Subranging ADCs Operate at High Speed with High Resolution
ADC-AN-5

Subranging A/D converters offer performance levels difficult to obtain with successive-approximation or flash converters. They can deliver higher conversion speed and resolution and suit such applications as digital signal processing. Part 1 of this 3-part series explores the architecture and operation of these devices. Part 2 will cover subranging-ADC parameters and specifications. Part 3 will conclude the series with test and measurement principles.

The subranging, or multipass, A/D converter has become increasingly popular in the last few years. A major reason for this popularity is digital signal processing, which demands high conversion speeds and resolution. The traditional successive-approximation converter has reached its speed-resolution limit (about 1 µsec for 12 bits) and can't meet the demands of many applications. Although flash converters offer high speeds, a practical limit exists to the resolution they can provide because the number of comparators rises exponentially with the number of bits. A 12-bit flash converter, for example, does not exist.

A flash converter, however, is an essential part of the subranging-ADC architecture. Designers have a wider choice of flash converters than they did a few years ago, and today's devices have significantly better 1 performance and lower prices (approximately $10 for an 8-bit flash converter in OEM quantities). Semicustom design has also helped spur subranging-ADC manufacturing because the devices inherently require more components than successive-approximation converters. Integrating the timing and correction logic on a single chip greatly reduces cost, the number of active components, and assembly and reliability problems.

The three types of subranging-ADC architectures are conventional, pipelined, and intermeshed; each type best suits certain applications. All subranging ADCs—whether the device is a hybrid IC or ICs and discrete components on a pc board—contain at least a sample-and-hold (S/H) circuit, a D/A converter, a scaling network, and timing and digital-correction logic.

The conventional subranging architecture (Fig 1) is a 2-stage A/D converter. With S_1 closed and S_2 open, the S/H circuit switches to the hold mode. The flash converter then quantizes the input signal, V_{IN}. After proper scaling, the D/A converter converts the digitized and latched signal back into an equivalent voltage. This voltage is subtracted from the original input signal at the summing junction, yielding the difference between the first conversion and the input signal. The closing of S_2 and the opening of S_1 feeds the difference signal back to the flash converter, which amplifies and digitizes the signal. After latching, the result of this conversion goes through the digital-correction logic to produce the output.

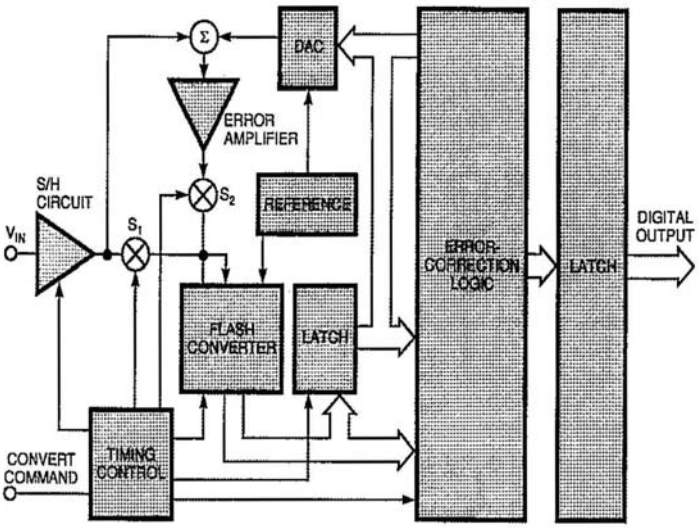

Fig 1. A conventional subranging ADC uses an S/H circuit, a flash converter, and a D/A converter.

The three different types of subranging-ADC architectures are conventional, pipelined, and intermeshed.

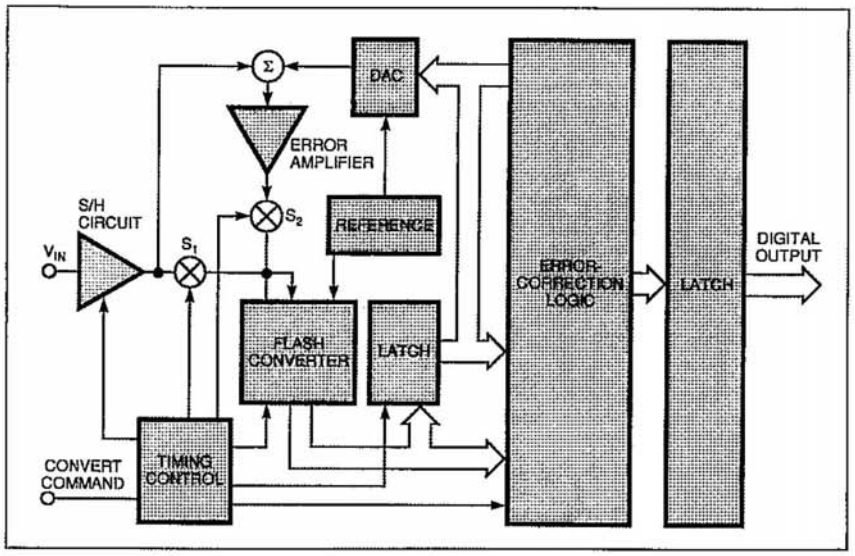

***Fig 1—A conventional subranging ADC** uses an S/H circuit, a flash converter, and a D/A converter.*

Designing and fabricating a good high-resolution flash converter is a major task. Therefore, in monolithic subranging ADCs, using a lower-resolution (3- to 4-bit) flash converter is preferable. Of course, the lower resolution means more passes through the flash converter, making the subranging ADC seem like a successive-approximation converter with a reduced number of trials. This variation of the conventional architecture, which combines flash and successive-approximation features, is the recursive subranging architecture.

The second type of subranging converter has a pipelined architecture. Compared with the other subranging-ADC types, the pipelined architecture offers a faster throughput rate because the circuit can initiate a new conversion before the previous conversion is finished. However, the conversion time is not significantly improved, and the digital output data corresponding to the present conversion is always delayed by at least one clock.

The 2-stage pipelined converter has an extra S/H circuit and an extra flash converter (**Fig 2**). In operation, the first S/H circuit switches to the hold mode after acquiring the input signal, V_{IN}. The first flash converter then quantizes the input signal, while the

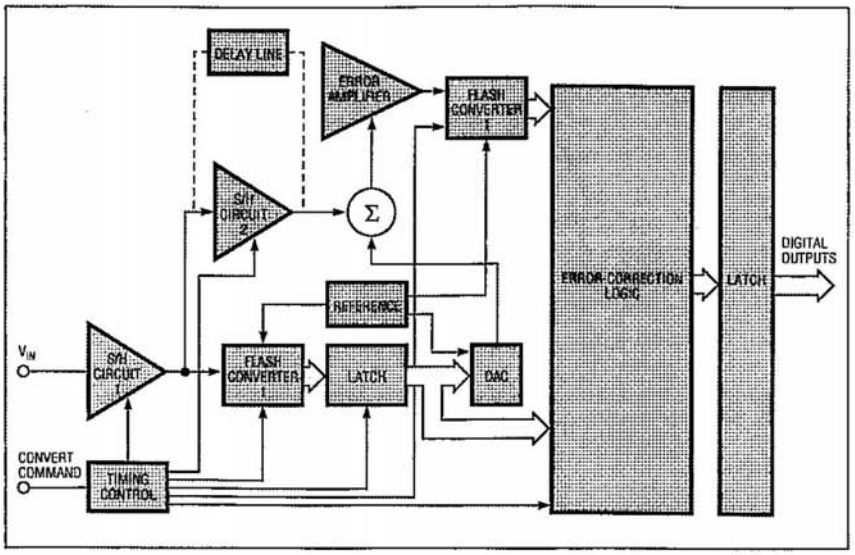

***Fig 2—A pipelined subranging ADC** uses two S/H circuits, two flash converters, and a D/A converter.*

second S/H circuit goes to the hold mode. The D/A converter latches and converts the digitized signal into an equivalent voltage. This voltage is then subtracted from the held output voltage of the second S/H circuit, which represents the input voltage at the time the first flash converter made its conversion. The second flash converter amplifies and digitizes the difference between the first conversion and the input voltage. After latching, the result of this conversion goes through the digital correction logic to produce the output.

Immediately after the second S/H circuit switches to the hold mode, the input S/H circuit can acquire a new signal, effectively increasing the throughput rate (the rate at which the converter can accept new convert commands). Another advantage of pipelining is that you can time the outputs of the S/H circuit and the A/D converter to change simultaneously. Thus, the outputs don't overdrive the error amplifier, resulting in only a very short transient switching glitch.

In place of the S/H circuit, you could also use a delay line. For the optimum throughput rate, the delay should not exceed the conversion time of the first flash converter plus the settling time of the D/A converter. The delay line itself must be of high fidelity and have a large bandwidth.

The third type of subranging converter uses an intermeshed architecture. **Fig 3** shows a block diagram of this type of converter, which operates as follows: After the S/H circuit acquires the signal, the MSB flash converter decides the range of the input. This input lies between two resistors in the ladder and determines the most significant bits. These two points on the resistor ladder then switch to the reference top and reference bottom of the second flash converter (the higher

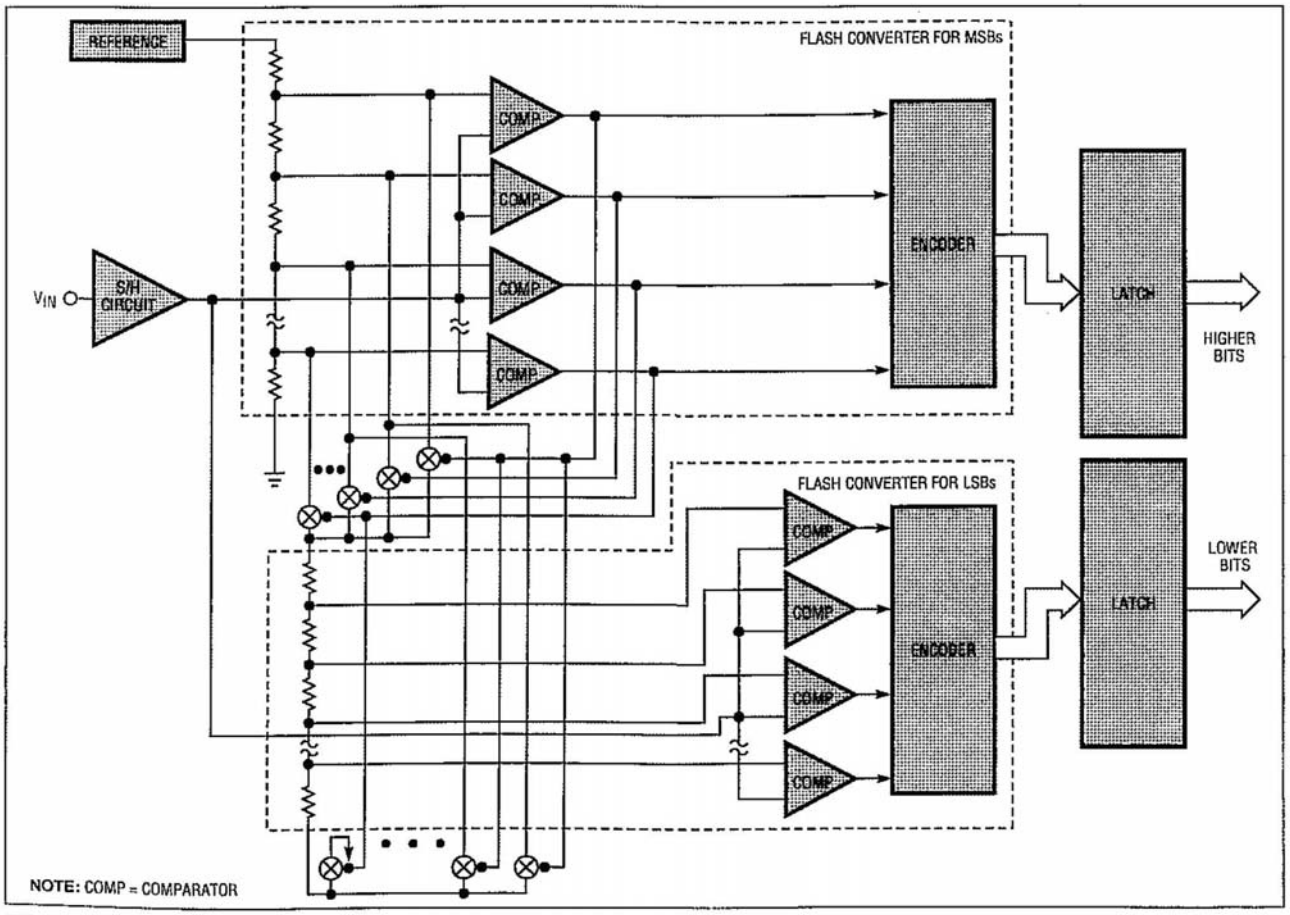

***Fig 3—An intermeshed subranging ADC** uses separate flash converters for the MSBs and LSBs. Note the absence of the D/A converter, the error amplifier, and the correction logic.*

Perhaps the most critical design decision is choosing the flash converter.

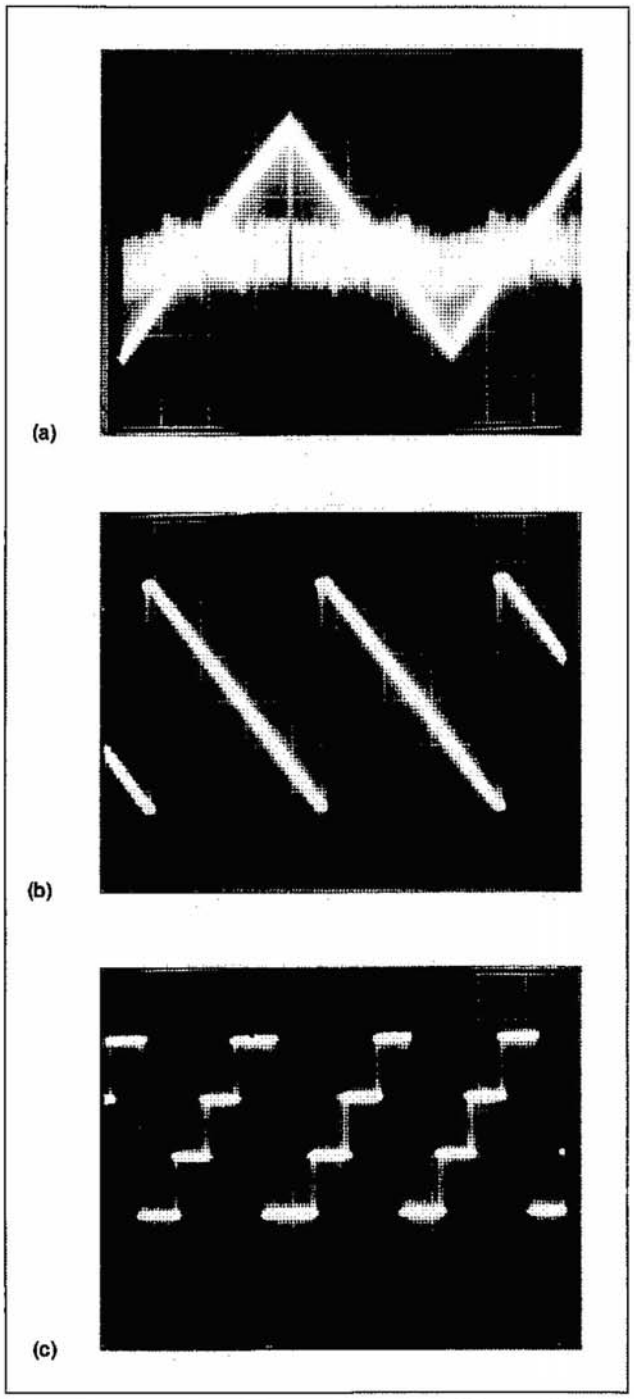

Fig 4—These waveforms illustrate first-pass problems at different test points. *Photo* ***a*** *shows the output of the error amplifier with a triangle-wave input. Photos* ***b*** *and* ***c*** *show the reconstructed output of the ADC using a 5-bit DAC (***b***) and a 2-bit DAC (***c***).*

voltage on the ladder will be the reference top and the lower voltage will be the reference bottom). Upon a convert command, the LSB flash converter digitizes the original input to produce the least significant bits.

Because this architecture has no correction logic, the MSB flash converter must be as linear as the intended linearity of the overall A/D converter. Note the absence of the D/A converter, the error amplifier, and the correction logic in the block diagram. This relative simplicity, plus the fact that the circuitry is repetitive, suits the intermeshed architecture for monolithic applications.

To the uninitiated, the block diagram of a subranging A/D converter can appear to be deceptively simple, consisting merely of a few building blocks. On the contrary, subranging A/D converters are the most challenging ADCs to design and manufacture. Because many sources of error exist in a subranging ADC, engineers should be aware of each one and pay attention to the smallest details. Establishing an error budget is the only practical way designers can achieve a design goal systematically.

The vast majority of errors occur in the first conversion because that conversion is only as accurate as the first-pass flash converter (5, 6, 7, or 8 bits). The best test point to detect first-pass problems is at the output of the error amplifier, where you can look at the difference output **(Fig 4a)**. The best analog input to the A/D converter for this observation is a triangular wave. You can also observe the effects of a particular error source at the reconstructed output of the ADC. A D/A converter provides this reconstruction. **Fig 4b** shows the reconstructed output using a 5-bit DAC; **Fig 4c** is the result using a 2-bit DAC.

Perhaps the most critical design decision is choosing the flash converter. Designers must be careful to match the flash converter with the application. The linearity of the flash converter dictates the number of correction, or overlap, bits. For example, a 12-bit subranging ADC that had a 7-bit flash converter with 12-bit linearity would require no error correction, assuming no other sources of error. However, if you use a typical flash converter that is accurate, or linear, to six bits, you'll need at least one bit of correction. Because there are always other sources of error, such as offset and gain drift, having two bits of correction is advisable.

A lower number of bits in the first-pass conversion translates to a lower amplification factor in the second pass. As a result, the amplifier settles faster, and you can get by with a lower-resolution D/A converter,

which is easier to design. However, this lower-resolution DAC mandates a higher-resolution flash converter in the second pass, which is more expensive than a lower-resolution converter.

CMOS flash converters are attractive because of their low power consumption. Converters that operate from a single supply have especially low power consumption. However, because most CMOS flash converters use a design scheme in which the comparators switch back and forth between the reference ladder and the input depending on the clock level (**Fig 5**), large glitches that are synchronized to the convert command appear at the converters' inputs. As a result, CMOS converters distort the input signal, making the devices hard to drive. To overcome this problem, use a high-speed, wide-bandwidth buffer with low output impedance.

CMOS flash converters cause large spikes on the power lines, which not only degrade the performance of the converter, but also create problems for the overall system. Heavily bypassing the reference voltages and power lines right at the flash converter helps lessen, or prevent, spikes. Also, avoid external HCMOS logic when the ADC's sampling rate is greater than 5 MHz. Because of its high-speed switching, HCMOS logic also creates large spikes on the power lines.

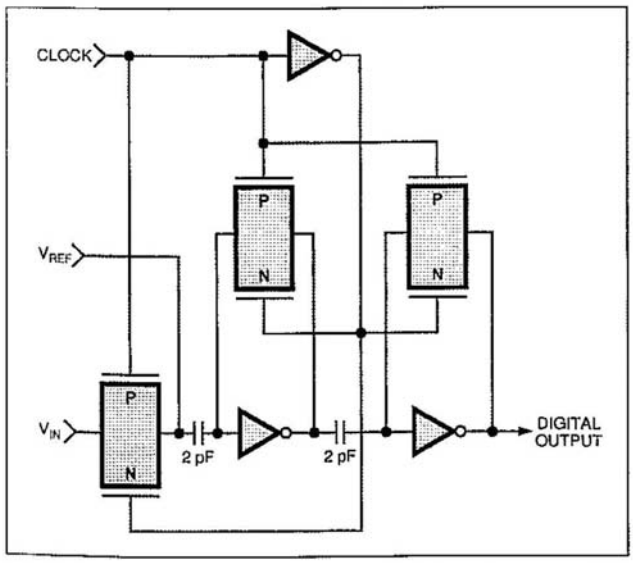

*Fig 5—**Most CMOS flash converters use a design scheme** in which the comparators switch back and forth between the reference ladder and the input. This switching can cause undesirable glitches synchronized to the convert command.*

Compared with CMOS types, bipolar flash converters, have fewer problems. However, most bipolar flash ADCs require dual supplies (often +5 and −5.2V) and usually consume more power than CMOS converters.

In a subranging A/D converter, the bit resolution of the D/A converter does not need to be more than the resolution of the flash converter used in the first-pass conversion. However, the DAC's differential and integral linearity must be considerably better than the desired differential and integral linearity of the A/D converter. The integral nonlinearity of the DAC not only causes integral nonlinearities but creates nonmonotonicity and differential nonlinearities every time the input of the DAC changes. This effect is called overlap; **Fig 6** illustrates overlap at the ADC's reconstructed output. The DAC should have an accuracy

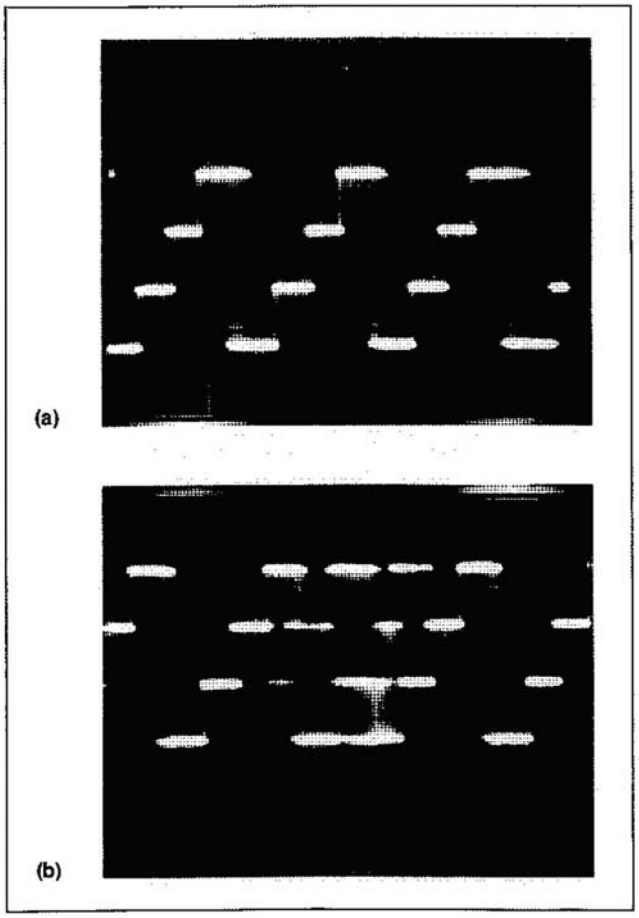

*Fig 6—**These scope photos compare** the reconstructed output of the D/A converter in the typical case **(a)** and when the DAC has integral error **(b)**. This error is called overlap.*

The resolution of the DAC in a subranging A/D converter does not need to be greater than that of the first-pass flash converter.

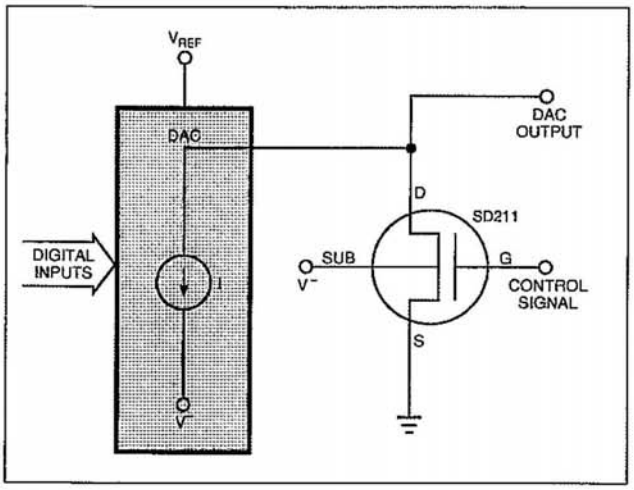

***Fig 7—By using a MOSFET to switch** the output of the DAC to ground while it's settling, you can minimize the load impedance at the output of the DAC and optimize its settling time.*

at least 1-bit greater than the desired accuracy for the total converter.

For subranging A/D converters, current-output DACs are a good choice because they settle faster than voltage-output DACs. For optimum settling time, however, you should minimize the load impedance at the output of the DAC. The best way to minimize this impedance is to switch the output of the DAC to ground with a MOSFET transistor while the DAC is settling (**Fig 7**). Most current-output DACs have application resistors you can use to set the best gain and gain-drift performance. If the DAC does not have such resistors you must design them into the same resistor network that sets the reference current of the DAC.

The error amplifier

The error amplifier scales the difference between the first-conversion output and the input signal. The characteristics of the first-pass flash converter determine the gain of this amplifier. This amplifier does not have to settle to better than the accuracy of the second flash converter; therefore, most commercially available monolithic high-speed op amps will do. However, because the amplifier's closed-loop gain requirements are normally high, the device is subject to saturation while the error signal is settling. To eliminate this problem, make sure the error signal settles before it switches to the amplifier, as **Fig 8** shows.

Another problem to watch for is overvoltage-recovery time. If the input of the A/D converter exceeds the analog input range even slightly, an undesirable overvoltage recovery time will occur. Usually the input capacitance of the flash converter increases the amplifier's overshoot and, as a result, increases the settling time. In this case, a small-value resistor in series with the output of the amplifier can help.

By far the most untamed errors for any system, particularly for subranging A/D converters, are those

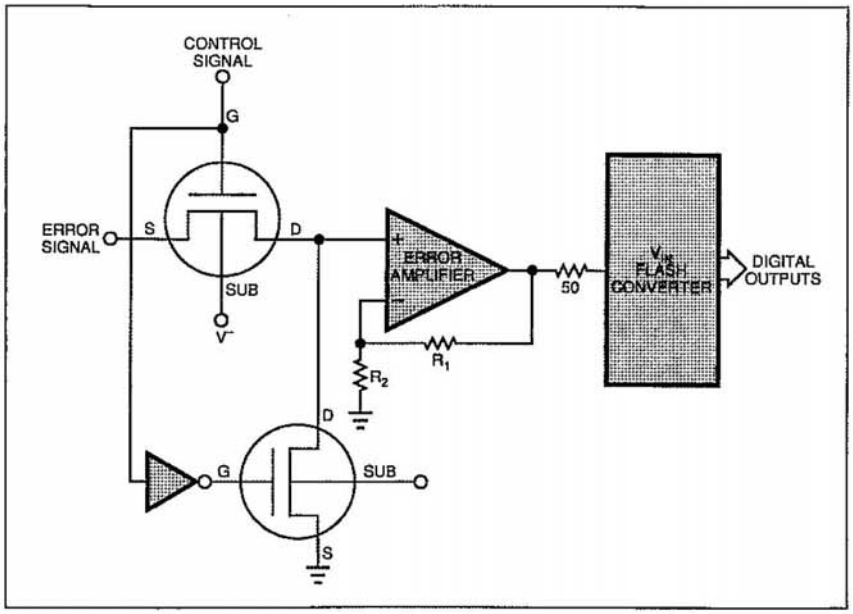

***Fig 8—Switching the input of the error amplifier** to ground while the error is settling can prevent the amplifier from saturating.*

The characteristics of the flash converter determine the gain of the error amplifier. Most monolithic high-speed op amps are suitable.

that improper grounding introduces. Designers often think that their breadboard is functioning properly, but when they change the setup or the timing (repetition rate, duty cycle, or the fall and rise times of the convert command), major errors occur because of ground loops or poor grounding. In some cases, the breadboard stops functioning. **Fig 9** shows the transients on the ground of a subranging ADC with poor grounding. You can see the effects of this poor grounding by looking at the error-amplifier output and the ADC's reconstructed output.

Because the number of overlap bits limits the digital correction, poor grounding could cause an error in the first pass that exceeds the digital-correction limit. This first-pass error can cause the remainder to go out of the correction window **(Fig 10)**.

You can solve most grounding problems by separating the analog and digital grounds and connecting them only to heavy ground planes underneath or close to the flash converter. Making ground runs as wide as possible and decoupling the power supplies will also help. Note that CMOS logic and CMOS flash converters tend to magnify any grounding problems because of the large transients they cause whenever they switch. In such cases, you need to take additional care.

A major factor in the proper functioning of a subranging A/D converter is the error-correction logic, which is the Boolean algebra performed on the outputs of the first- and second-pass conversions to produce the ADC's output. This logic corrects for any first-pass

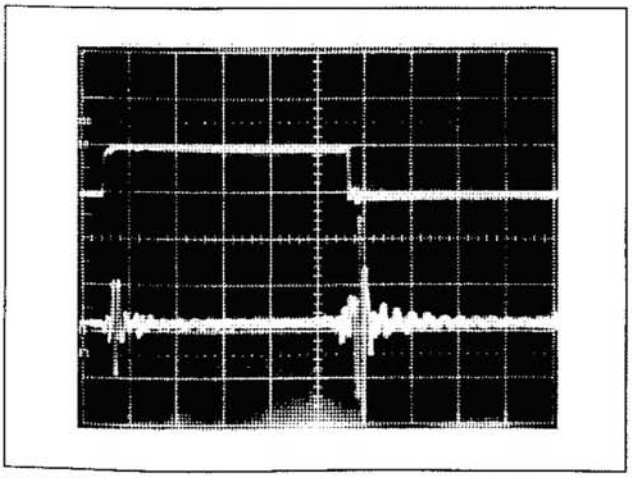

***Fig 9—This photo shows the transients** on the ground of a subranging ADC that has poor grounding. The upper trace shows the signal; the bottom trace shows the analog ground with respect to the same ground at a different physical location.*

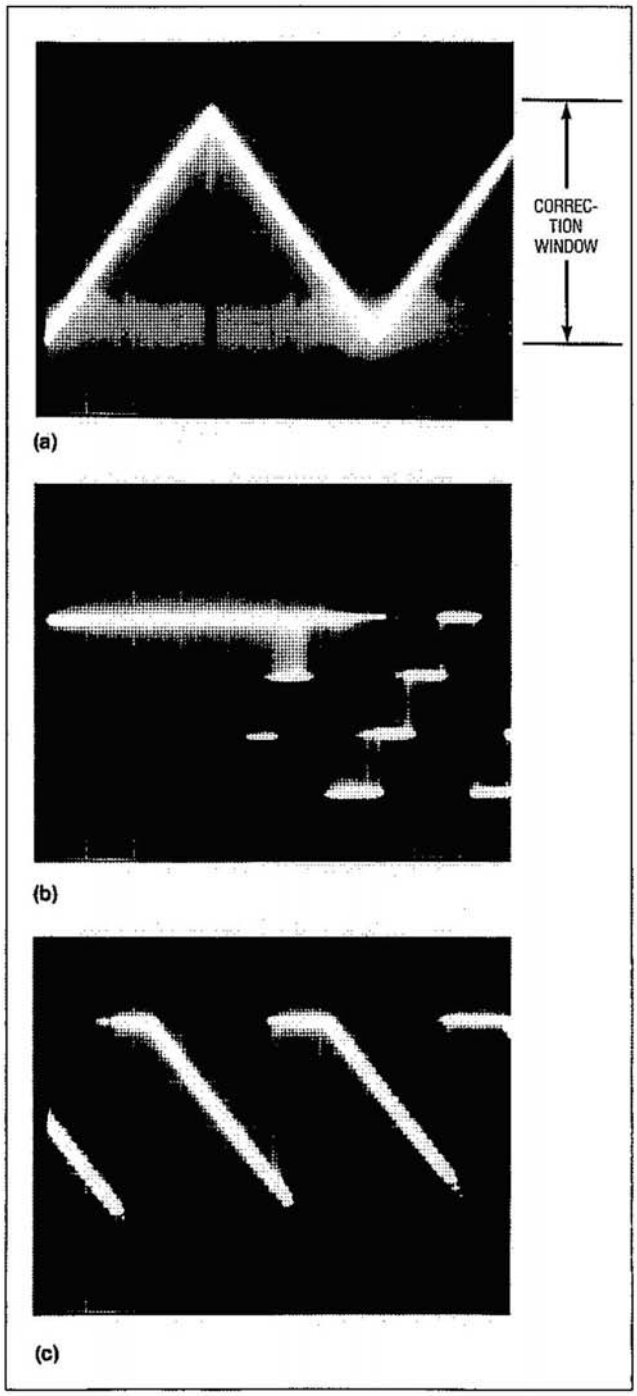

***Fig 10—In this ADC, a slight change in the start-convert pulse width** caused the remainder to go out of the correction window (a). Traces b and c show the resultant missing codes at the reconstructed output for a 2-bit DAC and a 5-bit DAC, respectively.*

You can solve most grounding problems by separating the analog and digital grounds and connecting them to heavy ground planes underneath or close to the flash converter.

errors that create remainders that stay within the correction window (**Fig 11**).

Generally, the device's biasing offsets the input negatively in the first-pass conversion to ensure that the error, or remainder, is positive. A positive remainder simplifies the correction logic to an addition, which means the outputs of the first and second conversion are added to get the final output of the ADC. The way these outputs are added depends on the number of correction bits, or overlaps. For example, in the case of a 7-bit flash converter with 1 bit of correction, the ADC's output is

First-pass data:	D1 D2 D3 D4 D5 D6 [D7]	→ overlap bit
Second-pass data:	+ [D1] D2 D3 D4 D5 D6 D7	
The ADC output:	B1 B2 B3 B4 B5 B6 B7 B8 B9 B10 B11 B12 B13.	

D7 in the first pass and D1 in the second pass must have the same weights to be added. If you use a 7-bit flash converter with 2 bits of correction, the ADC's output is

First-pass data:	D1 D2 D3 D4 D5 [D6 D7]	→ overlap bits
Second-pass data:	+ [D1 D2] D3 D4 D5 D6 D7	
The ADC output:	B1 B2 B3 B4 B5 B6 B7 B8 B9 B10 B11 B12.	

In this case, D6 and D7 in the first pass must have the same respective weight as D1 and D2 in the second pass. To make sure the overlap bits have the same weight, the gain of the error amplifier must be $K = G \times 2^{N-M}$, where G is the first-pass gain, N is the resolution of the first-pass flash converter, and M is the number of correction bits. In the first example, $K = G \times 2^{7-1} =$

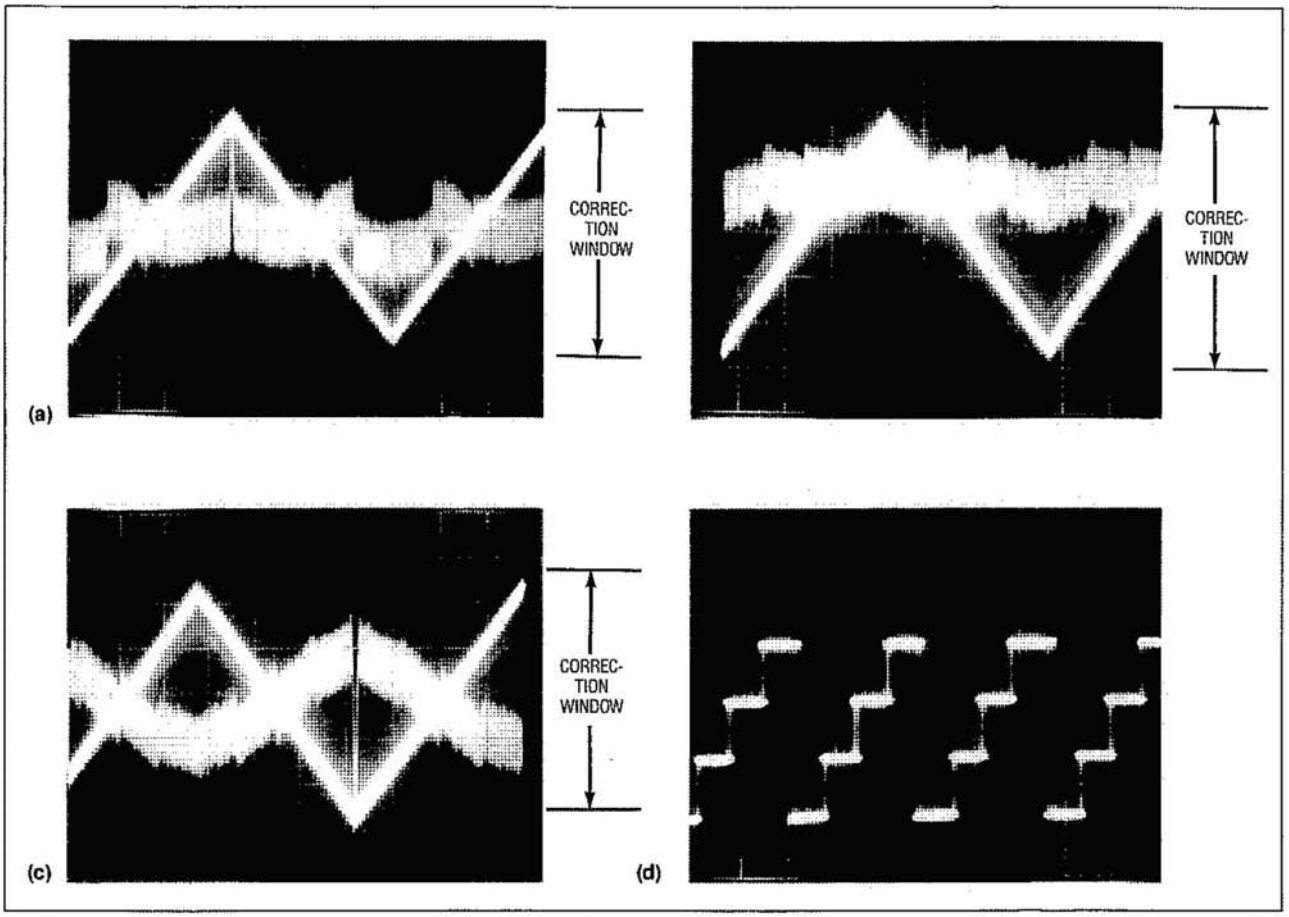

***Fig 11—Correction logic in a subranging ADC** corrects for first-pass errors that stay within the correction window. Photos **a**, **b**, and **c** show the effects of first-pass nonlinearity, offset, and gain errors, respectively. Photo **d** shows that these errors have no effect on the reconstructed output.*

Subranging ADCs use correction logic to guard against first-pass errors.

G×64. In the second example, $K = G \times 2^{7-2} = G \times 32$.

Note that the LSBs of the flash converter in the second pass are the LSBs of the output. As a result, the differential nonlinearity of the output is the flash converter's differential nonlinearity. The only time this statement is not true is if you do not adjust K to the right value. In this case, whenever the LSB of the first pass changes from 1 to 0 or 0 to 1, discontinuities, or noisy codes, will appear in the output of the A/D converter.

One common problem with the addition logic is that when the full-scale output carries, the output of the A/D converter rolls over, thereby generating an erroneous code. For example,

To avoid this problem, use a simple OR gate that ORs the carry with each output bit. This action forces the output to stay at all 1s for any input exceeding full scale. **Fig 12** shows an error-correcting circuit for a 12-bit, 2-pass converter with two bits of correction.

If the first-pass data is	1 1 1 1 1 1 1
and the second-pass data is	0 1 0 0 0 0 0,
then the output will be	0 0 0 0 0 0 0 0 0 0 0 0 0 with a carry.
The correct output is	1 1 1 1 1 1 1 1 1 1 1 1 1.

The S/H circuit

Another key element in the performance of any high-speed A/D converter is the S/H circuit. Because of the usually low input impedance in a subranging A/D converter, the designer must pay close attention to designing or selecting the S/H circuit. This essential element consists of an input buffer, a switch, a hold capacitor, and an output buffer that drives the A/D converter. When an S/H circuit is used in front of a subranging ADC, the circuit's dynamic output impedance and aperture uncertainty are of prime importance.

The dynamic output impedance of the S/H circuit, which is a function of the bandwidth of the output buffer, determines how fast the device responds to the

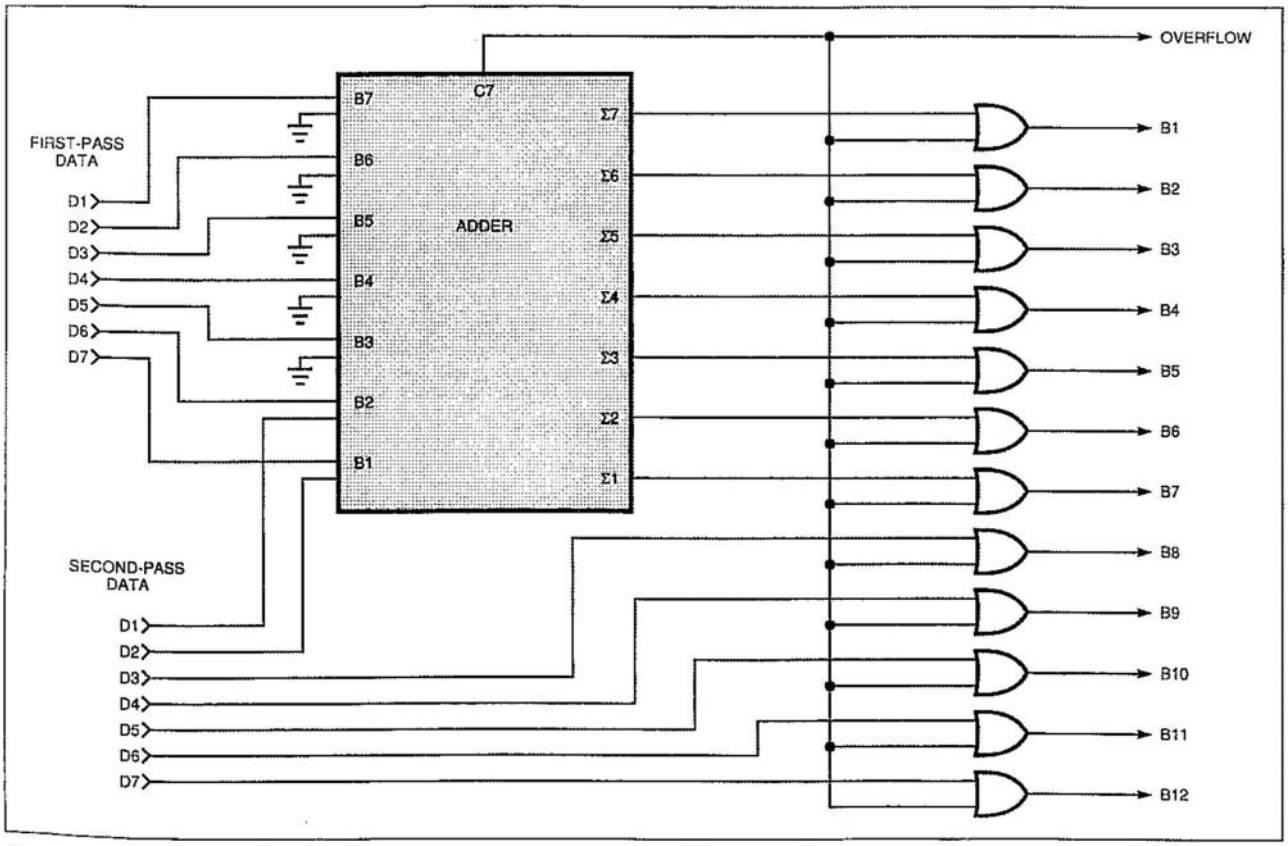

***Fig 12—This diagram shows the error-correction logic** for a 12-bit, 2-pass converter with two bits of correction.*

Subranging ADCs are the most challenging ADCs to design and manufacture.

dynamic load it drives. Because of subranging ADCs' low input impedance and the switching that takes place on the input, the devices require a S/H circuit with a low dynamic output impedance to attain an optimal conversion time.

For precision A/D converters, try to avoid open-loop S/H circuits. The gain of open-loop S/H circuits changes drastically with input frequency and amplitude, and most noticeably with the output load. Any of these factors can cause large static and dynamic errors.

Another error source is the aperture uncertainty, which is the uncertainty period associated with the closing of a switch. Many factors contribute to this error source. The most common are timing jitter caused by random noise, 60-Hz line frequency or other frequencies modulating the power lines, and the uncertainty of the sample command. The ADC's output takes on the spectral characteristics of the error source. The aperture uncertainty also limits the input frequency that the circuit can sample within the specified error budget. The relationship between aperture uncertainty (T_A) and the input frequency is as follows, where A is the amplitude of the input signal in volts:

$$\text{If } V_{IN}=A\cdot\sin(\omega t)=A\cdot\sin(2\pi f_t)$$

then the maximum rate of change for V_{IN} is

$$\frac{dV_{IN}}{dt}=\frac{dA\,\sin(2\pi f_t)}{dt}\Bigg|_{t=0}=2A\pi f\cos(2\pi f_t)\Bigg|_{t=0}$$

$$\left(\frac{dV_{IN}}{dt}\right)_{MAX}=2A\pi f$$

Aperture-uncertainty noise generally follows a Gaussian distribution similar to white noise, which means that the rms aperture uncertainty corresponds to the distribution's σ value. The distribution's 2σ point thus becomes the proper choice for the maximum value. The maximum aperture uncertainty equals $2T_A$. To determine the maximum full-scale sine-wave frequency (f_{MAX}) that produces ½-LSB error, first calculate the error arising from $2T_A$:

$$2T_A\text{ error}=2(T_A)_{MAX}(dV_{IN}/dt)=2T_AA2\pi f_{MAX}$$

Thus,

$$f_{MAX}=1/T_A\pi 2^{n+2}$$

The error resulting from the aperture uncertainty is primarily random, which makes the noise additive. The general expression for the noise produced by Gaussian time jitter is

$$\text{S/N ratio}=-20\log(2\pi fT_A)$$