Introduction

Integrating A/D converters have two characteristics in common. First, as the name implies, their output represents the integral or average of an input voltage over a fixed period of time. Compared with techniques which require that the input is "frozen" with a sample-and-hold, the integrating converter will give repeatable results in the presence of high frequency noise (relative to the measurement period). A second and equally important characteristic is that they use time to quantize the answer, resulting in extremely small nonlinearity errors and no possibility of missing output codes. Furthermore, the integrating converter has very good rejection of frequencies whose periods are an integral multiple of the measurement period. This feature can be used to advantage in reducing line frequency noise, for example in laboratory instruments (Figure 1).
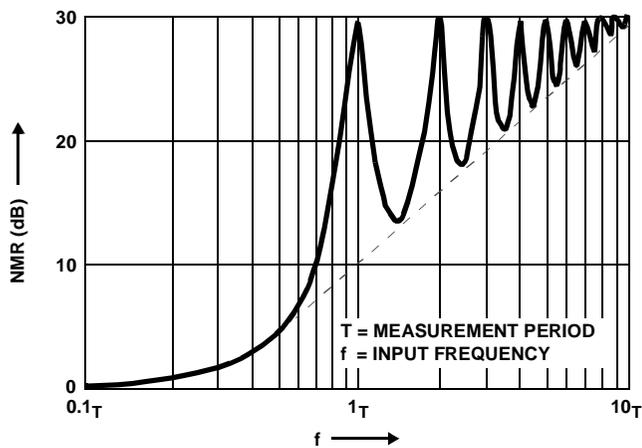
FIGURE 1. NORMAL MODE REJECTION OF DUAL-SLOPE CONVERTER AS A FUNCTION OF FREQUENCY

In addition, a competitive instrument-quality product should have the following features:

 1. Single Reference Voltage. This is strictly a convenience to the user; but since many designs are available with single references that contribute negligible error, products requiring dual references are rapidly becoming obsolete.

 2. Auto-Zero. This eliminates one trim-pot and a troublesome calibration step. Furthermore, it allows the manufacturer to use op amps with up to 10mV offset while still achieving system offsets of only a few microvolts.

 3. High Input Impedance. Recently developed monolithic FET technology allows input impedances of 1000MΩ and leakages of a few pico amps to be achieved fairly readily.

The unique characteristics of the integrating converter have made it the natural choice for panel meters and digital voltmeter applications. For this reason, overall usage of integrating converters exceeds the combined total of all other conversion methods. Furthermore, the availability of low cost one chip converters will encourage digitizing at the sensor in applications such as process control. This represents a radical departure from traditional data logging techniques which in the past have relied heavily on the transmission of analog signals. The availability of one chip microprocessor system (with ROM and RAM on chip) will give a further boost to the 'conversion at the sensor' concept by facilitating local data processing. The advantage of local processing is that only essential data, such as significant changes or danger signals, will be transmitted to the central processor.

The Dual Slope Technique - Theory and Practice

The most popular integrating converter is the "dual-slope" type, the basic operating principles of which will be described briefly. However, most of the comments relating to linearity, noise rejection, auto-zero capability, etc., apply to the whole family of integrating designs including charge balancing, triple ramps, and the 101 other techniques that have appeared in the literature. A simplified dual-slope converter is shown in Figure 2.

The conversion takes place in three distinct phases (Figure 3).
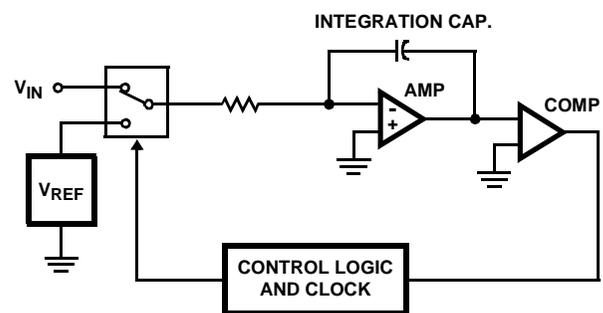
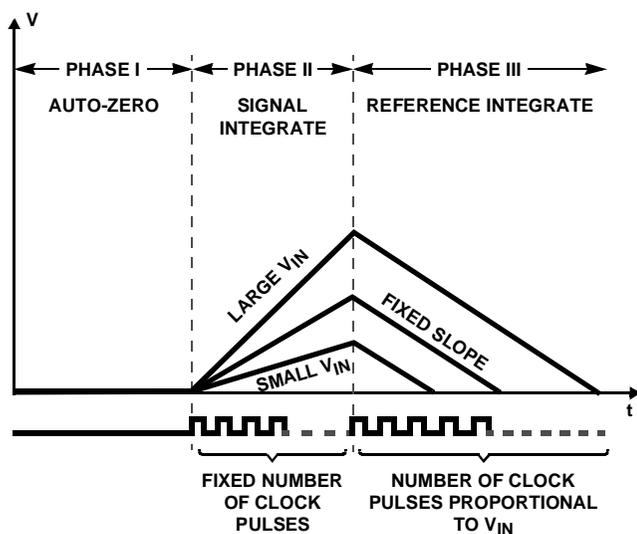**FIGURE 2. SIMPLIFIED DUAL-SLOPE CONVERTER**

FIGURE 3. THE THREE PHASES OF DUAL-SLOPE CONVERSION

Phase 1, Auto-Zero. During auto-zero, the errors in the analog components (buffer offset voltages, etc.) will be automatically nulled out by grounding the input and closing a feedback loop such that error information is stored on an "auto-zero" capacitor.

Phase 2, Signal Integrate. The input signal is integrated for a fixed number of clock pulses. For a $3^1/_2$-digit converter, 1,000 pulses is the usual count; for a $4^1/_2$-digit converter, 10,000 is typical. On completion of the integration period, the voltage V in Figure 3 is directly proportional to the input signal.

Phase 3, Reference Integrate. At the beginning of this phase, the integrator input is switched from V_{IN} to V_{REF}. The polarity of the reference is determined during Phase 2 such that the integrator discharges back towards zero. The number of clock pulses counted between the beginning of this cycle and the time when the integrator output passes through zero is a digital measure of the magnitude of V_{IN}.

The beauty of the dual slope technique is that the theoretical accuracy depends only on the absolute value of the reference and the equality of the individual clock pulses within a given conversion cycle. The latter can easily be held to 1 part in 106, so in practical terms the only critical component is the reference. Changes in the value of other components such as the integration capacitor or the comparator input offset voltage have no effect, provided they don't change during an individual conversion cycle. This is in contrast to Successive Approximation converters which rely on matching a whole string of resistor values for quantization.

In a very real sense the designer is presented with a near perfect system; his job is to avoid introducing additional error

sources in turning this textbook circuit into a real piece of hardware.

From the foregoing discussion, it might be assumed that designing a high performance dual-slope converter is as easy as falling off the proverbial log. This is not true, however, because in a practical circuit a host of pitfalls must be avoided. These include the non-ideal characteristics of FET switches and capacitors, and the switching delay in the zero crossing detector.

Analyzing the Errors

At this point it is instructive to perform a detailed error analysis of a representative dual slope circuit, Intersil's ICL7135. This is a $4^1/_2$-digit design, as shown in Figure 4. The error analysis which follows relates to this chip - however, the principles behind the analysis apply to most integrating converters.

The analog section of the converter is shown in Figure 5. Typical values are shown for 120kHz clock and 3 measurements/second. Each measurement is divided into three parts. In part 1, the auto-zero FET switches 1, 2 and 3 are closed for 10,000 clock pulses. The reference capacitor is charged to V_{REF} and the auto-zero capacitor is charged to the voltage that makes dV/dt of the integrator equal to zero. In each instance the capacitors are charged for 20 or more time constants such that the voltage across them is only limited by noise.

In the second phase, signal integrate, switches 1, 2 and 3 are opened and switch 4 is closed for 10,000 clock pulses. The integrator capacitor will ramp up at a rate that is proportional to V_{IN}. In the final phase, de-integrate, switch 4 is opened and, depending on the polarity of the input signal, switch 5 or 6 is closed. In either case the integrator will ramp down at a rate that is proportional to V_{REF}. The amount of time, or number of clock pulses, required to bring the integrator back to its auto-zero value is 10,000 $\left(\dfrac{V_{IN}}{V_{REF}}\right)$.

Of course, this is a description of the "ideal" cycle. Errors from this ideal cycle are caused by:

1. Capacitor droop due to leakage.
2. Capacitor voltage change due to charge "suck-out" (the reverse of charge injection) when the switches turn off.
3. Nonlinearity of buffer and integrator.
4. High-frequency limitations of buffer, integrator and comparator.
5. Integrating capacitor nonlinearity (dielectric absorption).
6. Charge lost by C_{REF} in charging C_{STRAY}.

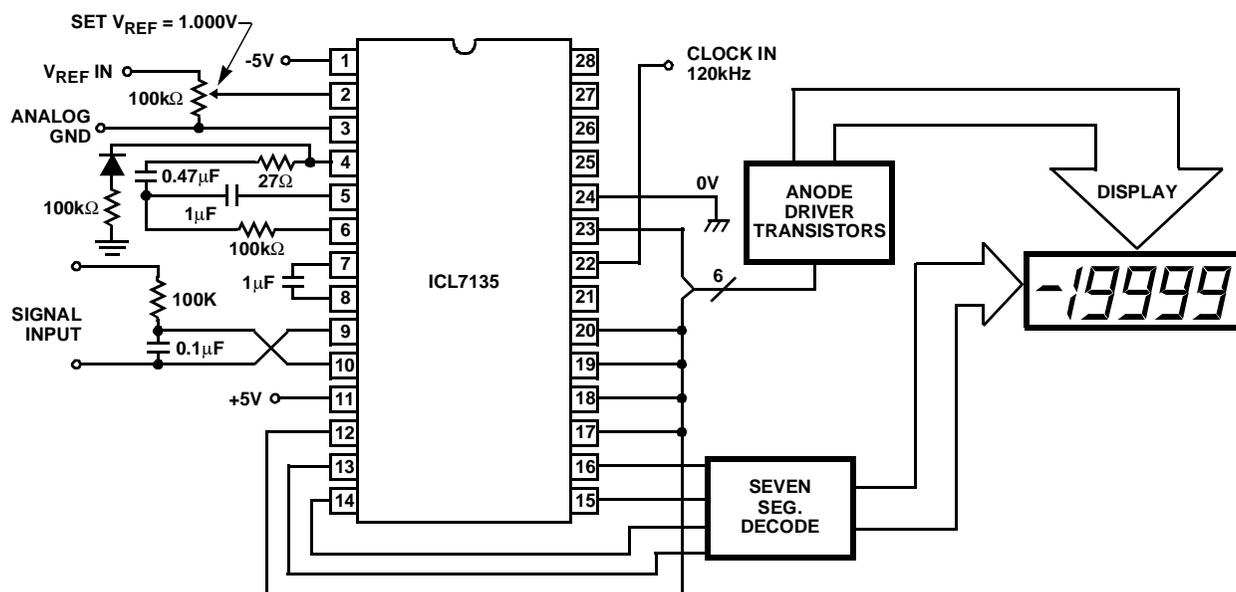Each of these errors will be analyzed for its error contribution to the converter.

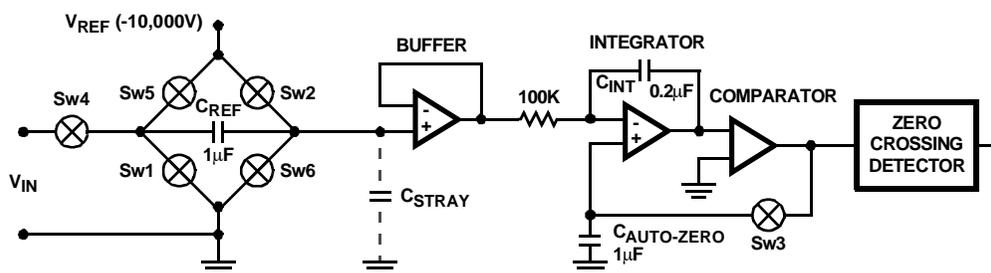**FIGURE 4. THE ICL7135 FUNCTIONAL DIAGRAM**

FIGURE 5. ANALOG SECTION OF A DUAL SLOPE CONVERTER

Capacitor Droop Due to Leakage

Typical leakage (I_{DOFF}) of the switches at normal operating voltage is 1pA each and 2pA at each input of the buffer and integrator op amps. In terms of offset voltage caused by capacitor droop, the effect of the auto-zero and reference capacitors is differential, i.e., there is no offset if they droop an equal amount. A conservative typical effect of droop on offset would be 2pA discharging 1μF for 83ms (10,000 clock periods), which amounts to an averaged equivalent of 083μV referred to the input. The effect of the droop on rollover error (difference between equal positive and negative voltages near full scale) is slightly different. For a negative input voltage, switch 5 is closed for the de-integrate cycle. Thus the reference capacitor and auto-zero capacitor operate differentially for the entire measurement cycle. For a positive voltage, switch 6 is closed and the differential compensation of the reference capacitor is lost during de-integrate. A typical contribution to rollover error is 3pA discharging 1μF capacitor for 166ms, equivalent to 0.249μV when averaged. These numbers are certainly insignificant for room

temperature leakages but even at 100°C the contributions should be only 15μV and 45μV respectively. A rollover error of 45μV is less than 0.5 counts on this 20,000 count instrument.

Charge "Suck-Out" When the Switches Turn-Off

There is no problem in charging the capacitors to the correct value when the switches are on. The problem is getting the switches off without changing this value. As the gate is driven off, the gate-to-drain capacitance of the switch injects a charge on the reference or auto-zero capacitor, changing its value. The net charge injection of switch 3 turning off can be measured indirectly by noting the offset resulting by using a 0.01μF auto-zero capacitor instead of 1.0μF. For this condition the offset is typically 250μV, and since the signal ramp is a straight line instead of a parabola the main error is due to charge injection rather than leakage. This given a net injected charge of 2.5pC or an equivalent C_{gd} of 0.16pF. The effect of switches 1, 2, 4, 5 and 6 are more complicated since they depend on timing and some switches are going on while others are going off. A substitution of an 0.01μF

capacitor for reference capacitor gives less than 100µV offset error. Thus, a conservative typical offset error for a 1.0µF capacitor is 2.5µV. There is no contribution to rollover error (independent of offset). Also this value does not change significantly with temperature.

Nonlinearity of Buffer and Integrator

In this converter, since the signal and reference are injected at the same point, the gain of the buffer and integrator are not of first-order importance in determining accuracy. This means that the buffer can have a very poor CMRR over the input range and still contribute zero error as long as it is constant, i.e., offset changes linearly with common mode voltage. The first error term is the nonlinear component of CMRR. Careful measurement of CMRR on 30 buffers indicated rollover errors from 5µV to 30µV. The contribution of integrator nonlinearity is less than 1µV in each case.

High Frequency Limitations of Amplifiers

For a zero input signal, the buffer output will switch from zero to V_{REF} (1.0V) in 0.5µs with an approximately linear response. The net result is to lose 0.25µs of de-integrate period. For a 120kHz clock, this is 3% of a clock pulse or 3µV. This is not an offset error since the delay is equal for both positive and negative references. The net result is the converter would switch from 0 to 1 at 97µV instead of 100µV in the ideal case.

A much larger source of delay is the comparator which contributes 3µs. At first glance, this sounds absolutely ridiculous compared to the few tens of nanoseconds delay of modern IC comparators. However, they are specified with 2 to 10mV of overdrive. By the time the ICL7135 comparator gets 10mV of overdrive, the integrator will have been through zero-crossing or 20 clock pulses! Actually, the comparator has a 300MHz gain-bandwidth product which is comparable to the best ICs. The problem is that it must operate on 30µV of overdrive instead of 10mV. Again, this delay causes no offset error but means the converter switches from 0 to 1 at 60µV, from 1 to 2 at 160µV, etc. Most users consider this switching at approximately $^1/_2$ LSB more desirable than the "so-called ideal" case of switching at 100µV. If it is important that switching occur at 100µV, the comparator delay may be compensated by including a small value resistor ($\cong$20Ω) in series with the integration capacitor. (Further details of this technique are given on page 5 under the heading "Maximum Clock Frequency".) The integrator time delay is less than 200ns and contributes no measurable error.

Integrating Capacitor Dielectric Absorption

Any integrating A/D assumes that the voltage change across the capacitor is exactly proportional to the integral of the current into it. Actually, a very small percentage of this charge is "used up" in rearranging charges within the capacitor and does not appear as a voltage across the capacitor. This is dielectric absorption. Probably the most accurate means of measuring dielectric absorption is to use it in a dual-slope A/D converter with $V_{IN} \equiv V_{REF}$. In this mode, the instrument should read 1.0000 independent of other component values. In very careful measurements where zero-crossings were observed in order to extrapolate a fifth digit and all delay errors were calculated out, polypropylene capacitors gave the best results. Their equivalent readings were 0.99998. In the same test polycarbonate capacitors typically read 0.9992, polystyrene 0.9997. Thus, polypropylene is an excellent choice since they are not expensive and their increased temperature coefficient is of no consequence in this circuit. The dielectric absorption of the reference and auto-zero capacitors are only important at power-on or when the circuit is recovering from an overload. Thus, smaller or cheaper capacitors can be used if very accurate readings are not required for the first few seconds of recovery.

Charge Lost by C_{REF} in Charging C_{STRAY}

In addition to leakage and switching charge injection, the reference capacitor has a third method of losing charge and, therefore, voltage. It must charge C_{STRAY} as it swings from 0 to V_{IN} to V_{REF} (Figure 5). However, C_{STRAY} only causes an error for positive inputs. To see why, let's look firstly at the sequence of events which occurs for negative inputs. During auto-zero C_{REF} and C_{STRAY} are both charged through the switches. When the negative signal is applied, C_{REF} and C_{STRAY} are in series and act as a capacitance divider. For C_{STRAY} = 15pF, the divider ratio is 0.999985. When the positive reference is applied through switch #5, the same divider operates. As mentioned previously, a constant gain network contributes no error and, thus, negative inputs are measured exactly.

For positive inputs, the divider operates as before when switching from auto-zero to V_{IN}, but the negative reference is applied by closing switch #6. The reference capacitor is not used, and therefore the equivalent divider network is 1.0000 instead of 0.999985. At full scale, this 15µV/V error gives a 30µV rollover error with the negative reading being 30µV too low. Of course for smaller C_{STRAY}, the error is proportionally less.

Summary

Error analysis of the circuit using typical values shows four types of errors. They are (1) an offset error of 2.5µV due to charge injection, (2) a full scale rollover error of 30µV due to C_{STRAY}, (3) a full scale rollover error of 5 to 30µV due to buffer nonlinearity and (4) a delay error of 40µV for the first count. These numbers are in good agreement with actual results observed for the ICL7135. Due to peak-to-peak noise of 20µV around zero, it is possible only to say that any offsets are less than 10µV. Also, the observed rollover error is typically $^1/_2$ count (50µV) with the negative reading larger than the positive. Finally, the transition from a reading of 0000 to 0001 occurs at 50µV.

These figures illustrate the very high performance which can be expected from a well designed dual-slope circuit - performance figures which can be achieved with no tricky 'tweaking' of component values. Furthermore, the circuit includes desirable convenience features such as auto-zero, auto-polarity and a single reference.

Maximum Clock Frequency

Because of the 3µs delay in the comparator, the maximum recommended clock frequency is 160kHz. In the error analysis it was shown that under these conditions half of the first reference integrate period is lost in delay. This means that the meter reading will change from 0 to 1 at 50µV, from 1 to 2 at 150µV, etc. As was noted earlier, most users consider this transition at midpoint to be desirable. However, if the clock frequency is increased appreciably above 160kHz, the instrument will flash 1 on noise peaks even when the input is shorted.

The clock frequency may be extended above 160kHz, however, by using a low value resistor in series with the integration capacitor. The effect of the resistor is to introduce a small pedestal voltage on to the integrator output at the beginning of the reference integrate phase (Figure 6). By careful selection of the ratio between this resistor and the integrating resistor (a few tens of ohms in the recommended circuit), the comparator delay can be compensated and the maximum clock frequency extended by approximately a factor of 3. At higher frequencies, ringing and second order breaks will cause significant nonlinearities in the first few counts of the instrument.
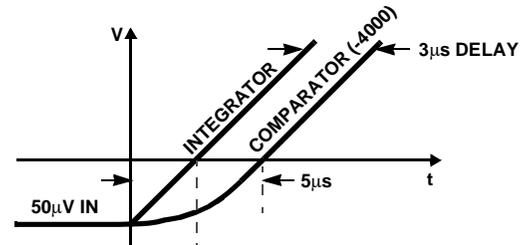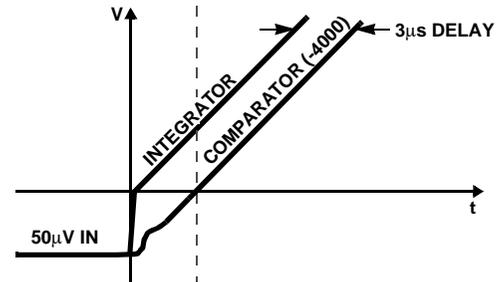
FIGURE 6A. UNCOMPENSATED

FIGURE 6B. COMPENSATED

FIGURE 6. INTEGRATOR AND COMPARATOR OUTPUTS FOR UNCOMPENSATED AND COMPENSATED SYSTEMS

Noise

The peak-to-peak noise around zero is approximately 20µV (peak-to-peak value not exceeded 95% of the time). Near full scale, this value increases to approximately 40µV.

Since much of the noise originates in the auto-zero loop, some improvement in noise can be achieved by putting gain in the buffer. A gain of about 5X is optimum. Too much gain will cause the auto-zero switch to misbehave, because the amplified V_{OS} of the buffer will exceed the switch operating range.

intersil