DALLAS SEMICONDUCTOR **/VI/IXI/VI**

Sep 01, 2000

INL/DNL Measurements for High-Speed Analog-to-Digital Converters (ADCs)

Although integral and differential nonlinearity may not be the most important parameters for high-speed, high dynamic performance data converters, they gain significance when it comes to high-resolution imaging applications. The following application note serves as a refresher course for their definitions and details two different, yet commonly used techniques to measure INL and DNL in high-speed analog-to-digital converters (ADCs).

Manufacturers have recently introduced high-performance analog-to-digital converters (ADCs) that feature outstanding static and dynamic performance. You might ask, "How do they measure this performance, and what equipment is used?" The following discussion should shed some light on techniques for testing two of the accuracy parameters important for ADCs: integral nonlinearity (INL) and differential nonlinearity (DNL).

Although INL and DNL are not among the most important electrical characteristics that specify the high-performance data converters used in communications and fast data-acquisition applications, they gain significance in the higher-resolution imaging applications. However, unless you work with ADCs on a regular basis, you can easily forget the exact definitions and importance of these parameters. The next section therefore serves as a brief refresher course.

INL and DNL Definitions

DNL error is defined as the difference between an actual step width and the ideal value of 1LSB (see Figure 1a). For an ideal ADC, in which the differential nonlinearity coincides with DNL = 0LSB, each analog step equals 1LSB ($1LSB = V_{FSR}/2^N$, where V_{FSR} is the full-scale range and N is the resolution of the ADC) and the transition values are spaced exactly 1LSB apart. A DNL error specification of less than or equal to 1LSB guarantees a monotonic transfer function with no missing codes. An ADC's monotonicity is guaranteed when its digital output increases (or remains constant) with an increasing input signal, thereby avoiding sign changes in the slope of the transfer curve. DNL is specified after the static gain error has been removed. It is defined as follows:

$$DNL = |[(V_{D+1} - V_D)/V_{LSB-IDEAL} - 1]|, \text{ where } 0 < D < 2^N - 2.$$

V_D is the physical value corresponding to the digital output code D, N is the ADC resolution, and $V_{LSB-IDEAL}$ is the ideal spacing for two adjacent digital codes. By adding noise and spurious components beyond the effects of quantization, higher values of DNL usually limit the ADC's performance in terms of signal-to-noise ratio (SNR) and spurious-free dynamic range (SFDR).
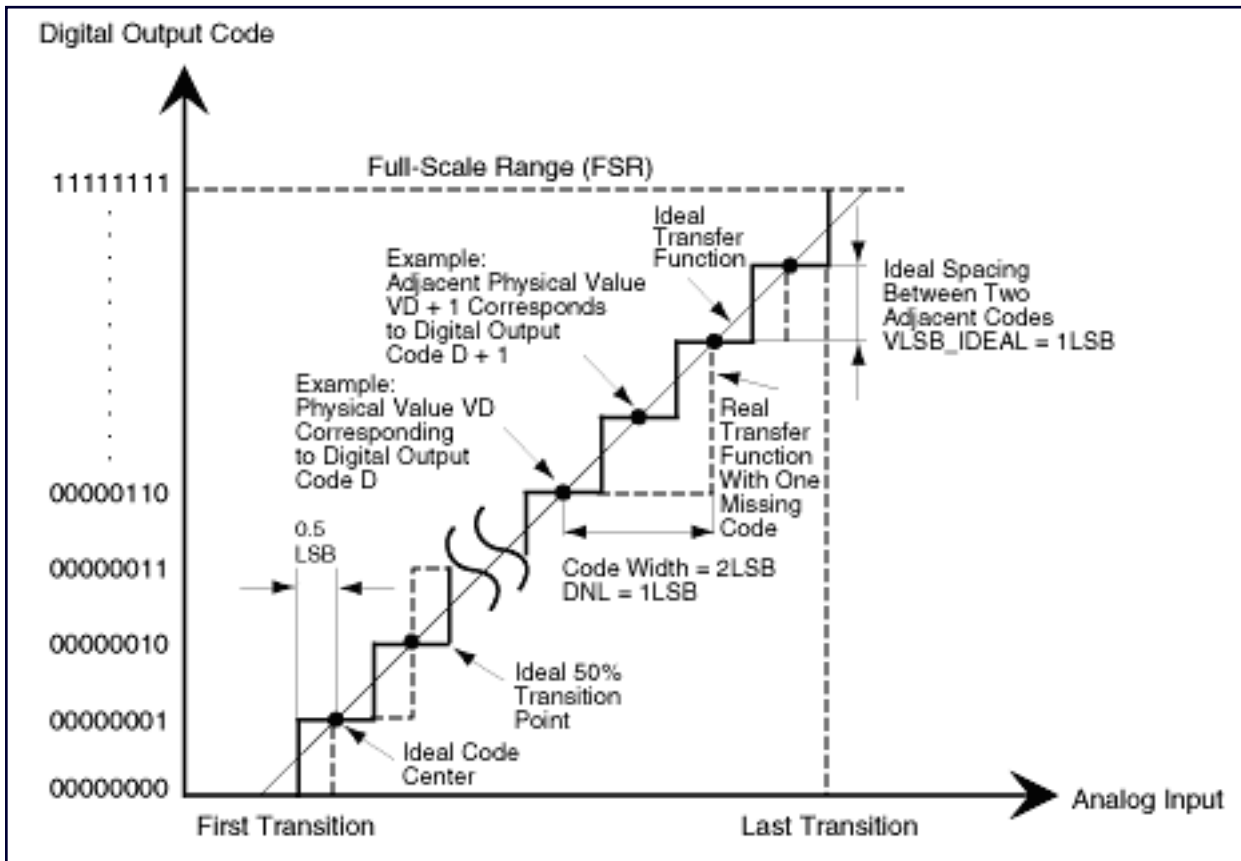

Figure 1a. To guarantee no missing codes and a monotonic transfer function, an ADC's DNL must be less than 1LSB.

INL error is described as the deviation, in LSB or percent of full-scale range (FSR), of an actual transfer function from a straight line. The INL-error magnitude then depends directly on the position chosen for this straight line. At least two definitions are common: "best straight-line INL" and "end-point INL" (see Figure 1b):

- **Best straight-line INL** provides information about offset (intercept) and gain (slope) error, plus the position of the transfer function (discussed below). It determines, in the form of a straight line, the closest approximation to the ADC's actual transfer function. The exact position of the line is not clearly defined, but this approach yields the best repeatability, and it serves as a true representation of linearity.

- **End-point INL** passes the straight line through end points of the converter's transfer function, thereby defining a precise position for the line. Thus, the straight line for an N-bit ADC is defined by its zero (all zeros) and its full-scale (all ones) outputs.

The best straight-line approach is generally preferred, because it produces better results. The INL specification is measured after both static offset and gain errors have been nullified, and can be described as follows:

$$INL = \left| [(V_D - V_{ZERO})/V_{LSB-IDEAL}] - D \right|, \text{ where } 0 < D < 2^N\text{-}1.$$

V_D is the analog value represented by the digital output code D, N is the ADC's resolution, V_{ZERO} is the minimum analog input corresponding to an all-zero output code, and $V_{LSB-IDEAL}$ is the ideal spacing for two adjacent output codes.
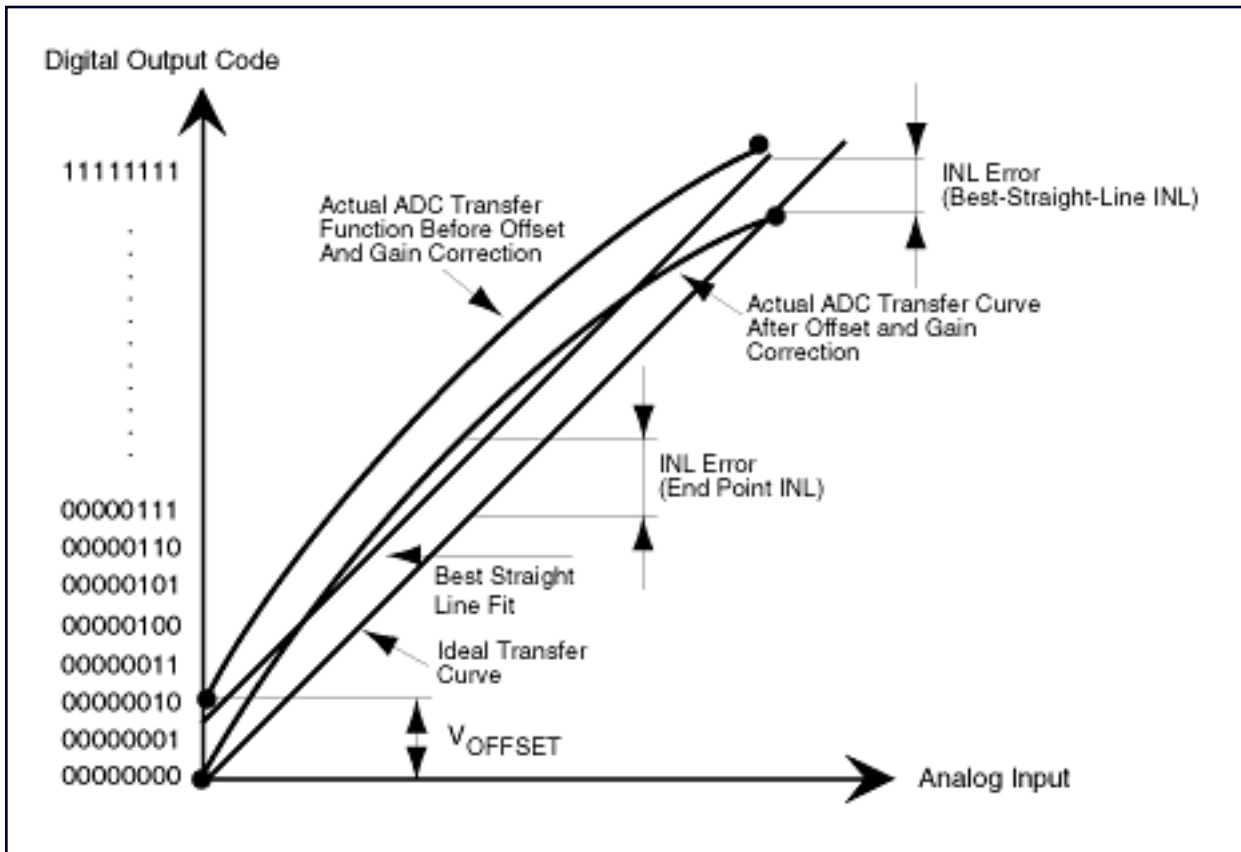

Figure 1b. Best straight-line and end-point fit are two possible ways to define the linearity characteristic of an ADC.

Transfer Function

The transfer function for an ideal ADC is a staircase in which each tread represents a particular digital output code and each riser represents a transition between adjacent codes. The input voltages corresponding to these transitions must be located to specify many of an ADC's performance parameters. This chore can be complicated, especially for the noisy transitions found in high-speed converters and for digital codes that are near the final result and changing slowly.

Transitions are not sharply defined, as shown in Figure 1b, but are more realistically presented as a probability function. As the slowly increasing input voltage passes through a transition, the ADC converts more and more frequently to the next adjacent code. By definition, the transition corresponds to that input voltage for which the ADC converts with equal probability to each of the flanking codes.

The Right Transition

A transition voltage is defined as the input voltage that has equal probabilities of generating either of the two adjacent codes. The nominal analog value, corresponding to the digital output code generated by an analog input in the range between a pair of adjacent transitions, is defined as the midpoint (50% point) of this range. If the limits of the transition interval are known, this 50% point is calculated easily. The transition point can be determined at test by measuring the limits of the transition interval, and then dividing the interval by the number of times each of the adjacent codes appears within it.

Generic Setup for Testing Static INL and DNL

INL and DNL can be measured with either a quasi-DC voltage ramp or a low-frequency sine wave as the input. A simple DC (ramp) test can incorporate a logic analyzer, a high-accuracy DAC (optional), a high-precision DC source for sweeping the input range of the device under test (DUT), and a control interface to a nearby PC or X-Y plotter.

If the setup includes a high-accuracy DAC (much higher than that of the DUT), the logic analyzer can monitor offset and gain errors by processing the ADC's output data directly. The precision signal source creates test voltages for the DUT by sweeping slowly through the input range of the ADC from zero scale to full scale. Once reconstructed by the DAC, each test voltage at the ADC input is subtracted from its corresponding DC level at the DAC output, producing a small voltage difference (V_{DIFF}) that can be displayed with an X-Y plotter and linked to the INL and DNL errors. A change in quantization level indicates differential nonlinearity, and a deviation of V_{DIFF} from zero indicates the presence of integral nonlinearity.

Analog Integrating Servo Loop

Another way to determine static linearity parameters for an ADC, similar to the preceding but more sophisticated, is using an analog integrating servo loop. This method is usually reserved for test setups that focus on high-precision measurements rather than speed.

A typical analog servo loop (see Figure 2) consists of an integrator and two current sources connected to the ADC input. One source forces a current into the integrator, and the other serves as a current sink. A digital magnitude comparator connected to the ADC output controls both current sources. The other input of the magnitude comparator is controlled by a PC, which

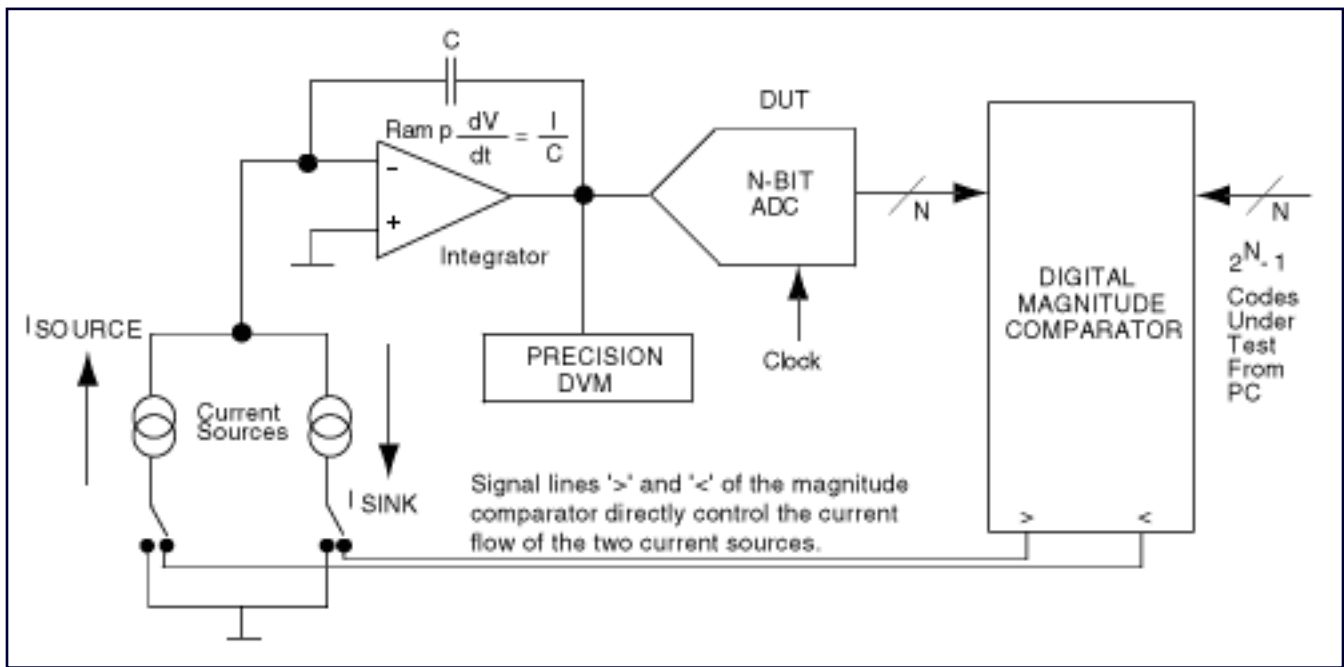sweeps it through the $2^N - 1$ test codes for an N-bit converter.

Figure 2. This circuit configuration is an analog integrating servo loop.

If the polarity of feedback around the loop is correct, the magnitude comparator causes the current sources to servo the analog input around a given code transition. Ideally, this action produces a small triangular wave at the analog inputs. The magnitude comparator controls both rate and direction for these ramps. The integrator's ramp rate must be fast when approaching a transition, yet sufficiently slow to minimize peak excursions of the superimposed triangular wave when measuring with a precision digital voltmeter (DVM).

For INL/DNL tests on the MAX108, the servo-loop board connects to the evaluation board through two headers (see Figure 3). One header establishes a connection between the MAX108's primary (or auxiliary) output port and the magnitude comparator's latchable input port (P). The second header ensures a connection between the servo loop (the magnitude comparator's Q port) and a computer-generated digital reference code.
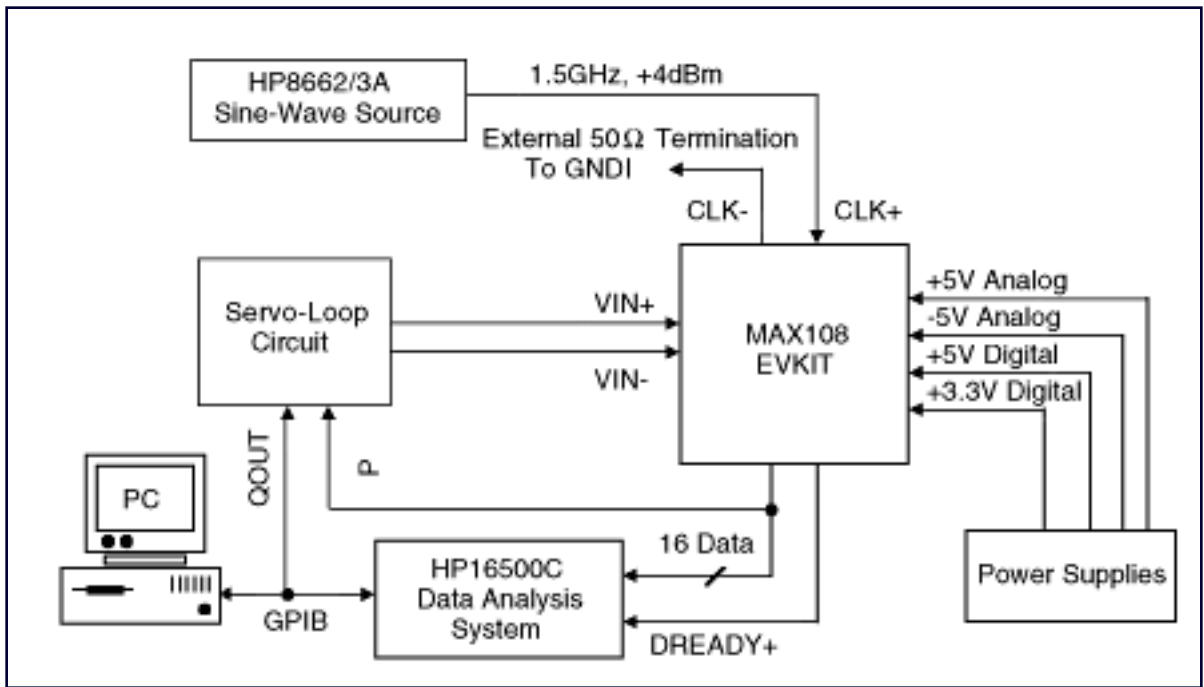
Figure 3. With the aid of the MAX108EVKIT and an analog integrating servo loop, this test setup determines the MAX108's INL and DNL characteristics.

The fully decoded decision resulting from this comparison is available at the comparator output P > QOUT, and is then passed on to the integrator configurations. Each comparator result controls the logic input of the switch independently and generates voltage ramps as required to drive succeeding integrator circuits for both inputs of the DUT. This approach has its advantages, but it also has several drawbacks:

- The triangular ramp should have low dV/dt to minimize noise. This condition generates repeatable numbers, but it results in long integration times for the precision meter.

- Positive and negative ramp rates must be matched to arrive at the 50% point, and the low-level triangular waves must be averaged to achieve the desired DC level.

- Integrator designs usually require careful selection of the charge capacitors. To minimize potential errors due to the capacitors' "memory effect," for instance, select integrator capacitors with low dielectric absorption.

- Accuracy is proportional to the integration period and inversely proportional to the settling time.

A DVM connected to the analog integrated servo loop measures the INL/DNL error versus output code (Figures 4a and 4b). Note that a parabolic or bow shape in the plot of "INL vs. output code" indicates the predominance of even-order harmonics, and an "S shape" indicates the predominance of odd-order harmonics.
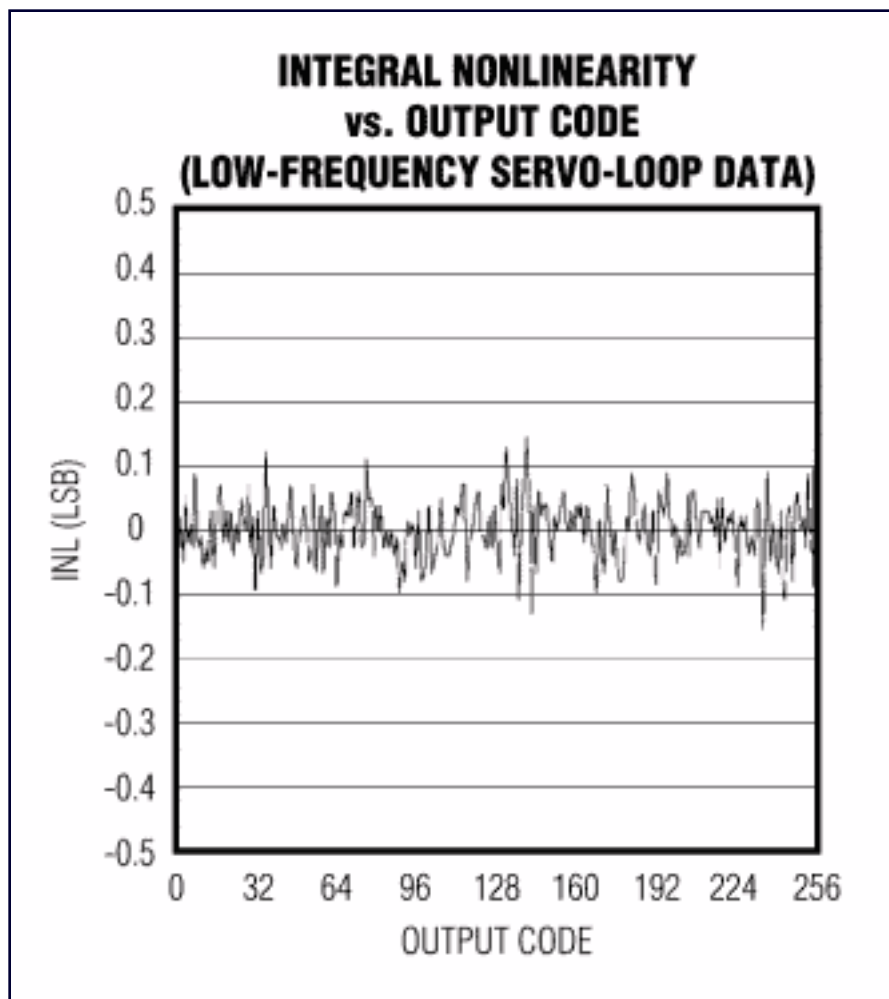
Figure 4a. This plot shows typical integral nonlinearity for the MAX108 ADC, captured with the analog integrating servo loop.
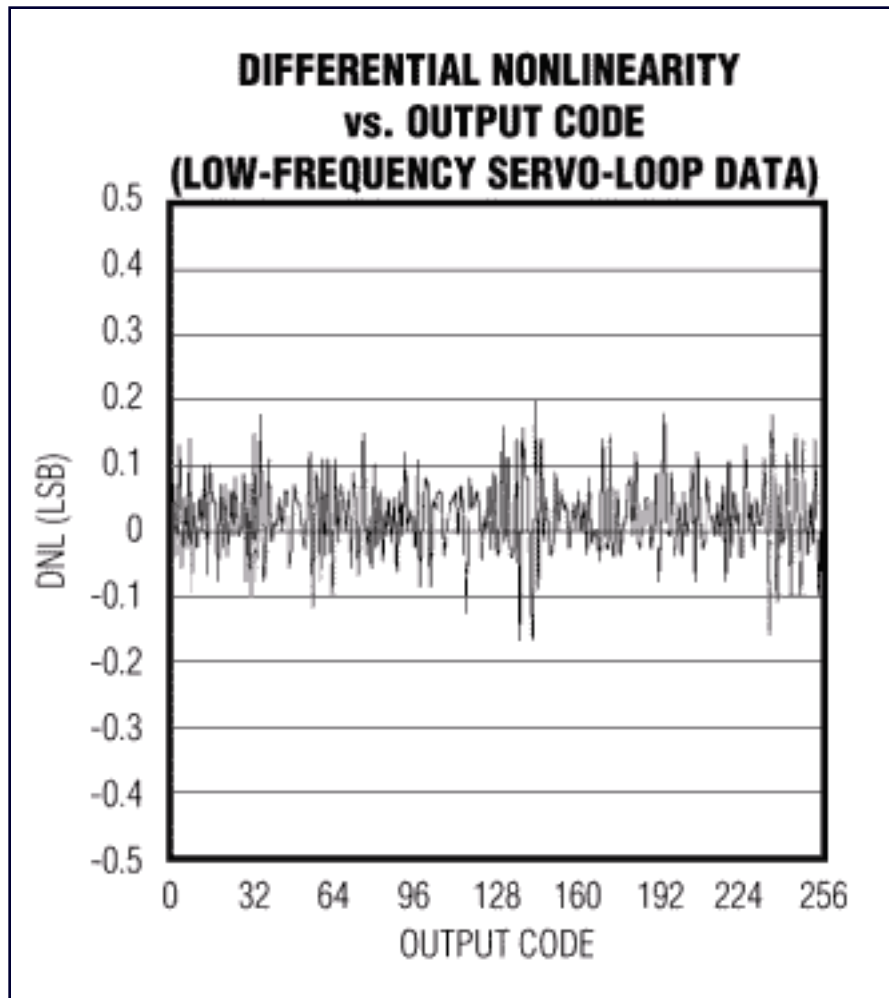
Figure 4b. This plot shows typical differential nonlinearity for the MAX108, captured with the analog integrating servo loop.

To eliminate negative effects in the previous approach, you can replace the servo loop's integrator section with an L-bit successive-approximation register (SAR) that captures the DUT's output codes, an L-bit DAC, and a simple averaging circuit. Together with the magnitude comparator, this circuit forms a SAR-type converter configuration (see Figure 5 and "SAR Converter" discussion below), in which the magnitude comparator programs the DAC, reads its outputs, and performs a successive approximation. Meanwhile, the DAC presents a high-resolution DC level to the input of the N-bit ADC under test. In this case, a 16-bit DAC was chosen to trim the ADC to 1/8LSB accuracy and obtain the best possible transfer curve.
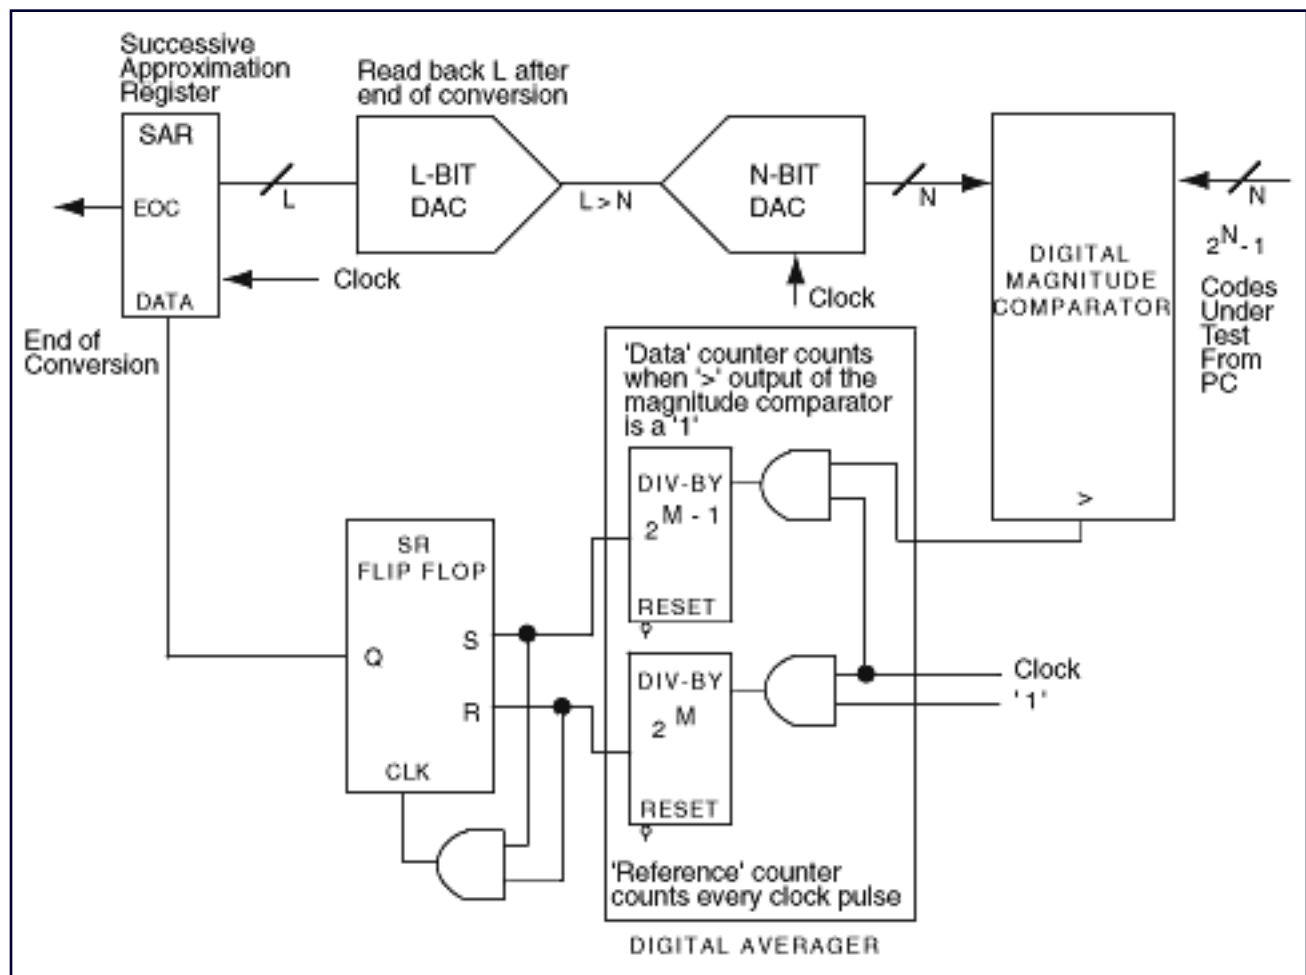
Figure 5. Successive approximation and a DAC configuration replace the integrator section of the analog servo loop.

The advantage of an averaging circuit is apparent when noise causes the magnitude comparator to toggle and become unstable, as it does on approaching its final result. Two divide-by counters are included in the averaging circuit. The "reference" counter has a period of 2^M clock cycles, where M is a programmable integer governing the period (and hence the test time). A "data" counter, which increments only when the magnitude comparator output is high, has a period equal to one-half of the first 2^{M-1} cycles.

Together, the reference and data counters average the number of highs and lows, store the result in a flip-flop, and pass it on to the SAR register. This procedure is repeated 16 times (in this case) to generate the complete output code word. Like the previous method, this one has advantages and disadvantages:

- The test setup's input voltage is defined digitally, allowing easy modification of the number of samples over which the result is to be averaged.

- The SAR approach provides a DC level rather than a ramp at the DUT's analog input.

- As a disadvantage, the DAC in the feedback loop sets a finite limit on resolution of the

input voltage.

SAR Converter

A SAR converter works like the old-fashioned chemist's balance. On one side is the unknown input sample, and on the other is the first weight generated by the SAR/DAC configuration (the most significant bit, which equals half of the full-scale output). If the unknown weight is larger than 1/2FSR, this first weight remains on the balance and is augmented by 1/4FSR. If the unknown weight is smaller, the weight is removed and replaced by a weight of 1/4FSR.

The SAR converter then determines the desired output code by repeating this procedure N times, progressing from the MSB to the LSB. N is the resolution of the DAC in the SAR configuration, and each weight represents 1 binary bit.

Dynamic Testing of INL and DNL

To assess an ADC's dynamic nonlinearity, you can apply a full-scale sinusoidal input and measure the converter's signal-to-noise ratio (SNR) over its entire full-power input bandwidth. The theoretical SNR for an ideal N-bit converter (subject only to quantization noise, with no distortion) is as follows:

$$SNR \text{ (in dB)} = N*6.02 +1.76.$$

Embedded in this figure of merit are the effects of glitches, integral nonlinearity, and sampling-time uncertainty. You can obtain additional linearity information by performing the SNR measurement at a constant frequency and as a function of the signal amplitude. Sweeping the entire amplitude range, for example, from zero to full scale and vice versa, produces large deviations from the source signal, as source amplitude approaches the converter's full-scale limit. To determine the cause of these deviations, while ruling out the effects of distortion and clock instability, use a spectrum analyzer to analyze the quantization-error signal as a function of frequency.

Countless other approaches are available for testing the static and dynamic INL and DNL of both high- and low-speed data converters. The intent here has been to give you a better understanding of the generation of powerful TOCs (typical operating characteristics) using tools and techniques that are simple but still smart and precise.

References

1. MAX108 data sheet, Rev. 1, 5/99, Maxim Integrated Products.
2. MAX108EVKIT data sheet, Rev. 0, 6/99, Maxim Integrated Products.
3. *Analog Integrated Circuit Design,* D. Johns & K. Martin, John Wiley & Sons Inc., 1997.
4. *Low-Voltage/Low-Power Integrated Circuits and Systems, Low-Voltage Mixed-Signal Circuits,* E. Sanchez-

Sinencio & A. G. Andreou, IEEE Press Marketing, 1999.

5. *Integrated Analog-to-Digital and Digital-to-Analog Converters,* R. van de Plasche, Kluwer Academic Publishers, 1994.

A177, September 2000

More Information

MAX104: QuickView -- Full (PDF) Data Sheet

MAX105: QuickView -- Full (PDF) Data Sheet -- Free Samples

MAX106: QuickView -- Full (PDF) Data Sheet

MAX107: QuickView -- Full (PDF) Data Sheet -- Free Samples

MAX108: QuickView -- Full (PDF) Data Sheet