted in any fashion suitable to the application at hand. The analog signal may be converted into a frequency by straightforward application of the VFC; it is also possible to improve the accuracy of the transmitted signal by adding the auto-gain and auto-zero circuits described above.

The high operating frequency and excellent linearity of a 1-MHz VFC permits very accurate transmission of dynamic signals. However, the operation of a VFC at high frequency is not without concerns; remember that 1 MHz is a frequency usually classified as RF (it is central to the AM broadcast band). Most opto-isolators are not fast enough to relay the very narrow pulses produced by VFCs with high-frequency outputs. Figure 15.11 is an example of a circuit, using a 6N137 opto-isolator, which has fast-enough response to keep pace with the VFC.




Figure 15.11. VFC applied with opto-isolator.

Another concern is the difficulty of relaying digital pulses over a long length of cable. Simply connecting a TTL output to a very long cable without regard for proper termination will not work. The logic signal cannot properly drive the cable capacitance, and also the pulses may be distorted by reflections caused by improper termination.

Indeed, it is possible that a single pulse launched from the transmitter may be reflected at the receiver and return to the transmitter. At the transmitter, it may be reflected again and then return to the receiver where it will register as a pulse of the signal frequency for a second time. This echo effect may cause considerable errors in the transmission of the signal frequency.

A circuit is shown in Figure 15.12 to drive a 100-foot length of shielded twisted-pair cable, using the open-collector output of a VFC. The circuit operates in a current mode, and no voltage pulses are transmitted over the cable. The LED of the opto-isolator presents a load which is essentially a very low impedance in series with a 1.5 volt DC source. That is, there can be large changes in current through the LED, with only very small voltage changes, but it needs 1.5 volts to turn on.

Since there are no large voltage signals on the cable, stray capacitance is not a problem. With VFCs like the AD650, the cable can be conveniently driven directly, with no extra components. And, as an added bonus, a frequency-output voltage signal is available at the output-transistor collector for local use. At the receiving end of the 100-foot cable, only a local ground is required to power the opto-isolator's detector circuit. Thus this circuit provides complete galvanic isolation, in addition to transmission of the signal over a great distance. Ground loops and power-line faults can all be avoided with this simple and effective circuit.

ately. The voltage across the integrating capacitor is stabilized when the feedback leakage current through R_{IN} equals the average current being switched into the integrator. The net result is that the average output voltage is proportional to the input frequency.

The output voltage waveform is a series of exponential sections; at high frequencies, it will look somewhat like the inset of Figure 15.13, but at very low input frequencies it will look a lot more like a series of "spikes," i.e., a small d.c. level with high ripple content. The ripple can be reduced by increasing the value of the integrator capacitor, but at the expense of reducing the speed of response of the tachometer circuit. The ripple can also be reduced by filtering the output with a lowpass filter; however, care must be taken not to contaminate the signal by introducing offsets or gain errors in the filter.

The circuit shown can accommodate almost any input signal waveform. With the 1nF coupling capacitor and the 2.2-k Ω resistor, a TTL input creates a clean negative going spike that triggers the one-shot on each falling edge. For input signals with slower edges, higher capacitance or resistance may be used, as long as the comparator is never exposed to a voltage lower than $-0.6V$ for longer than the one-shot's time period. If this happens, the one-shot will trigger itself more than once per cycle, creating discontinuities in the F/V transfer function.

The tachometer has the advantage of being a very simple circuit, but it has difficulty handling rapidly changing information. Where the bandwidth required of the voltage output signal is low, such as might be expected from thermocouple signals, the economy provided by this circuit can be attractive. Analog signals at higher frequency, however, are more efficiently reconstructed with a phase-locked-loop circuit.

15.4.3 PHASE-LOCKED LOOP F/V CONVERSION

A phase-locked loop (PLL) provides a much more favorable compromise between speed of response and output ripple than the simple tachometer circuit, at the price of added complexity. In a phase-locked servo loop, a local VFC is driven to produce an output frequency which exactly matches the frequency and phase of an input signal. Since the voltage that drives the VFC (derived from the phase or frequency error signal) is proportional to the input frequency, it is taken as the output voltage of the frequency-to-voltage converter circuit.

The linearity of the frequency-to-voltage conversion depends almost entirely on the linearity of the VFC used in the loop. In certain circumstances, it is possible to achieve cancellation of linearity errors, if the nonlinearity of the VFC used in the receiver is almost identical to that of the VFC used in the transmitter. In order to do this effectively, though, the user must select matched pairs of converters, since random selection will be inadequate.

In some PLL applications, linearity is not an issue; for example, in a frequency synthesizer. In such applications, the oscillator output frequency is first processed through a programmable “divide by N” before being applied to the phase detector as feedback. Here the oscillator frequency is forced to be equal to “N times” the reference frequency, and it is this frequency output from the VFC that is the desired output signal, not the input voltage.

A very simple PLL circuit is shown in Figure 15.14. The phase detector consists of two D-type flip-flops and a NAND gate, while the charge pump and loop filter consists of an exclusive-or gate, a DMOS field-effect transistor, and an op amp. The charge pump relies upon power supply voltages and logic levels to establish the phase detector gain; among designers, this is usually considered a heinous crime—but the spirit of this circuit is simplicity.

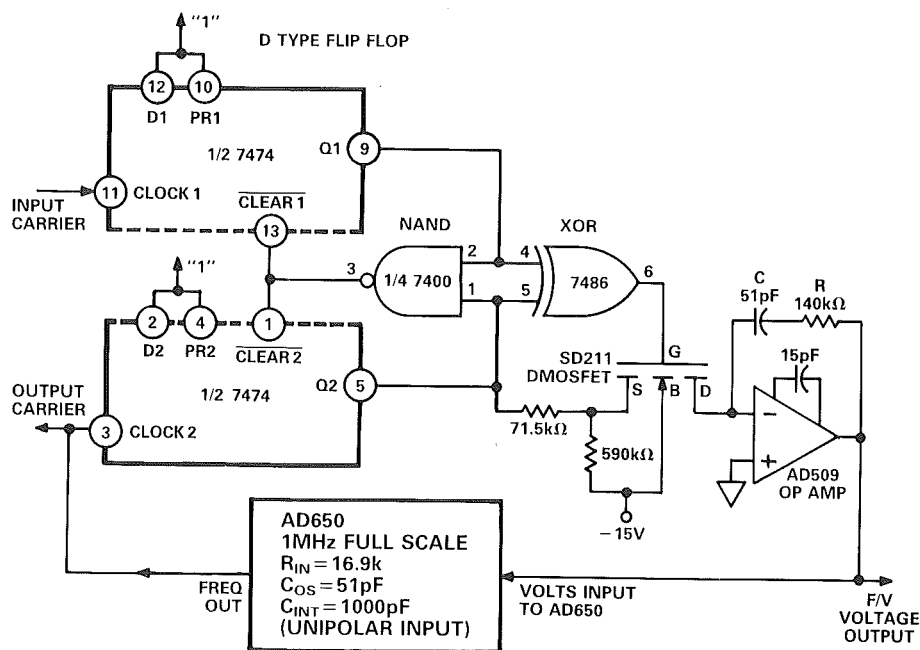


Figure 15.14. Phase-locked-loop in frequency-to-voltage conversion.

The phase detector is actually a so-called phase-frequency detector (PFD), which locks on edges. It provides proper feedback in the event of unequal frequencies; this means that the loop can never lose lock (unlike the situation when an analog multiplier is used as a phase detector; despite its superior noise rejection, a multiplier has only a limited range over which the output is a measure of the relative phase of the two inputs; if the phase error between the two inputs exceeds this range, the signal applied as feedback will not be correct and the loop will thrash about).²

²For more information on phase detectors and a full discussion of phase-locked-loop circuits, see F. M. Gardner, *Phase Lock Techniques*, 2nd ed. (New York: John Wiley and Sons, 1979).

For an analysis of the loop of Figure 15.14, start with the 7474 dual-D-type flip flop. When the input carrier matches the output carrier in both phase and frequency, the Q outputs of the flip flops will rise at exactly the same time. With the inputs of the exclusive-or (XOR) gate first at two zero's, then at two one's, the output will remain low, keeping the DMOS FET switched off. Also, the NAND gate will go low, resetting the flip-flops to zero.

Throughout the entire cycle just described, the DMOS integrator gate remains off, allowing the voltage at the output of the op-amp-connected-as-integrator to remain unchanged from the previous cycle. However, if the input carrier leads the output carrier by just a few degrees, the XOR gate will be turned on for the small time span that the two signals are mismatched. Since Q2 will be low during the mismatch time, a negative current will be fed into the integrator, causing its output voltage to rise. This will, in turn, increase the frequency of the VFC slightly, driving the system towards synchronization. In like manner, if the input carrier lags the output carrier, the integrator output will be forced down slightly to synchronize the two signals.

The $\pm 25\text{-}\mu\text{A}$ pulses from the phase detector are incorporated into the phase-detector gain expression (for K_d):

$$K_d = \frac{25 \mu\text{A}}{2\pi \text{ rad}} \approx 4 \times 10^{-6} \frac{\text{amperes}}{\text{radian}} \quad (15.11)$$

Also, the V/F converter is configured to produce 1 MHz in response to a 10-volt input, so its gain, K_o , is:

$$K_o = \frac{2\pi \times 1 \times 10^6 \text{ Hz}}{10 \text{ V}} = 6.3 \times 10^5 \frac{\text{radians}}{\text{volt-sec}} \quad (15.12)$$

The dynamics of the phase relationship between the input and output signals can be characterized as a second order system with natural frequency, ω_n :

$$\omega_n = \sqrt{\frac{K_o K_d}{C}} \quad (15.13)$$

and damping factor

$$\zeta = \frac{R}{2} \sqrt{C K_o K_d} \quad (15.14)$$

For the values shown in Figure 15.14, these relations simplify to a natural frequency of 35 kHz, with a damping factor of 0.8.

As a simple approach to determining component values for other PLL frequencies and VFC full-scale voltage, the following cookbook steps can be used:

1. Determine K_o (in radians per volt second) from the maximum input carrier frequency, f_{max} (in hertz), and the maximum output voltage, V_{max} :

$$K_o = \frac{2\pi f_{\max}}{V_{\max}} \quad (15.15)$$

2. Calculate a value for C based upon the desired loop bandwidth, f_n . This is the desired frequency range of the output signal. The loop bandwidth (f_n) is not the maximum carrier frequency (f_{\max}): the signal may be very narrow-band, even though it is transmitted over a 1-MHz carrier.

$$C = \frac{K_o}{f_n^2} \times 10^{-7} \frac{\text{VF}}{\text{rad-s}} \text{ farads} \quad (15.16)$$

where C is in farads, f_n is in Hz, and K_o is in radians/(volt-second).

3. Calculate R to yield a damping factor of approximately 0.8 using this equation:

$$R = \frac{f_n}{K_o} \times 2.5 \times 10^6 \frac{\text{rad}\Omega}{\text{V}} \text{ ohms} \quad (15.17)$$

where R is in ohms, f_n in hertz, and K_o in radians/volt-second.

If, in actual operation, the PLL overshoots or hunts excessively before reaching a final value, the damping factor may be raised by increasing the value of R . Conversely, if the PLL is overdamped, a smaller value of R should be used.

Phase-Locked-Loop Performance

The performance of the PLL circuit is demonstrated by the system shown in Figure 15.15. An analog signal is converted into a frequency, and then this

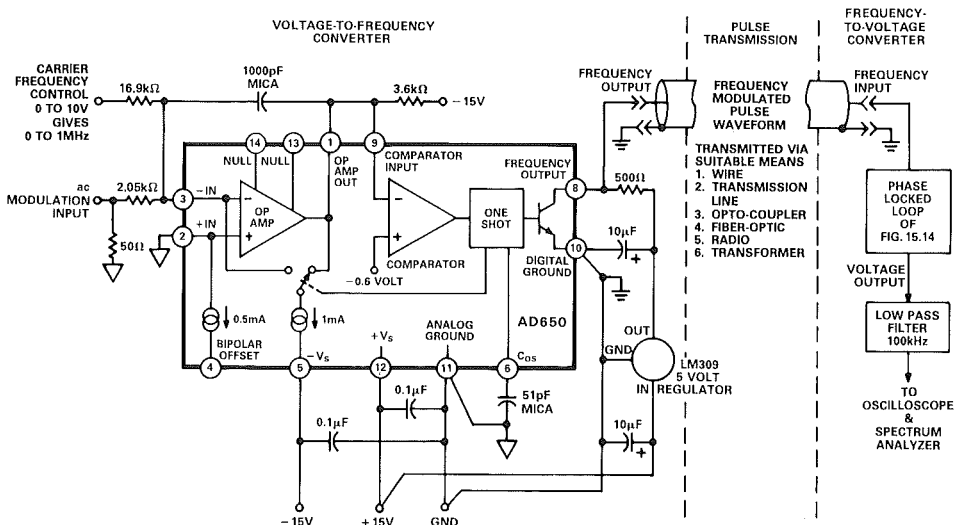



Figure 15.15. Circuit to demonstrate performance of PLL in frequency-to-voltage conversion.


frequency is converted back into an analog voltage by the PLL. The signal to be converted to analog is produced by a VFC with two additive inputs—one at dc, to set the carrier frequency, the other to establish an ac modulation signal. For this purpose, it is useful for the VFC's summing junction to be available.

The output frequency is then relayed to the PLL via a jumper cable (the signal at this point is a 5-volt digital pulse train, which may be transmitted in any fashion suitable to the application at hand). The filter on the output signal attenuates carrier feedthrough to allow easy interpretation of the signal with an oscilloscope and spectrum analyzer.


The response of the system to a step change of modulating frequency is shown in Figure 15.16a. The signal output is swinging between 5 volts and 10 volts, for an input step of 500 kHz to 1 MHz. Note that, despite the overshoot to 1.1 MHz, the response remains well-controlled. Note also the slight irregularity during the transition: it is caused by cycle-slipping during slewing, when feedback is lost temporarily and the PLL actually loses phase lock.




a. Step response.



b. Frequency amplitude response.



c. Zero-signal noise.



d. Harmonic distortion.

Figure 15.16. Performance of PLL as f/v converter.

The frequency response of the system, when driven with sine-wave excitation, is shown in Figure 15.16b. Here the output level is set to 2 volts peak-to-peak, and the carrier is 800 kHz. Note that the -3dB bandwidth is about 70 kHz, which is consistent with a damping factor of 0.8 and a natural frequency of 35 kHz*.

When an unmodulated carrier is applied to the PLL, the noise that appears at the output signifies the lower limit of the dynamic range of the system. The spectrum of the noise at the output of the PLL is shown in Figure 15.16c. By comparing this with the output levels of Figure 15.16b, the dynamic range of the system is seen to be about 80dB. The harmonic distortion of the system is shown in Figure 15.16d for a 2-volt peak-to-peak sinewave at 5 kHz; the amplitude of the first harmonic is seen to be 48dB below the fundamental. The harmonic distortion can be improved to the level of 60dB if the amplitude of the modulation is reduced, but this is at the expense of dynamic range, since the noise floor remains unchanged.

15.4.4 SOPHISTICATED PHASE-LOCKED LOOP

In some circumstances it is imperative that there be absolutely no feedthrough of the carrier into the output voltage of a PLL used in frequency-to-voltage conversion. Although it is possible to filter the voltage produced by the PLL to remove any residual carrier, it is not a trivial matter to design a filter with absolutely no d.c. offset voltage and no gain error. The post-filter may actually add errors worse than those to be eliminated.

A better strategy is to include the filter inside the phase-locked loop, where any offset or gain error will be compensated for by the feedback action of the loop. The voltage delivered to the VFC will always be forced to be the value required to produce an output signal of proper frequency and phase; hence, small gain and offset errors in the filter will not affect the output. Furthermore, since the voltage delivered to the VFC is a pure d-c level with no a-c components, there will be no phase noise, or incidental frequency modulation of the output frequency signal. This can be very important in such applications as frequency synthesis.³

Figure 15.17 shows a PLL circuit with a 5-pole Chebyshev active filter inside the loop. The phase detector is similar to the one used in the circuit of Figure 15.14. The charge pump operates in a differential mode and uses a separate reference voltage to establish the phase detector scale factor; also a differential integrator is used in the loop. The differential structure of the charge pump and integrator to some extent rejects digital noise.

*See page 13 of Gardner.²

³More details on a low-phase-noise synthesizer can be found in "Linear V-F Converter Chip Invades Module Territory," by L. DeVito, *Electronic Design*, February 23, 1984.

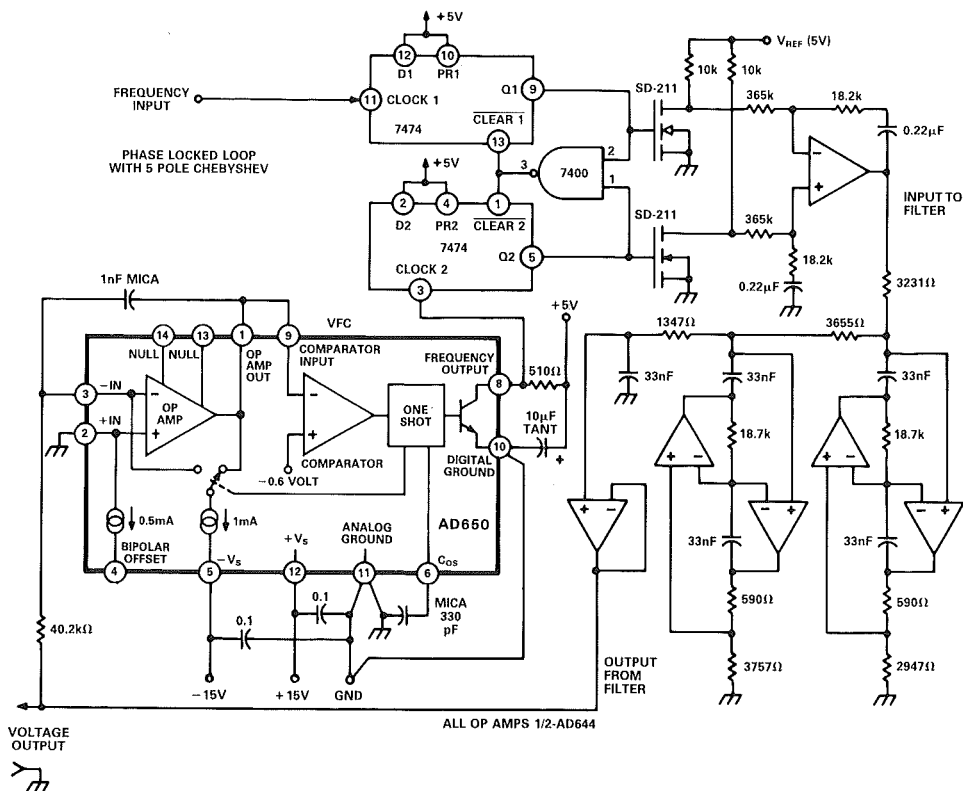




Figure 15.17. Phase-locked loop with 5-pole Chebyshev filter.

The low-pass filter has a 2.3-kHz cutoff frequency, based upon a frequency dependent negative resistor (FDNR) simulation of a doubly terminated ladder filter. The VFC used in this loop has a gain of 10^4 hertz per volt, or 100 kHz full scale for a 10-volt input. The performance of this PLL as a high quality F to V converter can be judged from the frequency spectra of Figure 15.18,



a. 30-kHz carrier.



b. 100-kHz carrier.

Figure 15.18. Output noise spectra of phase-locked loop of Figure 15.17, with and without filter.

which shows the output voltage in response to two constant-input carriers, 30 kHz and 100 kHz. The spectra in a and b show the performance of the circuit with and without the Chebyshev filter inside the loop. Note that the attenuation of carrier feedthrough and the concomitant improvement in dynamic range is as much as 45dB. The response shown is for a static carrier (the output voltage is constant); in the dynamic situation, the carrier feedthrough is more pronounced, and the improvement afforded by this circuit will be even more dramatic. As Figure 15.19 shows, the noise output voltage of the loop is less than 1 millivolt, peak to peak, and full-scale can be as much as 10 volts p-p; thus the dynamic range is 80dB.




Figure 15.19. Output noise of 5-pole-filtered PLL: waveforms at various input levels.

For analysis of the circuit, the block diagram of the PLL is shown in Figure 15.20. The loop transmission of the PLL is:

$$L(s) = \frac{5 \text{ V}}{2\pi \text{ rad}} \frac{sR_z C + 1}{sCR_i} K_f \frac{2\pi \cdot 10^5 \text{ Hz}}{s \cdot 10 \text{ V}} \quad (15.18)$$

The transfer function of the Chebyshev filter is designated as K_f ; it will be set equal to unity for now. The loop transmission then simplifies to:

$$L(s) = \frac{4.004 \times 10^{-3}s + 1}{1.6 \times 10^{-6}s^2} \quad (15.19)$$

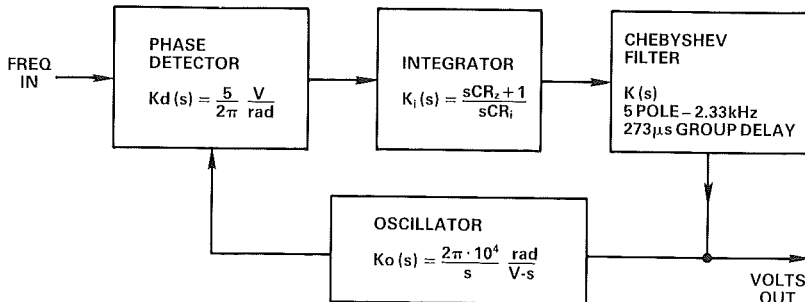


Figure 15.20. Block diagram of PLL for transfer-function analysis.

A Bode plot of the loop transmission is shown in Figure 15.21. Note that the loop crossover is at 400Hz and the frequency of the zero is 40Hz; this would provide a phase margin of 84° . The effects of the Chebyshev filter must now be considered. The corner frequency of the filter is 2.33 kHz; in the region of the 400Hz crossover, the magnitude of the filter response is unity, hence the loop crossover frequency is unaffected by the filter. The group delay of the Chebyshev filter at frequencies well below the corner frequency is $273\mu\text{s}$, which translates into a phase shift of 39° at 400Hz. Thus the phase margin of the loop is reduced to 45° by the presence of the filter.




Figure 15.21. Bode amplitude plot of PLL.

The dynamics of the PLL circuit are shown in Figure 15.22. The top photo shows the small-signal step response, with its rise time (10% to 90%) of $400\mu\text{s}$, which is consistent with a 400-Hz crossover. The bottom photo shows the response to a 1-volt step. Note that most of the response time is used for slewing, at about $30\text{ V}/\mu\text{s}$, to the vicinity of the final voltage level. While the integrator is capable of slewing faster, the error signal from the phase detector has only an average duty cycle of about 50% under this transient condition.

15.4.5 SYNCHRONOUS VFC APPLICATIONS

Two examples of synchronous v/f converter (SVFC) designs are the monolithic Analog Devices AD651 and the hybrid AD379. The AD651 operates at clock frequencies of up to 2 MHz, with linearity within 250 ppm; at lower clock frequencies, linearity approaches 20 ppm. With a 100-kHz clock,




Figure 15.22. Transient response of PLL.

its gain drift can be as low as 40 ppm per °C. It can operate with either a single supply—or in the more-traditional dual-supply mode, with supply voltages as low as 12 volts, or ± 6 volts. The hybrid AD379's gain and offset errors are about ten times smaller; it is also capable of handling bipolar input signals—it has a precision rectifier preceding the SVFC section and a sign bit, which indicates the polarity of the input signal.

Because the AD651 is a synchronous VFC, it doesn't require a one-shot in the conversion process. However, there is an on-chip one-shot to control the width of the frequency-output pulses; it drives the open-collector output. This is a useful feature in applications where the output pulse must be shorter than the clock period. For example, opto isolators require significant amounts of current; a very short pulse may be required to save power.

In the most elementary application of the SVFC, a stable clock signal is used, and output pulses are counted over a gated period, to measure frequency. However, if the gate period used to count the output pulses from the SVFC is derived from the clock signal by frequency division, a temperature-stable clock is not necessary. An example of such an application is shown in Figure 15.23.

Since the gain of the SVFC is proportional to the clock frequency, and the gate time is inversely proportional to the same clock frequency, the total




Figure 15.23. Synchronous VFC application. Accumulated count is a ratiometric measure of input voltage, independent of clock frequency, when counter clock and gate drive are derived from the same frequency source.

number of pulses during the gate time (for a given input voltage) will remain constant when the clock frequency changes.

Another application that exploits the proportionality between the SVFC gain and the clock frequency is the simple multiplier circuit of Figure 15.24. In this circuit, the clock frequency is established by the output of another VFC,




Figure 15.24. Multiplier application employing a free-running VFC to establish the scale factor of a synchronous VFC.

such as the AD650; since the VFC's output frequency is proportional to its input voltage, the frequency output of the SVFC is proportional to the product of the two voltages.

In an interesting system application, which benefits by the AD651's controllable-period one-shot, several data channels are multiplexed onto one digital line, as shown in Figure 15.25. Several SVFC devices, with very short output pulse-widths, are driven from different phases of a multi-phase clock, and all the open collectors are simply wire-or'd together. Thus, the presence or absence of an output pulse at each phase of the clock is identified with one of

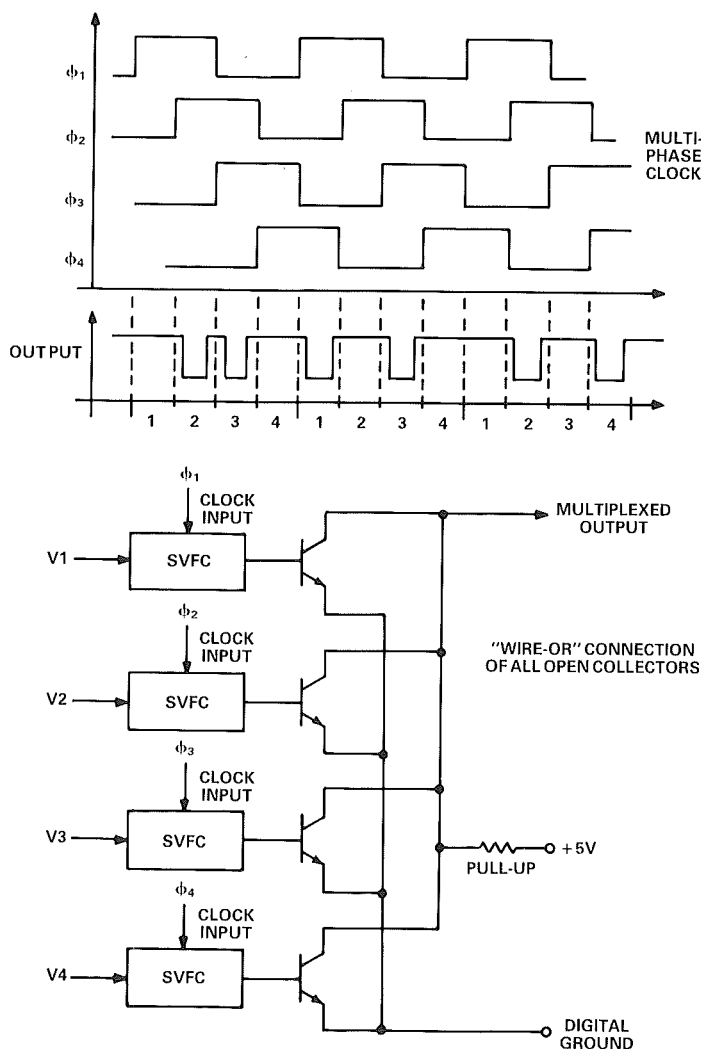


Figure 15.25. Multiplexed application. Each input channel is synchronized to a different clock phase.

the inputs. This arrangement can be quite cost-effective in a data-transmission type of application, such as that shown in Figure 15.12, since a single transmission line and opto-isolator can be shared among many data-channel outputs.

15.5 PRACTICAL MATTERS

15.5.1 DECOUPLING AND GROUNDING

It is good engineering practice to use bypass capacitors directly at the device's supply-voltage pins and to insert small-valued resistors (10 to 100 ohms) in the supply lines to provide a measure of decoupling between the various circuits in a system. Ceramic capacitors ($0.1\mu\text{F}$ to $1.0\mu\text{F}$) should be connected between the supply-voltage pins and analog signal ground for proper bypassing of the VFC. In addition, a higher-capacitance board-level decoupling capacitor ($1\mu\text{F}$ to $10\mu\text{F}$) should be located relatively close to the VFC on each power supply line.

Such precautions are imperative in high-resolution data acquisition applications, where one expects to exploit the full linearity and dynamic range of the VFC. Although some types of circuits may operate satisfactorily with power-supply decoupling at only one location on each circuit board, such practice should be strongly discouraged in high accuracy analog design. Devices such as the AD650 and the AD537 have separate digital and analog ground terminals. The emitter of the open-collector frequency-output transistor is the only signal node that should be returned to the digital ground. All other signals are referred to analog ground.

The purpose of the two separate grounds is to allow the high-precision analog signals to be isolated from the digital section of the circuitry; substantial amounts of noise on the digital ground can be tolerated without affecting the accuracy of the VFC. Such ground noise is inevitable when switching the large currents associated with the frequency-output signal.

For instance, at 1 MHz full-scale, it is necessary to use about 500 ohms of pull-up resistance in order to get a rise time fast enough to provide well-defined output pulses. This means that, with a 5-volt logic supply, the open-collector output will draw 10mA. The switching of this much current will surely cause ringing on long ground runs due to self inductance of the wires. For instance, #20 gauge wire has an inductance of about 20nH per inch; a current of 10 mA, switched in 50 ns, will produce a voltage spike of 50mV at the end of 12 inches of 20 gauge wire. With a separate digital ground, the VFC will easily handle these types of switching transients.

A remaining problem will be interference caused by radiation of electromagnetic energy by these fast transients. Typically, voltage spikes produced by inductive switching transients can capacitively couple into other

sections of the circuit. Another problem is ringing of ground lines and power supply lines, due to the distributed capacitance and inductance of the wires. Such ringing can also couple interference into sensitive analog circuits.

The best solution to these problems is proper bypassing of the logic supply at the VFC. A $1\mu\text{F}$ to $10\mu\text{F}$ tantalum capacitor should be connected directly between the supply side of the pull-up resistor and the digital ground pin. The pull-up resistor should be connected directly to the frequency output pin. The lead lengths of the bypass capacitor and the pull up resistor should be as short as possible. The capacitor will supply (or absorb) the current transients, and large ac signals will flow in a physically small loop through the capacitor, pull-up resistor, and frequency-output transistor. It is important that the loop be physically small for two reasons: first, short wires have less self-inductance, and second, smaller loops will not radiate radio-frequency interference as efficiently.

The digital ground should be separately connected to the power-supply ground. Since the leads to the digital power supply are only carrying dc current, they will not produce RFI.⁴

15.5.2 COMPONENT SELECTION

When applying typical monolithic VFCs, the user must select external parameter-setting and support components. Since the input resistor directly affects the scale factor, a high quality resistor must be used, especially if a low gain-temperature coefficient is desired. Do not use a carbon composition resistor here under any circumstances. Resistors of this type have a voltage-dependent nonlinearity which can be greater than the nonlinearity of the VFC, especially a high performance device such as the AD650.

The one-shot capacitor of the AD650 and the multivibrator capacitor of the AD537 also directly affect the scale factor and must be carefully chosen when low temperature drift is required. In general, a mica capacitor is good enough and quite cheap. However, in extreme cases, a ceramic C0G (formerly NP0) capacitor is necessary. The C0G ceramics also provide a bonus; they generally have low dielectric absorption compared to less-expensive ceramics. Teflon, polypropylene, and polystyrene capacitors also have low dielectric absorption, but their temperature coefficients are usually significant.

Dielectric absorption in the one-shot capacitor of charge-balance VFC types can cause difficulties. The problem occurs when the output frequency has been relatively constant for a long period of time and then changes. For some time after the change the capacitor will "remember" the previous average voltage it has seen and will have an apparent leakage current while it soaks

⁴More information on proper grounding and reduction of interference can be found in Ott, H. W., *Noise-Reduction Techniques in Electronic Systems* (New York, John Wiley and Sons, 1976).

at the new average voltage. This leakage current, not used to charge the main timing capacitor, produces an error. The result is that, after a change, the VFC will not immediately go to the proper frequency; it will slowly creep to the final value after a quick step to a nearby value. These effects can also cause excess nonlinearity of the voltage-to-frequency transfer relation, especially where the intrinsic linearity of the device is very good, such as at the low-frequency scales of the AD650.

Dielectric absorption can also create inaccuracies when a poor-quality capacitor is used for the integrating capacitor; however, the effect on linearity is not as severe as the degradation which can be caused by defects in the one-shot capacitor. The reason is that the average voltage seen by the integrating capacitor does not change as much, as a function of frequency, as the average voltage seen by the one-shot capacitor.

However, defects in the integrating capacitor can still create problems for the system designer. For example, consider the situation where one wishes to "squench" an AD650 (i.e., stop it from oscillating) by applying a negative input voltage, which causes the integrator output to rise up to the positive supply level and saturate. The integrating capacitor will then soak at a voltage equal to the supply potential; when the VFC is taken out of this "standby" mode, the integrating capacitor will remember its previous situation and slowly release absorbed charge as it soaks at its new average voltage. The charge being given up by the integrator capacitor is like a leakage and contributes directly as an input error. As a result the output frequency will not immediately be stable; it will slowly creep to a final value as the integrating capacitor forgets its past.⁵

⁵More information about dielectric absorption can be found in "Understand Capacitor Soakage to Optimize Analog Systems," by R. Pease, *EDN*, Oct. 13, 1982.

Chapter Sixteen

Intentionally Nonlinear Converters

D/A converters map the set of n -bit binary numbers into an equivalent set of 2^n voltages, currents, or gains. In most converters, for a given binary integer—of value, N —the nominal equivalent voltage is proportional (*linear*): In DACs, $V_{\text{OUT}} = V_{\text{FS}} N/2^n$, where V_{FS} is the nominal full-scale voltage. For example, if $n = 8$ and $N = 200$ (binary value: 11001000), then the nominal voltage corresponds to the value of the ratio, 200/256 of full-scale.

For such converters, the linear relationship is the key specification. It is in wide use because most measurements are linear measurements (we use linear measures for length, weight, voltage, current, power, etc.); but there do exist circumstances for which a nonlinear relationship between a physical quantity and a digital number is necessary or desirable. Some of these will be discussed below.

A nonlinear relationship can be achieved in three basic ways—through nonlinear analog signal conditioning, through lookup tables or nonlinear digital processing, or by means of a nonlinear conversion. Some aspects of nonlinear *analog* signal conditioning were discussed briefly in Chapter 2; A venerable but still useful reference on this subject is the *Nonlinear Circuits Handbook*.¹ Nonlinear *digital* signal processing can be achieved through software and/or hardware operations (see Chapters 5 and 21). The present chapter discusses aspects of the third approach, nonlinear *conversion*, with emphasis on logarithmic multiplying d/a converters. Nonlinear conversions involving transducer variables in trigonometric form i.e., resolvers, are discussed in Chapter 14.

¹D. H. Sheingold, ed., *Nonlinear Circuits Handbook* (Norwood: Analog Devices, Inc., 1974).

16.1 WHY INTENTIONALLY NONLINEAR CONVERTERS?

Many important applications involve the handling of analog signals over a wide dynamic range. Two important classes of signals requiring such handling are: (1) signals requiring accurate correspondence between physical quantities and numbers, and (2) signals calling for wide dynamic range with a specified accuracy at any level. The difference can be seen by a discussion of two examples.

In numerical machine control (NMC) applications, a typical requirement might be to move a workbed over a total distance of 1 foot, with a minimum increment distance of 0.001". This implies that the controller has a resolution of 1 part in 12,000 over the working distance. Any point in the range must be accurately located. Thus, a linear-coded D/A converter, with at least 14-bit resolution (1 part in 16,384), would be needed to meet this requirement. A discussion of high-resolution conversion will be found in Chapter 17.

On the other hand, a DAC might be required to set a system gain, with a gain tolerance no better than $\pm 12\%$ (1 dB)*, but over a range of 10,000:1, or 80 dB. For this purpose, an attenuator with (say) 128 2/3-dB steps would seem suitable to map the entire gain function. Since only 128 steps are needed, the digital input to a multiplying d/a converter with an appropriate conversion relationship would need a resolution of only 7 bits ($2^7 = 128$).

In this case, the appropriate conversion relationship is exponential; and the digital input is proportional to the logarithm of the desired gain ratio. Note that, if the conversion were linear, the accuracy would have to be $(1/8)/10,000$, or 1:18,000, i.e., 17 bits—and considerable accuracy would be wasted at the higher gains.

The ability to perform the desired function at a much lower resolution level provides savings in cost, package size, and system complexity; however, there may be a need for software or a lookup table to provide the digital data in the proper form for conversion, if it does not already exist in that form.

16.2 UNDERSTANDING LOGDACs

A LOGDACTM is a multiplying d/a converter with gain proportional to the exponential of the digital input. Equal changes of digital input produce equal ratios of analog gain change. When used in the forward path of an op amp (its specified mode of operation), it produces attenuation. Used in the feedback path, it provides gain. Figure 16.1 is the block diagram of an attenuator using a typical LOGDAC; it has an 8-bit μ P-bus-compatible digital data input, and it employs decoding logic and a high-resolution DAC to provide nominal 0.375-dB gain steps for 1-bit input changes.

*Using the relationship, $\text{dB} = 20 \log_{10} (\text{ratio})$, if the ratio is 1.12, the corresponding logarithmic quantity is 0.98 dB.

LOGDACTM is a trademark of Analog Devices, Inc.




Figure 16.1. Block diagram of 88.5-dB μ P-compatible LOGDAC.

To understand the operation of a LOGDACTM, consider first the operation of a linear-coded CMOS D/A converter as an attenuator. The simple block diagram shown in Figure 16.2a treats the D/A converter as a black box with analog input V_{IN} , digital input N and analog output, V_{OUT} . The transfer function is simply

$$V_{OUT} = \alpha V_{IN} \quad (16.1)$$

where α is the attenuation factor.

In a linear DAC, the attenuation factor's relationship to the value, N , of the digital input word is

$$\alpha = \frac{N}{2^n} \quad (16.2)$$

where N is the base-10 integer equivalent to the digital input code and n is the resolution (in bits) of the linear converter. In a LOGDAC, α is an exponential function of the digital input code. This nonlinear relationship affects only the *gain*; the analog input and output are always linearly related, with low distortion.

Figure 16.2b shows the basic current-switching circuitry employed in CMOS DACs (see Chapter 7). Because the input resistance of the R-2R ladder is constant and equal to the ladder resistance, R , regardless of digital input code or signal level, we can define an input signal current, $I_{IN} = V_{IN}/R$. If the lower end of each shunt arm is assumed to be at zero volts, regardless of switch position (in normal operation this assumption is valid), the input current, I_{IN} , will be subdivided equally at each ladder node, as indicated in the diagram. Thus, the current in any shunt arm is one half the current in the next left shunt arm; the input current is divided equally between the first shunt arm and the series arm leading to the rest of the ladder. Each attenuation of one-half is equivalent to -6 dB (actually, -6.02 dB), or $20 \log_{10}(1/2)$.

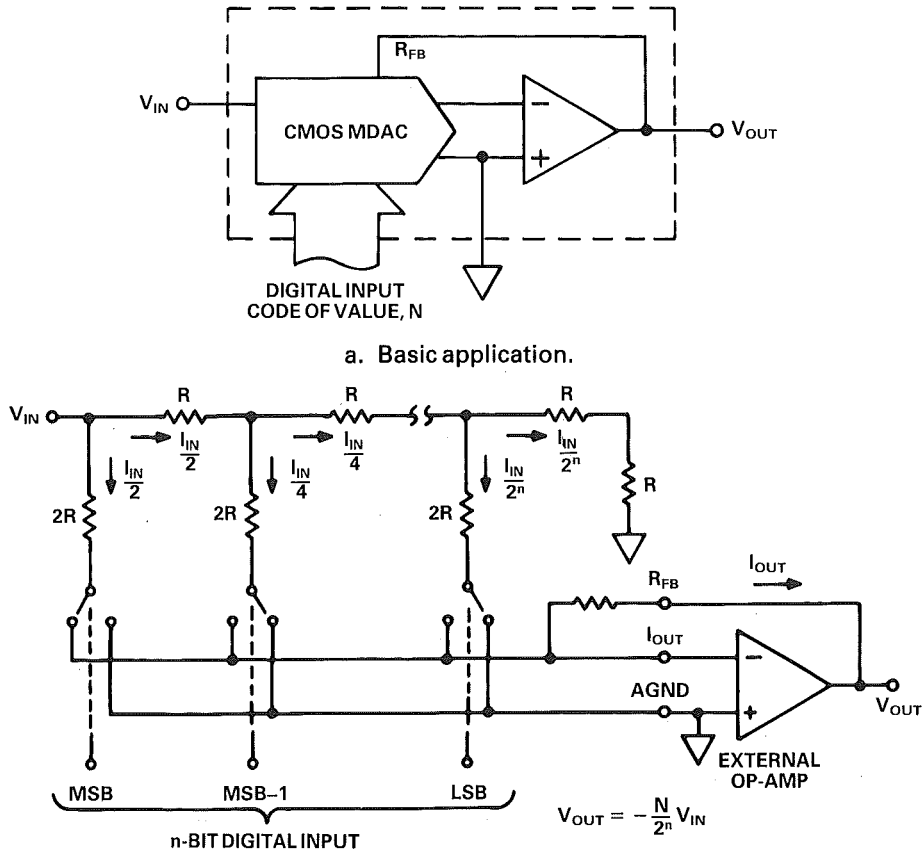


Figure 16.2. Review of linear multiplying DACs.

A linear converter has a logarithmic relationship requiring no computation or code manipulation, if one considers just the individual bits; since each bit is a power of two, its bit number, starting with the MSB, is simply the negative of the log of its attenuation. Thus, the logarithm of the MSB (ratio to full scale = $1/2$), to base 2, is -1 , \log_2 of the next bit is -2 , etc. \log_2 of the LSB is $-n$.*

For any linear-coded converter with n -bit resolution, there are 2^n possible input codes. If the all-zeros, or all-bits-off, input code is excluded,† there remain $2^n - 1$ input codes, corresponding to a choice from among $2^n - 1$ output attenuation values. These levels of attenuation are linearly related to the input voltage; each level is separated from the adjoining level by one least significant

*To transform \log_2 to \log_{10} , multiply by $\log_{10}(2)$, or 0.3010. Thus, $\log_{10}(\text{MSB}) = -0.3010$, $\log_{10}(\text{Bit } 2) = -0.6020$, etc. Using the definition, $\text{dB} = 20 \log_{10}(\text{ratio})$, the dB measure of the MSB is -6.02 (approximately 6), of Bit 2, -12 dB, etc., down to the LSB, with an approximate dB measure of $-6n$.

†In a logarithmic device—which is the goal of this transitional discussion—a gain of zero (0-V output), or perfect muting, is not a step in the logarithmic sequence; however, it is often available in actual devices.

bit (LSB), which is a constant percentage of input full-scale range (FSR), since

$$1 \text{ LSB} = \frac{\text{FSR}}{2^n} = \frac{V_{\text{IN}}}{2^n} \quad (16.3)$$

For an n -bit converter of this type, the maximum digitally controlled attenuation ratio, from all ones to 1 LSB is

$$20 \log_{10}(2^n - 1) \quad (16.4)$$

For a hypothetical 6-bit device (Figure 16.3), chosen as an example, the range is approximately 36 dB. Each 1-bit increase in the resolution of the converter increases the attenuation range by 6 dB. A 10-bit device has a maximum attenuation range of $10 \times 6 \text{ dB} = 60 \text{ dB}$.

Thus, it is possible to use just the individual bits of a linear DAC for logarithmic operations, without further errors or decoding, if one is willing to put up with the rather coarse 6-dB spacing. However, as Figure 16.3 shows, there are a great many codes available for interpolation, if properly chosen—but they occur mainly at the lower attenuation values.

To achieve constant dB step sizes, i.e., output gain changes which are equal percentages of reading, it is necessary to select the closest available individual codes which will yield the proper values of gain.




Figure 16.3. Log-scale plot of attenuation at individual bit codes for a linear 6-bit DAC.

For example, suppose we want 3-dB spacing, and are willing to choose the nearest code that interpolates between each pair of accurate 6-dB steps.

Since -3 dB corresponds to a ratio of 0.707, the 6-bit code corresponding to -3 dB would be about $45/63$ ($= 0.714$), or 101101. The difference in the actual ratios, which is an error due to limited resolution, must be added to errors in the accuracy of the DAC (assume that they are within $\pm \frac{1}{2}$ LSB) to establish the expected overall accuracy.

As Table 16.1 suggests, and Figure 16.4 illustrates, the errors, which are quite reasonable at the upper end of the range, and conservative for the 3-dB resolution, become rather large at the lower end of the range.

Required Attenuation dB	Possible Input Code		Actual Attenuation dB	Accuracy Actual Step dB	Possible Output Attenuation Range in dB with Respect to V_{IN}
	MSB	LSB			
0	11	11 11	-0.137	+0.069 -0.069	-0.068 to -0.206
-3	10	11 01	-3.06	+0.097 -0.096	-2.963 to -3.156
-6	10	00 00	-6.021	+0.135 -0.137	-5.886 to -6.158
-9	01	01 11	-8.889	+0.187 -0.191	-8.7 to -9.08
-12	01	00 00	-12.041	+0.267 -0.276	-11.774 to -12.317
-15	00	10 11	-15.296	+0.386 -0.404	-14.91 to -15.7
-18	00	10 00	-18.062	+0.527 -0.561	-17.535 to -18.623
-21	00	01 01	-22.144	+0.828 -0.915	-21.316 to -23.059
-24	00	01 00	-24.082	+1.023 -1.16	-23.059 to -25.242
-27	00	00 11	-26.581	+1.339 -1.584	-25.242 to -28.165
-30	00	00 10	-30.103	+1.938 -2.499	-28.165 to -32.602
-33	00	00 01	-36.124	+3.522 -6.021	-32.602 to -42.145

Table 16.1. Output attenuation and errors when 6-bit linear DAC is employed as logarithmic DAC with 3-dB gain steps.

Some interesting observations can be drawn from these results:

- 1) Highest step accuracy occurs at the smallest value of attenuation; accuracy decreases, expressed in increasing dB of error, with increasing attenuation. At the 6-bit converter's maximum required attenuation level of -33 dB, the actual output can range from -32.6 dB to -42.1 dB , a spread of almost 10 dB.




Figure 16.4. Attenuation of linear DAC with key input codes interpolated at 3-dB intervals. Error bands are shown.

2) From attenuation levels of -21 dB and lower, the bottom of one accuracy band aligns with the top of the next and so on. This is because, for steps of -21 dB, -24 dB, -27 dB, -30 dB and -33 dB, so little resolution is available that the digital input code to the DAC can decrease by only one LSB per 3 dB step. Although the accuracy bands align they do not overlap, hence the attenuator, though inaccurate at the low end, is monotonic over the entire range in 3 dB steps.

This should not be entirely unexpected, since we are dealing with a DAC having 6-bit resolution *and* 6-bit accuracy. Such a DAC is inherently monotonic, whether the output is expressed in LSBs or dB. Note that if the DAC were only 5-bit accurate, the error bands—which align in Figure 16.4—could overlap, allowing possible non-monotonic operation at the high attenuation levels. Under those conditions, the attenuator would only be monotonic in 3 dB steps down to a maximum of -21 dB.

3) From Table 16.1, at the 0-dB setting, the range of possible output attenuations is from approximately -0.07 dB to -0.2 dB. This corresponds to the output signal level being (worst case) approximately 2.3% below the input signal level. Note also that the errors are all negative, which means that the circuit is always acting as an attenuator. The reason for this is that at maximum gain (minimum attenuation), in the all-1's condition, the nom-

inal gain is $1 - 2^{-n}$, which is less than unity. The dB output of the attenuator, using a linear 6-Bit DAC, is (where N is the integer value of the binary number)

$$20 \log_{10} \frac{V_{OUT}}{V_{IN}} = 20 \log_{10} \frac{N}{2^n} \quad (16.5)$$

Since $N/2^n$ will always be less than unity, the output signal level will always be less than the input signal level.

4) Only 12 codes out of a possible 63 (again neglecting the all-zeros code) are required for 3-dB spacing, using a 6-bit DAC. The small number of codes employed suggests the possibility of reducing the number of bits in the digital input word from 6 to 4 and using an internal lookup table to select the codes for the R-2R ladder.

These observations for a linear 6-bit DAC performing an attenuation function with equal dB step-sizes are central to an understanding of LOGDAC behavior. The design philosophy behind the LOGDAC is to use an R-2R ladder network of sufficient resolution, concomitant with the required dynamic range, and sufficient accuracy to allow monotonic operation over a desired dynamic range, using defined step sizes. It will be instructive to look at the main features of three different commercially available LOGDACs and to look more closely at the specifications of a representative one, the Analog Devices AD7111, for an interpretation of the important d-c specifications published in the data sheets.

	AD7111	AD7115	AD7118
Dynamic Range	0 to -88.5 dB	0 to -19.9 dB	0 to -85.5 dB
Step Size	0.375 dB	0.1 dB	1.5 dB
Gain Error	± 0.1 dB	± 0.1 dB	Not separately specified; included in accuracy specs.
Digital Input	8-bit parallel	2 1/2 BCD Coding	6-bit parallel

Both the AD7111 and the AD7118, based upon a 17-bit-resolution R-2R ladder network, have dynamic ranges comparable with 16-bit-linear DACs, but they require only 8- and 6- bits, respectively, of digital input control. To simplify the internal code-conversion logic, which drives the R-2R ladder switches, a scheme is used that differs slightly from the basic one-to-one lookup table mentioned in the 6-bit linear DAC example.² However, from the user's point of view, the internal logic design of the devices is likely to be of little consequence.

The AD7115 high gain-resolution attenuator accepts a 2 1/2-digit binary-coded-decimal (BCD) word to provide a dynamic range of 0 to -19.9 dB, in

²A full explanation of the actual scheme can be found in Hynes, Michael J. and Burton, D. Philip, "A CMOS Digitally Controlled Audio Attenuator for Hi-Fi Systems," *IEEE Journal of Solid-State Circuits*, Vol. SC-16, Feb. 1981, pp. 15-20.

0.1-dB steps. The AD7115's attenuation is produced by a 12-bit-resolution R-2R ladder network, which is accurate to 12 bits. This device uses a lookup-table technique to convert its BCD input to a 12-bit-wide binary word, which controls the gain.

16.2.1 WHAT THE SPECIFICATIONS MEAN

As mentioned earlier, the AD7111 provides a dynamic range of 0 to -88.5 dB, controlled by an 8-bit-wide digital input word. Nominal gain resolution, or step size, for this device is 0.375 dB. A tabulation of ideal attenuation, in dB, as a function of input code, is shown in Table 16.2. The four more-significant bits (D7 through D4, corresponding to the MSB through bit 4) select the rows; the four less-significant bits choose the columns. The all-zeros condition represents 0 dB of attenuation, *i.e.*, unity gain.* The input codes selected by the last row, 1111XXXX (FX_H) represent a mute condition: no throughput, or infinite attenuation; this is equivalent to an all-zero code in a conventional linear DAC, such as the 6-bit device in the earlier example.

D3-D0																	
D7-D4		0000	0001	0010	0011	0100	0101	0110	0111	1000	1001	1010	1011	1100	1101	1110	1111
0000	0.0	0.375	0.75	1.125	1.5	1.875	2.25	2.625	3.0	3.375	3.75	4.125	4.5	4.875	5.25	5.625	
0001	6.0	6.375	6.75	7.125	7.5	7.875	8.25	8.625	9.0	9.375	9.75	10.125	10.5	10.875	11.25	11.625	
0010	12.0	12.375	12.75	13.125	13.5	13.875	14.25	14.625	15.0	15.375	15.75	16.125	16.5	16.875	17.25	17.625	
0011	18.0	18.375	18.75	19.125	19.5	19.875	20.25	20.625	21.0	21.375	21.75	22.125	22.5	22.875	23.25	23.625	
0100	24.0	24.375	24.75	25.125	25.5	25.875	26.25	26.625	27.0	27.375	27.75	28.125	28.5	28.875	29.25	29.625	
0101	30.0	30.375	30.75	31.125	31.5	31.875	32.25	32.625	33.0	33.375	33.75	34.125	34.5	34.875	35.25	35.625	
0110	36.0	36.375	36.75	37.125	37.5	37.875	38.25	38.625	39.0	39.375	39.75	40.125	40.5	40.875	41.25	41.625	
0111	42.0	42.375	42.75	43.125	43.5	43.875	44.25	44.625	45.0	45.375	45.75	46.125	46.5	46.875	47.25	47.625	
1000	48.0	48.375	48.75	49.125	49.5	49.875	50.25	50.625	51.0	51.375	51.75	52.125	52.5	52.875	53.25	53.625	
1001	54.0	54.375	54.75	55.125	55.5	55.875	56.25	56.625	57.0	57.375	57.75	58.125	58.5	58.875	59.25	59.625	
1010	60.0	60.375	60.75	61.125	61.5	61.875	62.25	62.625	63.0	63.375	63.75	64.125	64.5	64.875	65.25	65.625	
1011	66.0	66.375	66.75	67.125	67.5	67.875	68.25	68.625	69.0	69.375	69.75	70.125	70.5	70.875	71.25	71.625	
1100	72.0	72.375	72.75	73.125	73.5	73.875	74.25	74.625	75.0	75.375	75.75	76.125	76.5	76.875	77.25	77.625	
1101	78.0	78.375	78.75	79.125	79.5	79.875	80.25	80.625	81.0	81.375	81.75	82.125	82.5	82.875	83.25	83.625	
1110	84.0	84.375	84.75	85.125	85.5	85.875	86.25	86.625	87.0	87.375	87.75	88.125	88.5	88.875	89.25	89.625	
1111	MUTE	MUTE	MUTE	MUTE	MUTE	MUTE	MUTE	MUTE	MUTE	MUTE	MUTE	MUTE	MUTE	MUTE	MUTE	MUTE	MUTE

Table 16.2. Theoretical attenuation in dB vs. digital code for AD7111 LOGDAC.

There are some significant differences between the ideal attenuator, of Table 16.2, and real attenuators. For example, the last three input codes of the AD7111 (1110 1101, 1110 1110, and 1110 111), representing attenuations from 88.875 dB to 89.625 dB, are not recommended for use because of possible nonmonotonic behavior of the associated output levels. The code corresponding to the greatest usable attenuation, short of muting, is 1110 1100 (EC_H) corresponding to an attenuation level of 88.5 dB. Other differences will be found in a comparison between the ideal performance of Table 16.2 and the device specifications, in Table 16.3.

Before performing the comparisons, it may be instructive to consider what resolution and accuracy would be required for an R-2R ladder network to

*The all-zeros condition for LOGDACs corresponds essentially to all-ones for linear DACs, since minimum attenuation means maximum gain.

meet the ideal performance figures in practice. Assume that the ideal accuracy of any step (or level) in Table 16.2 is to within one-half the step size, i.e., ± 0.17 dB. It can be shown that an R-2R ladder with 17-bit resolution and 22-bit accuracy is required. Such DACs are difficult to procure, especially if one seeks them in the form of 16-pin plastic DIPs! It should therefore come as no surprise to find that the actual performance of practical, low-cost logarithmic DACs falls somewhat short of ideal.

Parameter	AD7111L/C/U GRADES		AD7111K/B/T GRADES		Units	Conditions/Comments
	$T_A = +25^\circ\text{C}$	$T_A = T_{\min}, T_{\max}$	$T_A = +25^\circ\text{C}$	$T_A = T_{\min}, T_{\max}$		
NOMINAL RESOLUTION	0.375	0.375	0.375	0.375	dB	
ACCURACY RELATIVE TO 0dB ATTENUATION						
0.375dB Steps:						
Accuracy $\leq \pm 0.17$ dB	0 to 36	0 to 36	0 to 30	0 to 30	dB min	Guaranteed attenuation ranges for specified step sizes
Monotonic	0 to 54	0 to 54	0 to 48	0 to 48	dB min	
0.75dB Steps:						
Accuracy $\leq \pm 0.35$ dB	0 to 48	0 to 42	0 to 42	0 to 36	dB min	Full Range is from 0 to 88.5dB
Monotonic	0 to 72	0 to 66	0 to 72	0 to 60	dB min	
1.5dB Steps:						
Accuracy $\leq \pm 0.7$ dB	0 to 54	0 to 48	0 to 48	0 to 42	dB min	
Monotonic	Full Range	0 to 78	0 to 85.5	0 to 72	dB min	
3.0dB Steps:						
Accuracy $\leq \pm 1.4$ dB	0 to 66	0 to 54	0 to 60	0 to 48	dB min	
Monotonic	Full Range	Full Range	Full Range	Full Range		
6.0dB Steps:						
Accuracy $\leq \pm 2.7$ dB	0 to 72	0 to 60	0 to 60	0 to 48	dB min	
Monotonic	Full Range	Full Range	Full Range	Full Range		
GAIN ERROR	± 0.1	± 0.15	± 0.15	± 0.20	dB max	
V_{IN} INPUT RESISTANCE (PIN 15)	9/11/15	9/11/15	7/11/18	7/11/18	k Ω min/typ/max	
R_{FB} INPUT RESISTANCE (PIN 16)	9.3/11.5/15.7	9.3/11.5/15.7	7.3/11.5/18.8	7.3/11.5/18.8	k Ω min/typ/max	

Table 16.3. Accuracy specifications of a LOGDAC (AD7111).

Table 16.3 is the section of the Specifications page from the AD7111 Data Sheet pertaining to device accuracy. Other specifications, such as logic input levels and logic input timing for the data-input latches, are omitted here.

Note that the accuracy is specified as relative to 0 dB attenuation and not with respect to V_{IN} . The accuracy of the 0 dB setting is established by the gain-error specification of ± 0.1 dB at $+25^\circ\text{C}$ and ± 0.15 dB over the temperature range (L/C/U grades). This is similar to a linear DAC specification, in which accuracy and gain error are specified separately. For both LOGDACs and linear DACs, gain error is measured with full-scale (FS) output, i.e., at 0 dB attenuation for a LOGDAC and with all 1's on the digital inputs for a unipolar binary-coded DAC.

Gain error results from a mismatch between R_{FB} (the feedback resistance) and the R-2R ladder resistance; its effect in a LOGDAC is to produce a constant additive attenuation error in dB over the whole range. Since the gain error of CMOS multiplying DACs is normally less than 1%, the accuracy error contribution due to gain-error effects is typically less than 0.09 dB.

For the AD7111 and the AD7115, which have step sizes comparable to the gain-error contribution, it is especially desirable to separate the specifications of gain error at 0 dB—and device accuracy, relative to 0 dB attenuation. On

the other hand, for the AD7118, which has a nominal step size of 1.5 dB, much greater than any gain-error contribution, the accuracy is specified relative to V_{IN} , and gain-error is included in the accuracy specification. This is similar to a specification of total unadjusted error for a d/a converter.

For simplicity in the earlier discussion of the 6-bit DAC, gain error was not considered; the DAC was assumed to have negligible gain error (not an unwarranted assumption, at the 6-bit level). Thus, the attenuation data of Table 16.1 (and its graphical representation in Figure 16.4) is specified with respect to V_{IN} .

In Table 16.3, the accuracy specification over the full 0 to -88.5 -dB attenuation range is divided into various subgroups. Each subgroup specifies—for a given step size—two ranges of attenuation:

- the range for which the error is always less than the specified amount
- the range for which the device is always monotonic.

To illustrate the point, consider the first subgroup for the AD7111L/C/U grades at $+25^{\circ}\text{C}$; from 0 dB to 36 dB, the attenuation can be increased in 0.375-dB steps with a step accuracy, relative to 0 dB, of ± 0.17 dB. Increasing the attenuation beyond 36 dB, in 0.375 dB steps, can result in attenuation steps with accuracy bands wider than ± 0.17 dB relative to 0 dB. The associated monotonic range of 0 dB to 54 dB suggests that, from 36 dB to 54 dB, the accuracy of each increasing step is deteriorating, yet the output level will always respond in the proper direction to a 0.375 dB step increase or decrease. Within each subgroup, the step accuracy specification is chosen to be less than one half of the relevant step size.

If the attenuation is increased beyond 54 dB, in 0.375-dB steps, the accuracy bands of adjoining 0.375 dB steps can overlap each other, causing non-monotonic performance.

The difference between the V_{IN} and the R_{FB} input resistance specifications is worth noting. For a linear-coded DAC the feedback resistance, R_{FB} , and the R-2R ladder resistance, R , are made equal—or as nearly equal as possible. The equality was assumed for the linear 6-Bit DAC example and its importance can be easily seen from Figure 16.2 where:

$$\begin{aligned}
 V_{OUT} &= -I_{OUT} R_{FB} \\
 &= -\alpha I_{IN} R_{FB} \\
 &= -\alpha V_{IN} \frac{R_{FB}}{R}
 \end{aligned} \tag{16.6}$$

Equation (16.1) assumes that $R = R_{FB}$. It was noted earlier that, since α is always less than unity, because the gain at all-1's is not equal to full scale, an ideal ($R_{FB} = R$) 6-Bit linear DAC will produce an actual attenuation at its "0 dB" setting that can range from approximately -0.07 dB to -0.2 dB.

Equation (16.6) suggests that, by a deliberate increase in the ratio, R_{FB}/R , the gain can be increased, shifting the entire DAC transfer function towards 0 dB. With the right ratio, this corrects the negative skew of the 0-dB code and results in a distribution of the “0 dB” error band around 0 dB. The typical ratio of R_{FB} to R for Analog Devices LOGDACs is:

AD7111;	1.044 : 1
AD7115;	1.04 : 1
AD7118;	1.05 : 1

Since the all-1's code is nominally within 1 LSB of full scale, and the AD7111 and the AD7118 employ a 17-bit-resolution R-2R ladder, while the AD7115 is based upon a 12-bit-resolution R-2R ladder, one would expect the R_{FB} -to- R ratios for the AD7111 and the AD7118 to be much less than the equivalent ratio for the AD7115 if a similar internal decoding technique were used in all three. The ratios are actually greater because differing decoding techniques are used for the 17-bit and the 12-bit R-2R ladders.

16.3 APPLICATIONS OF LOGDACs

The LOGDAC is a multiplying d/a converter with gain that is exponentially related to the digital input. The analog input can be voltage or current, ac or dc, positive or negative in polarity. This flexibility suggests a variety of applications, including dB-programmable attenuators and amplifiers, logarithmic a/d converters, level-independent automatic gain control, wide-range programmable state-variable filters, digitally programmable oscillators, distortion generators, audio panners, etc. Dynamic range, at a given accuracy level, can be extended by using fixed blocks of attenuation with the programmable attenuation provided by the LOGDAC.

In this section, we will consider some basic configurations of LOGDACs, from which can be derived many circuits with wider ramifications. Treated here will be exponential attenuators and amplifiers, logarithmic a/d converters, and range extension for attenuators. Considerable additional information is available from manufacturers in the form of data sheets, application notes, and applications-engineering advice.

16.3.1 BASIC ATTENUATOR OR DIGITAL POTENTIOMETER

Figure 16.5a shows the basic connection for using the LOGDAC as a digitally controlled attenuator. The nominal response of this circuit is

$$V_{OUT} = -(10)^{-\frac{NS}{20}} V_{IN} \quad (16.7)$$

where S is the step size, in dB, and N is the integer value of the binary or BCD digital input, over the specified range for which device performance is valid.

Expressed as a digitally controllable gain for an analog signal, V_{IN} ,

$$\text{GAIN} = \frac{V_{OUT}}{V_{IN}} = - (10)^{-\frac{NS}{20}} \quad (16.8)$$

For example, if the step size is 0.1 dB, as in the AD7115, and $N = 100$, then $-V_{OUT}/V_{IN} = 10^{-1/2} = 0.316$. In general, the response is of the form shown in Figure 16.5b; equal increases of N result in equal (increasing attenuation or decreasing gain) ratios of output.





Figure 16.5. Application of LOGDAC as attenuator.

16.3.2 EXPONENTIAL AMPLIFIER OR POT WITH FEEDBACK


If the LOGDAC is used as the feedback element, as in Figure 16.6a, the output and input change sides of the equation; now the output must generate whatever voltage is necessary such that, when attenuated by NS dB, it produces a current equal to the input current. The result is that the polarity of the exponent becomes positive, and each step of N produces an increment of gain, starting from 1.00 (i.e., 0 dB) at $N = 0$:

$$V_{OUT} = - (10)^{+\frac{NS}{20}} V_{IN} \quad (16.9)$$

Using the same example as in the case of the attenuator, with a 0.1-dB step size and $N = 100$, then $-V_{OUT}/V_{IN} = 10^{1/2} = 3.16$. A typical response function is shown in Figure 16.6b. Again, equal changes in N result in equal (but increasing) ratios of output gain.



a. Block diagram.



b. Gain as a function of digital code (linear plot).

Figure 16.6. Application of LOGDAC as amplifier with digitally controlled gain in dB steps.

By using fixed gain or attenuation in conjunction with the circuits of Figures 16.5 and 16.6, the overall attenuation or gain can be programmed to start from a level different from 0 dB. For example, in the case of Figure 16.5, an associated gain (say) of 10 volts/volt in cascade with the LOGDAC and its amplifier will result in a maximum gain of 20 dB, occurring at the 0-dB code, with

the normal complement of S -dB steps. Thus, the dynamic range of gain remains the same, but its starting point is 20 dB higher. In such applications, it may be tempting to consider economizing on amplifiers by using series or parallel resistance with R_{FB} , but with an extra amplifier stage, you can avoid the consideration of matching and temperature-tracking problems that is necessary when the on-chip resistors are connected with external resistors.

Naturally, in all cases, with or without additional gain or attenuation, the limits to dynamic range presented by amplifier output bounds and input noise levels must be considered. Within these ranges, however, the LOGDACs, like all CMOS multiplying DACs, are able to handle a wide range of positive and negative analog signals with wide bandwidth and low distortion and noise—once the gain is digitally set.

16.3.3 LOGARITHMIC A/D CONVERTER

The LOGDAC can be used to replace the linear DAC in a successive-approximation or tracking-type converter. In each case, the conversion loop chooses a code that reduces the unbalance to within 1 LSB of the analog value; when that state is reached, the digital input to the LOGDAC is proportional to the log of the input, that is

$$V_{IN} = -(10)^{-\frac{SN}{20}} V_{REF}$$

Therefore, the code represents the digital value,

$$N = -\frac{20}{S} \log_{10} \left(-\frac{V_{IN}}{V_{REF}} \right) \quad (16.10)$$




Figure 16.7. Application of LOGDAC as logarithmic (log-ratio) a/d converter.

Again, within the system limitations of dynamic range and accuracy, this type of a/d conversion allows a wide dynamic range of input to be compressed into a relatively small digital resolution. Advantages might include single-precision 8-bit bus operation and fast floating-point conversion without the need for programmable-gain amplifiers and repeated conversions (logarithmic conversion is inherently a limited floating-point operation).

Logarithmic DACs can also be used in automatic gain-control loops, which function essentially as tracking a/d converters.³

16.3.4 RANGE EXTENSION

A popular requirement is to extend the attenuation range. The purpose may be to extend a limited range, as in the case of the AD7115—perhaps to double its inherent range of 20 dB to 40 dB, in 0.1-dB steps. Or one may wish to extend the most-accurate range—for example, to double the 0 to 36-dB range of the AD7111 to 0 to 72 dB, while maintaining the same basic accuracy specification for 0.375-dB steps.

To achieve this, one or more stages of switched fixed-gain attenuation or programmable-gain amplification are made available in cascade with the LOGDAC to provide the range of total attenuation required, when combined with the attenuation range of the LOGDAC. At low levels of system attenuation (within the basic LOGDAC's accurate attenuation range), the precision attenuator is switched out of the circuit and contributes 0 dB attenuation to the signal path, all attenuation being controlled by the LOGDAC. At some user-defined input code, the precision attenuator switches in an attenuation equal to that defined by the input code, while the LOGDAC code is reset to all 0's to give 0 dB attenuation. As input codes increase further, to call for increased attenuation, the LOGDAC again controls the incremental attenuation as required.

This technique extends the required range of accurate performance by the amount introduced by the precision attenuator, assuming that it introduces no additional errors itself. For example, the AD7118 (L/C/U grades), with a step size of 1.5 dB, has a step accuracy relative to V_{IN} of ± 0.35 dB, from 0 dB to -30 dB, increasing to ± 0.7 dB from -31.5 dB to -48 dB. A table of ideal attenuation vs. input code for the AD7118 is shown in Table 16.4.

To extend the specified dynamic range by, say, 24 dB, the input code change from $N = 15$ to $N = 16$, i.e., 001111 ($0F_H$) to 010000 (10_H), corresponding to an attenuation level change from 22.5 dB to 24 dB, provides an opportunity to implement the scheme. If the AD7118 is driven only by the 4 less significant bits, and a range switch (switching in a 24-dB precision attenuator on "1") is driven by the next largest bit, the desired performance will be obtained for

³John Wynne, "Level-Independent Automatic Gain Control," *Analog Dialogue* 17-1 (1983), pp. 16-17.

N	Digital Input		Attenuation dB	N	Digital Input		Attenuation
	D5	D0					
0	00 00 00		0.0	31	01 11 11		46.5
1	00 00 01		1.5	32	10 00 00		48.0
2	00 00 10		3.0	33	10 00 01		49.5
3	00 00 11		4.5	34	10 00 10		51.0
4	00 01 00		6.0	35	10 00 11		52.5
5	00 01 01		7.5	36	10 01 00		54.0
6	00 01 10		9.0	37	10 01 01		55.5
7	00 01 11		10.5	38	10 01 10		57.0
8	00 10 00		12.0	39	10 01 11		58.5
9	00 10 01		13.5	40	10 10 00		60.0
10	00 10 10		15.0	41	10 10 01		61.5
11	00 10 11		16.5	42	10 10 10		63.0
12	00 11 00		18.0	43	10 10 11		64.5
13	00 11 01		19.5	44	10 11 00		66.0
14	00 11 10		21.0	45	10 11 01		67.5
15	00 11 11		22.5	46	10 11 10		69.0
16	01 00 00		24.0	47	10 11 11		70.5
17	01 00 01		25.5	48	11 00 00		72.0
18	01 00 10		27.0	49	11 00 01		73.5
19	01 00 11		28.5	50	11 00 10		75.0
20	01 01 00		30.0	51	11 00 11		76.5
21	01 01 01		31.5	52	11 01 00		78.0
22	01 01 10		33.0	53	11 01 01		79.5
23	01 01 11		34.5	54	11 01 10		81.0
24	01 10 00		36.0	55	11 01 11		82.5
25	01 10 01		37.5	56	11 10 00		84.0
26	01 10 10		39.0	57	11 10 01		85.5
27	01 10 11		40.5	58	11 10 10		87.0
28	01 11 00		42.0	59	11 10 11		88.5
29	01 11 01		43.5	60	11 11 XX		∞
30	01 11 10		45.0				

NOTES

X = 1 or 0. Output is fully muted for $N \geq 60$ Monotonic operation is not guaranteed for $N = 58, 59$

Table 16.4. Ideal attenuation vs. input code for the AD7118.

precision attenuations of up to a maximum possible of 46.5 dB, since—at the next code change—the input code to the AD7118 is reset to all 0's and the precision 24 dB attenuator is switched out of the signal path. Maximum attenuation range has been traded for increased accuracy. Figure 16.8a illustrates the basic scheme and (b) illustrates a circuit that will implement the switched attenuator.

Differing circuit configurations can be used, depending upon the performance required (i.e., accuracy and total range), employing more or less of the attenuation range of the LOGDAC, one or more steps of switched precision attenuation, and magnitude of the external attenuation. For example, if a wider dynamic range desired, a dedicated control bit for the precision attenu-

Chapter Seventeen

High-Resolution Data Conversion

This chapter deals with the problems, challenges, and opportunities afforded by conversion devices and techniques that result in high resolutions and accuracies. High-resolution converters are a special breed, and—with the possible exception of very high-speed converters—they have been the most likely candidates for custom design to meet unusually demanding system requirements. However, converter manufacturers are rising to the challenge, and a growing number of high-resolution-and-accuracy converters are becoming available at competitive prices.

Converters meeting this definition guarantee resolutions beyond 17 bits or at least 16 bits of resolution with comparable accuracy. The principal requirement of a 16-bit-accurate converter is an integral nonlinearity specification of $\pm \frac{1}{2}$ LSB, or 7.6 parts per million (ppm). Data converters with resolutions of 17 bits and beyond, even without a commensurate improvement in accuracy, can still offer the wide dynamic range required for many signal processing applications.

The definition of “high resolution” has been shifting upward as data-converter design and IC process technology improve. High resolution, a dozen years ago, was considered to be 12-bits or greater. Sixteen bits of resolution was a laboratory curiosity for designers with undeniable need to brave the world of microvolts. Today, 12-bit resolution and accuracy are available in low-cost ICs (despite gloomy predictions by some hardshelled converter design experts), and 16-bit converters have inherited the mantle of “high resolution”.

At this writing, although 16-bit monolithic, hybrid and modular converters fill pages of manufacturers’ data books, few offer true 16-bit performance. Most are “16/14” bit converters with only fourteen bits of accuracy. The non-

monotonic (or missed codes) behavior of many of these devices falls so far short of 16 bits of real resolution that one might wonder why users want the extra pins. One reason may be the hope that—some day—a device with real 16-bit performance will evolve that can be plugged into the same socket; another is that the low cost of many of these devices has led to their use in applications where designers, not needing the sometimes illusory dynamic range that “16-bit” devices seem to offer, use them to get *true 12-bit* performance with less need to trim the converter.

17.1 APPLICATIONS

Applications for high resolution thus fall into two basic categories: high accuracy and wide dynamic range. High-accuracy applications include: automatic test equipment (ATE), calibration standards, instrumentation, process control, and precision positioning. High-accuracy applications usually involve low bandwidth; the key specifications are linearity, offset, and gain error. Designs calling for converters with wide dynamic range, without especial regard for dc or gain accuracy, include such applications as digital audio and waveform reconstruction; they require low total harmonic distortion (THD) and a high signal-to-noise ratio (S/N).

A major use for high-resolution converters is in testing lower-resolution converters. To test an a/d converter or d/a converter accurately requires another converter with at least two additional bits of resolution and accuracy. For example, in DAC testing, a digital code is fed to both the reference DAC and the converter being tested, and the outputs from both converters are connected to a differential amplifier. Its amplified voltage output, the difference between the output of the reference DAC and the converter being tested, corresponds to the error of the DAC being tested, since the reference converter is “known” to be accurate. As a rule, accurate 16-bit DACs are used to test 12-bit DACs, and accurate (or calibrated) 18-bit-or-better DACs are used to test 16-bit DACs.

Semiconductor equipment manufacturers use high-resolution-converters in computer-based IC processing machines. An example of this is a beam steering application in electron-beam lithography. To shrink device feature sizes, say from 3 to 1 microns, calls for a threefold improvement in the accuracy of the components used in the processing equipment. An electron beam is used, either to make the mask for standard photolithographic processing or—directly steered across the surface of the wafer—to write the appropriate patterns on the die.

A computer controls the shape, intensity, focus and deflection of the electron beam as it runs the gauntlet of lenses and electrostatic and electromagnetic deflection stages (Figure 17.1). High-resolution d/a converters provide analog signals to control the beam's deflection, in the same way as in a CRT. To resolve 1-micron (1- μm) features to 0.25 microns, for a die size of 5 millimeters,

the beam must be steered with a resolution and accuracy of better than 1:20,000. Since there will be additional error sources, the resolution of a 16-bit (or better) converter, 1:65,536, is needed. The wafer is positioned by an X-Y positioning table.




Figure 17.1. Electron-beam lithography uses a shaped electron beam to write patterns on integrated-circuit wafers. High-resolution d/a converters control the deflection of the electron beam.

Similarly, in numerical control, resolution of the a/d and d/a converters used in positioning determine the tolerance of machined parts. Using a DAC with 16 bits of accuracy and resolution means that a machine tool can mill a three-foot (about one-meter) precision steel part to within about 0.6 mils (15 micrometers). To get a feeling for how precise this is, it's worth noting that the same metal part will expand about 13 micrometers for each 1°C increase in temperature, which calls for great effort to control the temperature of the part while it is being machined (or at least while the measurement is made), as well as high accuracy for the other electronic components in the system.

17.2 HIGH-RESOLUTION D/A CONVERTERS

17.2.1 TESTING DAC INTEGRAL AND DIFFERENTIAL NONLINEARITY

High-resolution converters do not readily lend themselves to high-speed testing using automatic test equipment. Testing can be semiautomated or automated, but a custom test fixture is usually necessary; it will provide an electrically quiet, thermally stable environment, as well as a test-instrumentation front end comprising components having the required degree of accuracy, stability, and timeliness of calibration.

Differential nonlinearity, as defined in Chapter 11, is the difference between

the actual analog output change when the digital input is increased by 1 bit, and the theoretical value of 1 LSB for the device being tested, i.e., $FSR/2^n$, where FSR is the full-scale voltage range (output reference) and n is the resolution in bits.

A simple way to measure differential nonlinearity is to connect a parallel-load down counter to the device under test (Figure 17.2). The inputs to the counter are first preset to the initial digital value (for example, 1000 . . . 00). When the clock goes low, the digital input is asynchronously loaded into the counter, and impressed on the converter. Then, on the rising edge of the clock pulse, the counter counts down by one, and thereby decrements the code presented to the converter (e.g., to 0111 . . . 11). On the next cycle, the counter is preset to the initial value again, and the cycle repeats.




Figure 17.2. A simple scheme for checking differential linearity of a DAC without high-precision equipment.

Ideally, the resulting analog output will be a square wave with an amplitude of 1 LSB at the clock frequency, biased at a dc level equal to the initial value. The deviation of the square-wave's amplitude from the expected value of 1 LSB is the differential nonlinearity for the device at that particular digital input transition. Since a 1-LSB change is small ($10\text{ V}/65,536 = 153\text{ }\mu\text{V}$), the signal will have to be amplified with the dc level removed. The dc level can be removed by a-c coupling or by taking the difference between the DAC's output and the output of a reference DAC, set at the same initial code; the former technique is simpler, but the latter has the advantage of also indicating the absolute error of the initial code.

For 18-bit converters, there are 2^n —or 262,144—codes. Such a large number of codes makes it impractical to test each individual combination, unless a semi-automated test setup—and plenty of time—is available. Even with an automated tester that required as little as 10 ms per test, the entire sequence would take almost one hour. Fortunately, it is not necessary to test all input combinations to characterize a high resolution device with no summation er-

rors*—a total of only n tests is required. If summation errors are small, then a total of $2n$ tests may be sufficient. These tests are usually performed at the major carries, such transitions as 1000 . . . 00 to 0111 . . . 11; 0100 . . . 00 to 0011 . . . 11; 0000 . . . 100 to 0000 . . . 011, etc., and, if necessary, at second major carries, such as 1100 . . . 00 to 1011 . . . 11.

Integral nonlinearity, INL, also referred to as *relative accuracy*, is the deviation of the actual converter output from a straight line drawn between the end points of the converter's output-vs.-input function. INL includes both bit errors and summation errors; it is difficult to measure, because the wide dynamic range of the signal requires finding the difference between two large numbers. A $6\frac{1}{2}$ -digit meter, accurate to 0.0002% (i.e., 2 ppm), would be necessary to measure the integral nonlinearity of a 16-bit converter.

For nonlinearity measurements, a better approach is to compare the output of the converter under test to a voltage having a known accuracy. The errors are then read with a null meter. This test requires a high-precision 18-bit d/a converter to measure a 16-bit-accurate converter. A precision divider, traceable to the National Bureau of Standards, must be used to test an 18-bit converter.

17.2.2 MAINTAINING HIGH RESOLUTION AND ACCURACY

The high price paid for 18 bits of resolution makes it economical, and sometimes mandatory, to perform fine-tuning to improve or maintain the accuracy by adding correction circuitry. To improve the linearity of an 18-bit converter from 16 to 18 bits of integral nonlinearity, for example, the integral linearity error can be measured at all codes. Then a corresponding correction signal for each code is added using a converter having lower resolution and a full scale value equal to only a few LSBs (viz., the maximum error) of the device to be corrected.

In such a scheme, a programmable read-only memory (PROM), addressed by a common input bus, stores the digital input to the correction converter. The drawback to this approach is the large amount of PROM memory required: $256K \times 8$ -bits (assuming 8 bits, or 256 levels, of correction). However, the memory requirements can be reduced, with essentially the same results, by correcting for only the major error sources, rather than every single code.

For a converter with no summation errors, the worst-case integral linearity error will be less than or equal to one-half the of the worst case differential

**Summation errors*: For a linear device, superposition holds. This means that, for (say) eighteen terms with individual values equal to either b_i or zero, the output of a DAC with no offset and exactly unity gain—which sums those terms—will ideally be equal to their algebraic sum, for any combination of b_i 's and zeros. The difference between the actual output and the mathematical sum is a *summation error*. It is due to a nonlinear summation, such as might be caused by an amplifier with distortion or a nonlinear feedback resistor. In defining summation error, it doesn't matter what the individual terms, b_i , are; they are not necessarily exactly equal to $V_{FSR}/2^i$; if they aren't, their deviations from the ideal are called *bit errors*.

linearity error. Therefore, if the summation errors and differential linearity errors of a converter are corrected, this will result in correction for integral linearity errors as well. Information about the size and location of summation errors for any given converter must be determined experimentally.

Many high resolution converters, instead of having a single ladder network, comprise several independent internal d/a converter stages of 4, 8 or 12-bits. For example, the Analog Devices DAC1146, an 18-bit d/a converter, has three stages, with independent resolutions of 4, 12, and 2 bits. The internal architecture can be designed to guarantee that these stages do not interact to produce interstage summation errors. Summation errors in the less-significant stages are suppressed with respect to full scale and so should be negligible.

The only summation errors that are greater than $\frac{1}{4}$ LSB (at 18 bits) occur in the four MSBs. Since any value of DAC output must include one of the sixteen bit-combinations, to correct for the summation errors in these four MSBs requires only 16 bytes of memory. To ensure correction for interstage errors, the number of bytes can be increased to 32 or 64—still many orders of magnitude less than a full correction scheme. The only remaining task is to correct for the differential nonlinearity errors in the less-significant bits, if necessary. The net result is equivalent to an 18-bit-accurate converter.

Figure 17.3 shows the principles of a semiautomated calibration scheme to make a true 18-bit d/a converter. In this example, the major-carry and summation-error corrections for the first six bit combinations are established in RAM, and the corrections for the lower-order bits are set by means of adjustable trim resistors. The correction codes are different for each DAC used, and they should be expected to change over time and temperature. For this reason, the long-term correction memory should be EEPROM or non-volatile RAM.

The correction circuitry, shown at the lower left, consists of an 8-bit DAC, driven by the correction RAM, and the trim potentiometers for the lower bits. The circuitry employed to determine the errors and implement the correction, shown above and to the right, includes an 8-bit counter, driven by a clock and gated by a differential comparator with a sampling capacitor on one input, a precision instrumentation amplifier with a gain of 1000, to amplify the error, and a 12-bit DAC, to hold the bit level while two adjacent bits are compared.

Here's the principle: For each of the 64 major carries involving the first 6 bits—for example, from **11 0100 0000 0000 0000** to **11 0011 1111 1111 1111**—the DAC's output difference should be exactly 1 LSB. To check and adjust it, a 12-bit DAC is latched at the upper value, and its output is subtracted from that of the DAC under test (DUT), first at the upper value, then at the lower value, with the difference amplified by 1,000. The first of the differences is stored on a capacitor; the second is added to a current equal to 1,000 LSBs

the *on* bits to generate the appropriate correction code for the d/a converter. For example, with the desired output code at 3/4 full scale, i.e., the two MSBs *on* (11 0000 00 00), the processor adds the correction terms for bits 1 and 2 and loads the sum into the 8-bit correction DAC, which adds the analog value of the correction to the output.

Because the LSB is so small in high-resolution converters, few potential contributions to error can be neglected. Once a converter has been calibrated, the principal source of error is the variation of parameters with temperature. When a high-resolution-and-accuracy converter is used in an environment having wide ranges of temperature, many of the specifications may suffer serious degradation. This is easy to understand when you recognize that, for an 18-bit converter, 1 part-per-million is a large change.

For example, a DAC's worst-case monotonic temperature range or an ADC's worst-case no-missing-codes range is computed by subtracting the specified differential nonlinearity (DNL) from 1 LSB and dividing the result by the DNL temperature coefficient. This gives the minimum temperature deviation around the specified temperature (usually 25°C) at which the converter will remain monotonic or exhibit no missing codes. A 16-bit converter with 1/2-LSB DNL (7.5 ppm) and DNL tempco of 1 ppm/°C offers a monotonic temperature range of $\pm 7.5^\circ\text{C}$ around the initial 25°C operating temperature. An 18-bit d/a converter with DNL of 1/2 LSB (1.9 ppm) and a drift tempco of 0.4ppm/°C will remain 18-bits monotonic for a range of only $\pm 4.7^\circ\text{C}$ around 25°C or from 20.3°C to 29.7°C. These are minimum monotonic temperature ranges; the conservative assumption is that worst-case trims and worst-case tempcos occur at the same code. However, this is not necessarily the case; typical performance can be considerably better.

Differential nonlinearity is only one component of a converter's accuracy. Absolute accuracy error, at any output value, consists of gain error, zero error and integral linearity error. Gain and zero errors can be trimmed to zero at a specific ambient temperature. As the converter temperature deviates from the calibration temperature, errors accumulate and degrade the accuracy of the converter. For example, a unipolar 16-bit d/a converter, trimmed to initial accuracy of 1/2 LSB, with drift specifications such as $\pm 10\text{ppm}/^\circ\text{C}$ gain drift, and $\pm 0.5\text{ppm}/^\circ\text{C}$ offset drift, will have a maximum full-scale deviation from the initial accuracy of $\pm 10.5\text{ppm}/^\circ\text{C}$. (If the drift of the reference is not included in the gain drift specification, it must be added on.) The total $\pm 10.5\text{ppm}/^\circ\text{C}$ drift means that the converter is 16-bits accurate (± 1 LSB) for approximately $\pm 1.4^\circ\text{C}$ around the specified operating range.

However, one must retain the system perspective. The converter is not used alone, and a precision measurement system will usually require that the ambient temperature variation be kept small, in order that the performance of all analog system elements, including the converter, suffer as little degradation as possible. When necessary, analog and digital temperature-compensation

techniques may be used to reduce temperature-related drifts to some degree. In addition, one should consider that, in a great many cases, high-resolution converters are purchased in order to get improved resolution and accuracy for lower resolutions over wider temperature ranges; for these applications, *full* accuracy-and-resolution are not needed over temperature.

Analog compensation involves sensing the changes with temperature in the major error sources and injecting compensating currents or voltages into the circuit. Sources of errors, within the internal circuitry, include the reference (which has a major impact on determining the full-scale output, and thereby the gain), the offset (zero), and the bit current weights. Since the MSBs have the largest current weights, they are the primary determinants of errors in the resistor ladder network.

Digital correction involves measuring the deviation of the converter's output, using a known reference source, and either adding an analog current or voltage to provide the correct value, or summing appropriate corrections digitally so that the code which determines the output incorporates the correction. There are a number of ways to provide this. For traceable standards a system can provide an autocalibration cycle on a periodic basis.

The 16-bit d/a converter shown in Figure 17.4 is connected to two 8-bit d/a converters, which can be programmed to provide automatic calibration of offset and gain errors. During a calibration cycle, the output of the precision DAC is compared with the correct value for each calibration point, and comparator measures the errors. Offset error is measured as the difference between the the converter output at zero and "ground". After zeroing, gain error is measured as the the difference between actual output with an all-1's input code and (FS - 1 LSB) output. Once measured, the correction codes

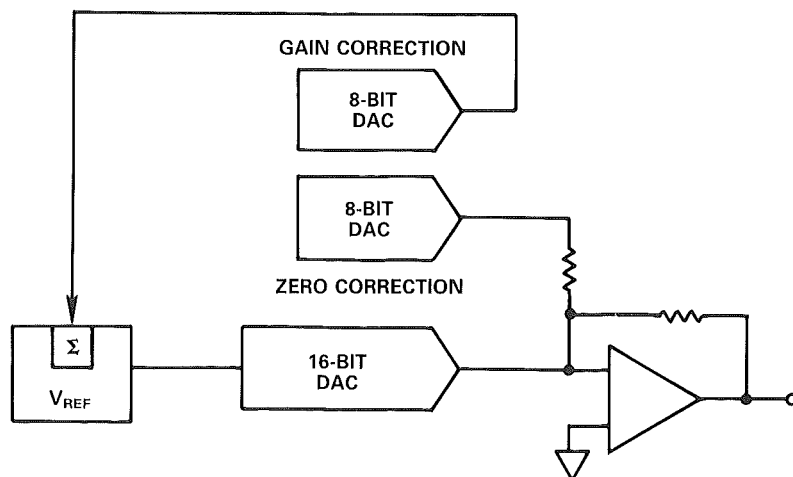


Figure 17.4. Use of auxiliary DACs for offset and gain correction of a high-resolution DAC.

are applied to the two correction DACs to compensate for the errors. The offset DAC injects the required current into the summing node of the output current-to-voltage amplifier. The gain DAC adjusts the reference voltage.

With a dedicated microprocessor, linearity—as well as gain and offset—can be corrected automatically. A typical scheme, which can be implemented as a system or a complex module, is shown in Figure 17.5. The elements of the scheme include a 16-bit d/a converter, a stable, temperature-compensated voltage reference, an error-measuring circuit, and a microcomputer to calculate correction factors for offset, gain and linearity. To trim the converter (initially within a few LSBs) it is adjusted, using the reference, offset, gain, and manual adjustments of (say) the four most-significant bits. The accuracy of the reference over temperature and time ultimately determines the accuracy of the d/a converter after calibration, since the reference is used as the “standard” to compensate the converter.




Figure 17.5. D/A converter with gain, offset, and linearity correction, using a dedicated microprocessor.

The error-measurement circuit compares a series of voltage pairs to derive the compensation factors. For example, the circuit compensates offset errors by measuring the zero output of the converter against ground, gain errors by comparing the DAC's full-range output (full scale—1 LSB) and a similar fraction of the reference, and nonlinearity errors by conducting a linearity test on the four MSBs, in much the same manner as described in Figure 17.3, except that an integrating ADC may be used to determine the error digitally, instead of the tracking measurement and conversion scheme. The correction factors are then stored in RAM and used as inputs to the linearity/offset-trim, and gain-trim d/a converters.

Some applications use high-resolution converters as programmable voltage or current sources. Since many high resolution d/a converters produce a current output, which can be applied to the inverting input terminal of an external op amp, with feedback via an on-chip application resistor—matched to other resistors in the device—the user can select from a universe of operational amplifiers to match a particular application's requirements. For example, a low-drift amplifier, with an inside-the-loop booster follower, can provide a large output drive for a programmable power supply.

It is also possible for some current-output DAC types to drive load resistance directly, with the voltage developed buffered by a follower op amp, in order to avoid the inherent sign inversion of inverting-amplifier configurations. However, a high-level inverting amplifier, used as an I-to-V converter, will make it possible to avoid voltage-compliance and summation errors. Voltage compliance is the maximum voltage that can appear at the current output terminal while maintaining a specified linearity. While some converters, such as the 10-bit AD561, have large compliance voltage ranges (e.g., -2 V to $+10\text{ V}$), others have low compliance voltage ranges; for example, the typical compliance for the DAC1146 is $\pm 500\text{ mV}$. Exceeding that specification leads to linearity errors.

Summation (integral-linearity) errors may occur in a voltage-output converter as the analog output increases from zero to full scale. The power dissipated by the feedback resistor increases, and the resistor heats up. This causes a change in the resistance value—and a corresponding change along the transfer function. In addition, if an external resistor is used with the internal feedback resistor for a non-standard gain, the external resistor will not track with internal resistors, resulting in larger than expected gain tempcos.

17.2.3 DYNAMIC APPLICATIONS OF DACs

Important Specifications

Dynamic applications take advantage of the wide dynamic range inherent in a high resolution converter. A 16-bit DAC, for example, offers 96 dB of dynamic range, versus the 72-dB range of a 12-bit converter. Although a number

of monolithic ICs do offer 16-bit resolution, the lower accuracy, 14-bits, means reduced signal-to-noise ratio, since differential linearity errors, especially around zero, reduce the signal-to-noise ratio of the converter.

In waveform reconstruction, the basic goal is to reconstruct the waveform as closely as possible to its original form. Specifications that impact on this accuracy include such results-oriented specs as total harmonic distortion, intermodulation distortion, noise, limited dynamic range, settling time, and aliasing. Users of high-resolution data converters in waveform reconstruction applications are generally not concerned with differential nonlinearity, as such, or such traditional precision-dc-measurement specifications as dc offset and gain.

Dynamic range is the ratio of the maximum output signal to the smallest output signal the converter can produce (1 LSB), expressed logarithmically, in decibels ($\text{dB} = 20 \log_{10}(\text{ratio})$). For an N-bit converter, the ratio is theoretically very nearly equal to 2^N (in dB, $20 N \log_{10}(2) = 6.02 N$). However, this theoretical value is degraded by converter noise and inaccuracies in the LSB weight.

The *signal-to-noise ratio* of a perfect digital waveform-reconstruction system is given by the ratio of the full-scale rms input to the rms quantization error.

For sine waves, the rms value is $1/2$ the peak-to-peak value, divided by $\sqrt{2}$; for an N-bit system, the full-scale rms value is thus $2^{(N-1)}/\sqrt{2}$. The quantization error, an expression of the fact that each quantized level represents an uncertainty as to the actual value of points in its neighborhood, is determined by the difference between a linear response (the expected output) and the characteristic analog-to-digital staircase function (the actual ideal digital output); the shape is a sawtooth oscillating once per LSB (or quantization interval, Q) between $\pm 1/2$ LSB (i.e., $\pm Q/2$). The rms value of this sawtooth wave is the peak output divided by $\sqrt{3}$, or: $(Q/2)/\sqrt{3}$. Thus, the signal-to-noise ratio is

$$S/N = \frac{2^{(N-1)} \sqrt{3}}{\sqrt{2} \cdot 1/2} = 2^N \frac{\sqrt{3}}{\sqrt{2}} = 1.225 \times 2^N \quad (17.1)$$

Expressed in dB,

$$\text{SNR} = 20(\log 1.225 + N \log 2) = 1.76 + 6.02 N \text{ dB} \quad (17.2)$$

Thus, the resolution and quantization level establish a noise floor; system noise and other sources of random noise will decrease the signal to-noise ratio.

17.2.4 TESTING HARMONIC DISTORTION OF DACs

Total harmonic distortion (THD) is the difference between an ideal sine wave and its reconstructed version using an a/d and a d/a converter. THD is the

ratio of the square-root-of-the-sum-of-the-squares of the RMS values of the harmonics to the RMS value of the fundamental. For a converter with a finite number of digital inputs, N , and associated output voltages, THD is customarily calculated from the formula:

$$\text{THD} = \frac{\text{RMS error}}{\text{RMS signal}} = \frac{\sqrt{\frac{1}{N} \sum_{i=1}^N [E_L(i) + E_Q(i)]^2}}{E_{\text{RMS}}} \quad (17.3)$$

Where $E_L(i)$ is the linearity error and $E_Q(i)$ is the quantization error of the converter at the sampling point i . Intermodulation distortion is caused by the additional error produced when the ideal output is composed of two sine waves of different frequency.

A computer-generated sine wave stored in a programmable read-only memory (PROM) can be used to demonstrate and test harmonic distortion. The PROM in Figure 17.6 contains one cycle of a computer-generated sine wave. A complete cycle is generated by counting through all address values; the frequency depends on the clock rate. If there are N values stored in the PROM to form a complete cycle, the sine-wave frequency will be C/N , where C is the clock rate. If a constant clock frequency is used, N must be reduced by skipping data points in order to increase frequency.

Thus, in the scheme of Figure 17.6, if the adder is set to increase the PROM address by one on each count, 4,096 inputs will be presented to the converter




Figure 17.6. Testing harmonic distortion as a function of frequency.

on each cycle of the PROM, for a very close approximation to the sine wave. If the adder is set to increase the PROM address by 1,024 on each count, four inputs will be presented on each cycle (an adequate number of samples). Frequency-select switches program the adder with the number of codes that the converter should skip on each count. In this way, any of 2,048 discrete frequencies between 12 Hz and 25 kHz (with corresponding count densities) can be generated by a constant 50-kHz sampling rate.

This series of quasi-sinusoidal digital codes is fed to the converter to generate staircase approximations of an analog sine wave. The DAC output is deglitched (see below) and displayed on a spectrum analyzer. Total harmonic distortion can be computed for each set of samples by comparing the amplitude of the fundamental frequency with the amplitudes of the harmonics.

17.2.5 IMPROVING DYNAMIC PERFORMANCE OF HIGH-RESOLUTION D/A CONVERTERS

Glitches and Deglitching

Because they distort the waveform, glitches in a d/a converter's output increase total harmonic distortion. Their spiky nature produces large distortion components at frequencies considerably greater than the fundamental, and their asymmetry introduces dc components, all of which makes glitches extremely difficult to filter with linear filters.

Glitches have two principal components, one introduced by charge coupling from the digital logic signals to the analog output, the other caused by asymmetry in timing as codes change. The former is less serious, because the pulses it produces, though variable, tend to change less in amplitude than the latter.

The asymmetry in timing, caused by skewed response in the up- and down-directions in both digital logic and analog switches, produces false intermediate codes, which can vary tremendously in amplitude.

For example, if the code is changing by 1 LSB at mid-scale, from 0111 . . . 11 to 1000 . . . 00, and the switches turn off (or receive their *off* drive signals) faster than they turn on (or receive their *on* drive signals), there will be a brief interval during which the output is seeking to respond to 0000 . . . 00; this will result in a large negative-going spike; on the other hand, the change to the next code, from 1000 . . . 00 to 1000 . . . 01, will only result in a less-than-1-LSB negative-going spike. Changes at other code transitions will result in spikes of differing amplitudes—largest at major carries—depending on which codes are turning off and which are turning on.

Converters using CMOS technology have additional glitch impulse introduced by the charge stored on the gate-to-source and gate-to-drain capacitance of CMOS switches during switching. The output capacitance, associated with the relatively large N-channel devices used for the DAC

switches, is highest when all the switches are on, and lowest when all the switches are off.

A deglitcher is essentially a sample-and-hold, which holds the signal from just prior to conversion until just after the signal settles. Although the deglitcher itself introduces a small glitch, it is small relative to the DAC's glitch, of constant amplitude, and independent of the digital code. Since it injects energy at the sampling frequency rather than the signal frequency, it introduces no harmonic distortion within the signal passband. The deglitcher can be band-limited with a linear filter to suppress distortion caused by slew-rate limiting of the output amplifier.

Most digital audio applications are stereo and require a right and left channel to be fed from a single fast, high-resolution 16-bit DAC, updated at 50 kHz. The L and R channel signals update the DAC alternately. A popular scheme for simultaneously updating both audio output channels—while eliminating the DAC glitch—employs a track-and-hold to store the right channel, while the left channel signal updates the DAC. Then two sample-and-holds (“deglitchers”), one on the DAC output, the other on the R track-and-hold, update the two output channels and hold the signal until the next joint update, while the R-channel track-and-hold is once again updated to start the new cycle.

The deglitcher response is band limited to eliminate distortion of the amplifier. A time constant of $3.4\text{ }\mu\text{s}$ ensures a full 20-kHz sine wave response without distortion. A 16-bit converter, track-and-hold, and deglitchers can produce an audio signal with distortion of less than 0.005%.

17.2.6 BIPOLAR-OUTPUT DAC PERFORMANCE

Dynamic signals, especially audio, are inherently bipolar. Performance around zero is important, since zero represents “silence” in a digital recording; many audio recordings consist of silences punctuated by sound. Performance around zero also determines the signal-to-noise ratio of the converter.

Signal polarity in d/a conversion is commonly handled in two basic ways—by offsetting (offset binary and two's complement) and by absolute value (sign-magnitude). Offsetting techniques employ a basic unipolar binary converter with an offset of $\frac{1}{2}$ scale. A sign-magnitude converter senses the polarity of the signal and switches the output circuit of the basic unipolar converter between the inverting and noninverting modes.

An offset-type bipolar converter exhibits its worst temperature performance around zero volts because it requires the MSB to track with the sum of all the less-significant bits, as well as with the bipolar offset resistor; it also suffers dynamically, because the offset zero is the point at which the major-carry glitch occurs. Large signals tend to mask error-induced noise, while noise buries small signals. A sign-magnitude converter exhibits its best performance around zero volts because all the bits are OFF, and transitions involve only the least-significant bit.

A simple unipolar binary converter is converted to sign-magnitude by the addition of a fast, linear, low-drift inverting amplifier with accurately set unity gain, a SPDT CMOS switch and an output buffer amplifier. The result is a maximum DNL drift of $\pm \frac{1}{2}$ ppm/°C for $\pm \frac{1}{8}$ -full-scale range and ± 1 or 2 ppm/°C over the full range.

17.2.7 EXTENDED RESOLUTION

A substantial number of high-resolution converters are converters with extended resolutions without commensurate increase in accuracy. These designs aim strictly at high (monotonic) resolution to obtain the wide dynamic range required in applications such as sonar, radar, optical signal processing, where the actual value of any data point is not required to great accuracy, but where the range of values is quite large, typically extending well beyond 80 dB. For example, high-quality photodiodes used in optical measurements produce linear current outputs over a 100-dB range, requiring at least a 17-bit a/d converter.

17.3 HIGH-RESOLUTION A/D CONVERTERS

High-resolution analog-to-digital conversion techniques are used to obtain high accuracy, improve system resolution, or boost the dynamic range. Goals for a system may specifically include lower noise and improved temperature stability.

17.3.1 FLOATING-POINT CONVERTERS

A floating-point converter improves dynamic range by acquiring data in two parts, usually employing a two-step approach. In a typical approach, a binary programmable-gain amplifier (PGA) scales the signal to the proper range, typically between 1 MSB and full scale; the input code to the PGA controls the exponent of 2 (i.e., $G = 2^X$). The scaled signal is then converted by an ADC whose output code forms the mantissa, while the digital signal used to set the gain prefaces it with the exponent of 2, viz., X. For example, a 12-bit converter with a 4-bit preface and the code, **0100 1100 0000 0000** would have an overall gain of 8×0.75 .

The dynamic range of a converter is the ratio of the full-scale input range to the smallest signal the converter can detect. With a floating-point converter, the smallest value of the LSB corresponds to the value of the LSB of the converter when the PGA is programmed for its highest gain.

$$\text{LSB} = \frac{\text{FSR}}{\text{GAIN} \times 2^N} \quad (17.4)$$

For a floating-point converter with full-scale range of 10 volts, 256 V/V maximum gain and a 12-bit ADC, this corresponds to:

$$\text{LSB} = 10\text{V}/(256 \times 4096)$$

or $9.5\mu\text{V}$. The dynamic range in decibels is based on the logarithm of the product of the amplifier gain and the converter's dynamic range, i.e.,

$$\begin{aligned}\text{dB} &= 20 \log (2^{12} \times 2^8) \\ &= 20 \log 2^{20} = 120 \text{ dB}\end{aligned}$$

This corresponds to a 12-bit resolution, with a dynamic range of 20 bits.

There are a number of different approaches to designing a floating-point converter. In some systems, software is used to set the gain of the the PGA. The signal is converted, and if the MSB is not equal to 1, the gain is increased by a number of binary steps (up to a specified maximum), essentially equal to the number of leading zeros. This maximizes the full-scale range of the conversion process and insures a wide dynamic range. Another approach is to use a logarithmic converter; it is fast, conceptually simple, maintains essentially constant percentage error over a major portion of the dynamic range, and uses a minimal bus width.

Another approach to floating-point converter design, with improved throughput rate, involves setting the gain of the PGA with a flash autoranger. The 20-bit converter diagrammed in Figure 17.7 consists of a pair of track-and-hold amplifiers, flash octave a/d converter (reference levels are at octave intervals, instead of equally spaced), a nine-range programmable gain amplifier, and a fast 12-bit a/d converter. The first track-and-hold amplifier holds the signal for the flash autoranger, which determines which binary quantum the input falls in, relative to full scale (e.g., $1/2$ to $1/4$, $1/64$ to $1/128$, etc.), and encodes the information as a 4-bit digital byte. Responding to it, the PGA




Figure 17.7. Floating-point a/d converter, employing a 12-bit ADC, with 20 bits of dynamic range represented by a 16-bit output word.

adjusts its gain to the appropriate level, and the track-and hold amplifier holds the amplified signal while the two-step a/d converter translates it to a binary number. A sixteen-bit latch holds the 4-bit output from the encoder and the 12-bit output from the converter.

There also exist modified floating-point converters, in which the prefix digits produce a selection of gains that are not in strict binary relation, for example, 1, 10, 100, or 1, 2, 5, 10, 20. An interesting example is a 12-bit a/d converter, with an amplifier having a maximum gain of 256, expressed by a prefix code based on the rule, 2^{2^x} (i.e., 1, 4, 16, 256); here, a wide range of gains can be controlled with a 16-bit word.

17.3.2 EXTENDED RESOLUTION ADCs

Stochastic techniques

Resolution can be improved beyond the actual resolution of the converter by means of multiple conversions. With stochastic conversion techniques, the resolution of an a/d converter (for example, a successive-approximation type) can be extended by several bits. The stochastic a/d technique (Figure 17.8) adds a pseudorandom dither with an average value of zero to the input signal, makes several conversions on the signal, and then computes the average. The time period of the conversion is set to be an integer number of power-line cycles to maximize the normal-mode noise rejection.

When pseudorandom noise is added to the signal input, the output of the converter will vary by some LSBs from a nominal value. As each conversion is completed, its output is summed with the previous ones and stored in an accumulator. When the desired number of conversions have been completed, the sum of all the data samples is averaged. In the example of Figure 17.8, an input of 0.3 is added to 8 samples of digitally generated dither and converted each time with 2-bit resolution. Without the dither, the digital value of the converter output is 0.25 (code 01); when dithered and averaged, using the eight terms and arbitrary precision, the result is 0.3125, corresponding to the code, 0101, representing the binary fraction, 0.0101 (instead of the lower-resolution 0.01).

As the simple 2-bit example of Figure 17.8 shows, the result will depend upon where the main analog input signal lies within a particular code width of the main a/d converter. For example, if the input lies precisely in the center of the code, the data samples will have equal positive and negative distributions about the nominal code value. If the analog input lies elsewhere between the limits of the quantum, the computed average will show either a positive or negative bias from the center of the code, and—since the sum is obtained digitally to arbitrary precision—the digital result is, in effect, a meaningful extension of resolution.

This approach requires time and software, but little or no additional hard-




Figure 17.8. Random dither improves effective converter resolution. Getting four-bit resolution from a 2-bit converter.

ware. The noise signal can be furnished by a low-resolution dither DAC, using noise numbers that are computed or retained in memory. The approach provides the optional tradeoff of either high-speed, medium-resolution conversion or slower, high-resolution conversion.

17.3.3 CHALLENGES IN DESIGNING A HIGH-RESOLUTION A/D CONVERTER

In a high-resolution, fast, successive-approximation a/d converter, all components—and in particular, the reference DAC and the comparator, must be accurate to considerably better than N bits (for example, 18 bits in an 18-bit converter) in order for the converter's overall accuracy to be better than N bits. We have already had some discussion of the difficulties of designing 18-bit DACs; since $\frac{1}{2}$ LSB of 18 bits is 1.9 ppm of full scale, a true 18-bit DAC would drift less than 1.9 ppm over the operating temperature range; for ordinary environments, this requirement implies resistor matching too close to contemplate. In addition, for an ADC, the comparator is an additional crucial component; it must respond quickly to an overdrive of $\frac{1}{2}$ LSB, or 19 microvolts out of 10 volts. It must also have input noise significantly smaller than this value.

No complete “off the shelf” comparator exhibits these specifications, so the designer must construct one from discrete components, consuming a lot of power to lower noise and increase speed, but also with very high gain to achieve the desired sensitivity. As the speed of the converter increases, however, the design starts to approach some quantum limitations.

For example, the charge involved in measuring one LSB of an 18-bit ADC in $10\mu\text{s}$ is only about 24,000 electrons. Determining the charge of a single electron, and the current flow of a single LSB demonstrates this. The charge on an electron is 1.6×10^{-19} coulombs, and thus $1 \text{ coulomb} = 6.2 \times 10^{18}$ electrons. Since, by definition, $1 \text{ ampere} = 1 \text{ coulomb/second}$, the current flowing from the LSB is 7.6 nA , or $= 2.4 \times 10^{10} \text{ e-/second}$. A $10\text{-}\mu\text{s}$ bit-conversion time, therefore, corresponds to about $\frac{1}{2}\mu\text{s/bit} \approx 24,000$ electrons.

The above discussion highlights a few of the problems associated with the design of a high resolution data converter. Although the design problems of the manufacturer are not the particular concerns of the user, many of the issues that face the designer also face the user. A key factor is the effect of noise on converter performance.

Once a converter has been selected, installed and tested, system designers often encounter the perplexing problem of the digital codes deviating from their expected values. This deviation is manifested as a number of unexpected codes over many conversions or over time with a steady analog input or output. These errors are caused by noise in the converter circuitry, in external circuitry or inherent in the analog signal itself. Noise produces code errors in the output code of an ADC or analog output errors in a DAC if the peak value of the noise, after filtering or integration exceeds one LSB. The percentage of time for which a given noise level will exceed the 1-bit threshold may be calculated.

17.3.4 A LITTLE NOISE THEORY

Noise in a/d converters arises from two principal sources, the “ideal” no-noise *quantization (or roundoff) error* that is inherent in the data conversion process, and *input noise*—which includes noise generated in the converter, noise arriving with the signal, and noise coupled in from the environment. If the analog signal is recovered in a d/a converter and compared with the analog input, there will be an error that depends on a combination of the error due to input noise and the quantization error.

In the same way that op-amp noise is referred to the op-amp input, all input noise in an ADC can be referred to the comparator input (where the decision is made as to which quantum an input belongs to), whether it is the comparator’s own noise, noise in the summing resistors at the input to the comparator, noise in the reference DAC output, noise that arrives with the signal, or noise that is coupled in from the power supply or the environment. All these noise sources can be summed to provide the total error signal the comparator

will see and that will determine the maximum realizable conversion error.

If the comparator were perfect, the arriving signal “clean,” and all interference eliminated, there would still be unavoidable noise due to thermal resistance noise (Johnson noise) of the signal-source input resistance and the DAC’s resistance ladder. Johnson noise voltage is given by

$$E_n = \sqrt{4kTRB} \quad (17.5)$$

Where E_n is the rms value of voltage generated in the effective source resistance, R (ohms); k is Boltzman’s constant (1.381×10^{-23} J/K); T is the temperature in kelvins ($^{\circ}\text{C} + 273.2^{\circ}$); and B is the effective bandwidth ($f_2 - f_1$), “brick wall,” in hertz. In more practical units,

$$E_n = 0.129 \sqrt{R \times B} \text{ microvolts rms} \quad (17.6)$$

where $R \times B$ is in $\text{M}\Omega\text{-Hz}$, $\text{k}\Omega\text{-kHz}$, or $\Omega\text{-MHz}$. For example, if $R = 1\text{ k}\Omega$ and $B = (5\text{ kHz} - 4\text{ kHz}) = 1\text{ kHz}$, then E_n is equal to $0.129\mu\text{V}$. If $R = 100\text{ ohms}$ and $B = 1\text{ kHz}$, then $E_n = 41\text{ nanovolts}$.

E_n is the theoretical minimum; in practical conversion systems, noise will be several times as large due to the other factors mentioned. Other additional noise sources will add to this minimum to further degrade the resolution of the converter. For noise sources that are independent of each other, a designer can calculate the overall noise from all known or estimated random noise sources as

$$E_n = \sqrt{(E_{n1})^2 + (E_{n2})^2 + \dots + (I_{n1} R_1)^2} \quad (17.7)$$

Where E_n is the total noise, E_{ni} are the various voltage noise sources, and I_{ni} are the various uncorrelated current-noise sources. In actual practice, the contributions of these noise sources can be as much as (or more than) 10 times the theoretical minimum Johnson noise.

The significance of the total noise in a high-resolution conversion system depends on the magnitude of the quantum step (1 LSB). If the a/d converter under consideration is assumed to be linear, the LSB will equal

$$\text{LSB} = \frac{\text{FSR}}{2^N} \quad (17.8a)$$

where FSR is the full-scale range of the converter and N is the resolution in bits. The equivalent input noise can be compared to this value (ratio = K) to determine its digital significance:

$$K = \frac{E_n}{\text{LSB}} = \frac{2^N E_n}{\text{FSR}} \quad (17.8b)$$

For example, in a 16-bit ADC, with 10-volt full-scale range, if the effective total input noise has an rms value of $25\mu\text{V}$, the value of K , in terms of the number of least-significant bits, is $65,536 \times 25 \times 10^{-6} / (10\text{ V}) = 0.16$.

If the noise is Gaussian, the rms value corresponds to the standard deviation, σ , of the distribution (Figure 17.9). Since the noise is presumed to be random, it contains all amplitudes, some very large, but the probability of large values decreases quite rapidly with amplitude. For example, the probability of a peak greater than 3σ is 0.27%, while the probability of a peak greater than 7σ is 2.6×10^{-12} .

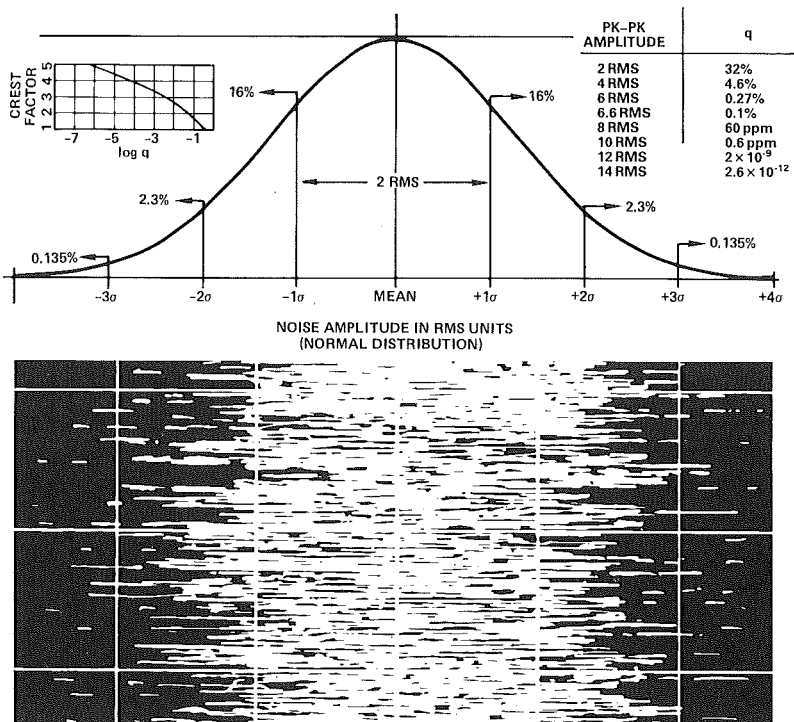


Figure 17.9. Relationship between rms value of Gaussian noise and probabilities of various peak amplitudes.

If we are concerned about the probability of missing codes (with error larger than $\frac{1}{2}$ LSB), we can relate that to the probability of a peak greater than $0.5/K$, since K corresponds to bit value of σ . For the above example, the probability of a peak larger than $\frac{1}{2}$ LSB will correspond to the probability of a peak greater than 3σ , or 0.27%.

Figure 17.9 provides information from which the probability of missing codes due to noise can be computed for any assumed value of input noise, converter resolution, and full-scale range, based on the above discussion and a Gaussian distribution.

It is important to note, however, that this discussion is addressed to *random* noise arising from natural phenomena. If the input noise is dominated by large spikes picked up from the power supply or the electrical environment, the fact that their rms value is minuscule will be of small comfort.

The size of the error, in LSBs, depends on the noise factor, K , and where the signal is located in the quantization interval. This value will be at a minimum when the actual code is centered in the interval and increase as the value deviates from the center value. For an expected random distribution of signals over time, the rms value of quantization noise can be treated as one of the inputs to equation 17.7.

Noise leads to random errors and missing codes in an ADC or nonmonotonic behavior in a DAC. Since these errors can lead to a closed-loop system making inappropriate responses to specific false data points, high-resolution systems, where random noise tends to limit resolution and to cause errors, can benefit by digital or analog filtering to smooth the data.

17.3.5 BOARD LAYOUT CONSIDERATIONS FOR HIGH-RES ADCs

Printed wiring board layout is extremely critical when using high-resolution analog-to-digital converters. High-speed logic level signals are present on the same board as low-level analog signals with microvolts of resolution. If signal conditioning or high-gain amplification is also included, the problems are compounded.

Figure 17.10 shows what can happen in an inadequate design. An amplifier with a gain of 1,000 is used to amplify a low-level 0-to-10 mV signal and present the resulting 0-to-10 V signal to the ADC. Suppose that, due to board strays, there are 1,000 megohms of resistance and 0.1 pF of capacitance between the summing node of A1 and one of the digital logic lines. The amount of dc pickup, relative to the analog input, is 10^{-6} (1,000 ohms/1,000 megohms), or 5 microvolts—0.05% of full scale. On the other hand, assuming feedback capacitance of 10 pF, a 5-volt logic edge would be attenuated by 0.1




Figure 17.10. Example of effects of stray capacitance and leakage resistance.

pF/10 pF, in the first stage, while the analog signal experienced a gain of 31.6, so 5 volts of logic would insert a 1.6 mV spike, referred to the input, or 16% of full-scale. However, it would be damped out within 2 μ s.

Effective solutions to this problem involve *distance*—keeping high-level and digital lines as far as possible from low-level analog lines, *isolation*—using shielding and guarding to isolate low-level signals, and *orientation*—where leads must cross, doing so at right angles, using twisted pairs to balance pick-up, etc.

If grounded guards are placed around the summing nodes of the amplifiers, stray capacitance from digital signal leads is to ground, rather than to the sensitive nodes. Not all guards are grounded guards; In order to be fully effective for low-frequency and dc common-mode pickup, as well as AC strays, guarding must be done at the same potential as the signal to be guarded.




Figure 17.11. A 16-bit sampling a/d converter.

Unfortunately, due to space limitations, optimum guarding and grounding practice in the neighborhood of high-resolution ADCs is sometimes difficult to achieve, and converter noise is the likely result. With the analog input grounded in the data-acquisition system of figure 17.11, noise could come from many sources. It is helpful to have a workable procedure for tracking down interference-noise problems. The internal architecture of a converter or data-acquisition system can be helpful in a rational search for noise sources.

17.4 HIGH-RESOLUTION SAMPLE/HOLDS

In high-resolution conversion, the dynamic characteristics, even of slowly varying signals must be considered. In order to faithfully digitize a signal, of frequency, f , and resolution, n , to 1 LSB, the conversion time (aperture) uncertainty, τ_a , must be less than:

$$\tau_a = \frac{2^{-n}}{\pi f} \quad (17.9)$$

If, for example, a 16-bit successive-approximation converter can complete a conversion within $35\mu\text{s}$, the highest frequency that can be converted with 16-bit resolution is 0.14 Hz . In order to convert at a sampling rate of 25 kHz, to handle, say, 10-kHz input signals, a sample-hold must be used ahead of the converter, with an aperture uncertainty better than 0.5 ns.

Sampling a 20-kHz signal to 16 bits requires the following specifications (actually, they should be considerably better when considering worst-case performance, but these are generally considered acceptable in the industry):

Aperture Jitter	0.25 ns
Slew Rate (20V pk-pk)	$1.26\text{V}/\mu\text{s}$
Feedthrough ($\frac{1}{2}$ LSB Max)	-102dB
Droop Rate ($\frac{1}{2}$ LSB Max in $15\mu\text{s}$)	$5.1\mu\text{V}/\mu\text{s}$
Acquisition time ($\pm \frac{1}{2}$ LSB max (with 20-kHz signal and $15/\mu\text{s}$ ADC)	$10\mu\text{s}$
Pedestal shift (max)	-96.3dB
Gain Tempco ($\pm 10^\circ\text{C}$ ambient)	$1.5\text{ppm}/^\circ\text{C}$
Thermal tail	0.3mV
Linearity error (max)	$\pm .0015\%\text{FSR}$

Aperture jitter is the uncertainty about when the sample is taken; it must be considered, even though the T-H control line is driven by a precise clock. All high-speed sampled-data systems depend on low aperture jitter for digitizing high-frequency signals for spectrum analysis and accurate signal reconstruction.

The T-H amplifier's slew rate determines the maximum switching rate when following changes between multiplexed input signals. The feedthrough from input to output while in the *hold* mode should be less than 1 LSB. The *hold*-

mode droop rate should be less than 1 LSB of droop in the output during the conversion time of the a/d converter. For a 16-bit ADC with a 15- μ s conversion time, for example, the maximum droop rate, as noted above, is $\frac{1}{2}$ LSB per 15 μ s; since $1 \text{ LSB} = 10/2^{16} \text{ V} = 152.6 \mu\text{V}$, the maximum droop rate is $5.1 \mu\text{V}/\mu\text{s}$.

The linearity error should be less than 1 LSB over the transfer function, as set by the relative accuracy of the a/d converter. The track-hold's acquisition time and settling time (t_{a+s}), along with the conversion time of the a/d converter (t_c), determine the highest sampling rate, f_s .

$$f_s = \frac{1}{t_{a+s} + t_c} \quad (17.10)$$

This, in turn, will determine the highest input signal frequency that can be sampled at a minimum of twice per cycle, according to Nyquist sampling theory.

The pedestal shift due to input signal changes should either be linear, to be seen as gain error, or negligible. Feedthrough should also be negligible. The temperature coefficients for drift should be low enough so that the full accuracy is maintained over some minimum temperature range. The droop rate and pedestal will shift more over temperature for temperature ranges above $+70^\circ\text{C}$, generally a considerably higher temperature range than these devices will experience for most of the applications in which they are used. An additional factor to consider with high-resolution converters is the noise in the T-H during the *track* mode, since it will affect the value that is sampled when the T-H is switched into *hold*. This noise must be added (root sum-of-squares) to converter noise when calculating the actual noise error in an ADC.

Minimal thermal-tail effects are another requirement for high-resolution applications. The self-heating-induced transient errors due to transient thermal imbalances resulting from the changing current levels in the output stages of T-H amps may cause more than 1 LSB of error because of the time required for the converter to settle to equilibrium temperature.

The profusion of T-H specifications and the need to carefully monitor the design might be better solved by turning to a high-resolution data acquisition subsystem or sampling a/d converter, which has been engineered to perform an overall task. The sampling a/d converter contains both a track-and-hold and an a/d converter. An example of a high-resolution sampling converter's structure can be seen in Figure 17.11.