D. Asta and F.H. Irons

Dynamic Error Compensation of Analog-to-Digital Converters

Analog-to-digital converters (ADC) produce nonlinear distortion in the digitization of an applied signal. Spurious signal components produced by ADC distortion mask low-amplitude applied signal components. As a result, ADC distortion limits the achievable dynamic range in many high-performance digital communication systems. Limitations imposed by current ADC technology dictate the bandwidth or dynamic range performance of both new and existing systems. We describe a dynamic error-correction technique that significantly reduces the level of nonlinear distortion observed in a digitized signal. The technique, called phase-plane compensation, corrects the ADC output samples by using a two-dimensional lookup table. The table accounts for both static and dynamic error sources within the ADC. Several applications of phase-plane compensation to commercially available ADCs are described. One example yields a dynamic range improvement of 15 dB (more than 2 bits) over an input frequency range spanning 4 MHz. The examples highlight the strengths of the technique and topics for further investigation.

Analog-to-digital converters (ADC) are an essential weak link in any digital system that processes analog signals. Receivers often perform in a signal environment where the interfering signals are much stronger than the desired signal. Spurious signal components produced by ADC distortion in the conversion mask low-amplitude applied signal components. The goal of an ideal ADC is to reproduce faithfully the desired signal without nonlinear distortion. At present, ADCs limit dynamic range (which is directly related to nonlinear distortion) in most high-performance systems, and much interest exists in ADCs that provide a large distortion-free dynamic range. This article focuses specifically on the use of ADCs in spectral analysis applications. The term *spectral analysis* refers to Fourier analysis of finite-length data records with a discrete Fourier transform (DFT). Spectral analysis applications include radar, power spectrum estimation, adaptive array processing, and spread spectrum communications.

A primary goal of spectral analysis is to identify and separate signals that are simultaneously present in a data segment. If the ADC distorts the applied signal before or during conversion, subsequent spectral analysis will indicate the presence of spurious signals that are not actually present in the applied signal. Spurious signals can mask or degrade the detection of weak components within an applied signal. Moreover, any signal processing following the spectral decomposition will treat the spurious products as valid signals. Spurious signals degrade the performance of spectral analysis systems because of the additional computation they require and because of the false alarms they produce.

To keep system degradation due to spurious signals at an acceptable level, an ADC must accurately quantize the applied signal. Because of nonlinear distortion, the ADC is not completely effective in performing quantization. Currently, in any digital communication system the ADC is viewed as the dominant limitation in linear performance and system bandwidth.

To improve performance, an additive correction applied to the ADC output samples after the quantization process attempts to compensate

for ADC nonlinearities. This article describes a compensation technique, called phase-plane compensation, that was developed at Lincoln Laboratory specifically for spectral analysis applications. The new technique, which is an extension of previously developed static correction techniques, corrects conversion error by using the instantaneous value and rate of change of the signal itself.

To provide the groundwork for a description of compensation techniques, the next section reviews key quantization principles. The review includes discussions of quantization principles, error mechanisms in real ADCs, performance measures, and test techniques for ADCs.

Quantization Principles

An ADC represents an analog signal as a sequence of binary words. While several existing methods represent analog data digitally, the predominant method employed in high-speed converters encodes the data as an approximation to a linear transfer characteristic. The input voltage range of the converter, V_{fs}, is subdivided into uniform intervals, called quantization intervals, q, where $q = V_{fs}/2^N$. Ideally, an ADC encodes any input voltage within a given quantization interval with the digital code corresponding to that interval. The number of bits used to represent the converter output codes determines the resolution of the ADC. A converter with N-bit resolution produces 2^N output codes. Figure 1(a) illustrates the quantization process, which can be modeled as a staircase-like input-output transfer characteristic. Positive input voltages that exceed $V_{fs}/2$ are represented by the most positive code, while negative voltages less than $-V_{fs}/2$ are represented by the most negative code. The limiting behavior that occurs for large positive and negative input voltages mimics saturation in amplifiers and produces clipped outputs in an analogous manner. The voltage V_{fs} is called the full-scale voltage range of the ADC. The power level P_{fs} associated with a full-scale sinusoidal input signal is $V_{fs}^2/8R$, where R is the resistance seen at the ADC input.

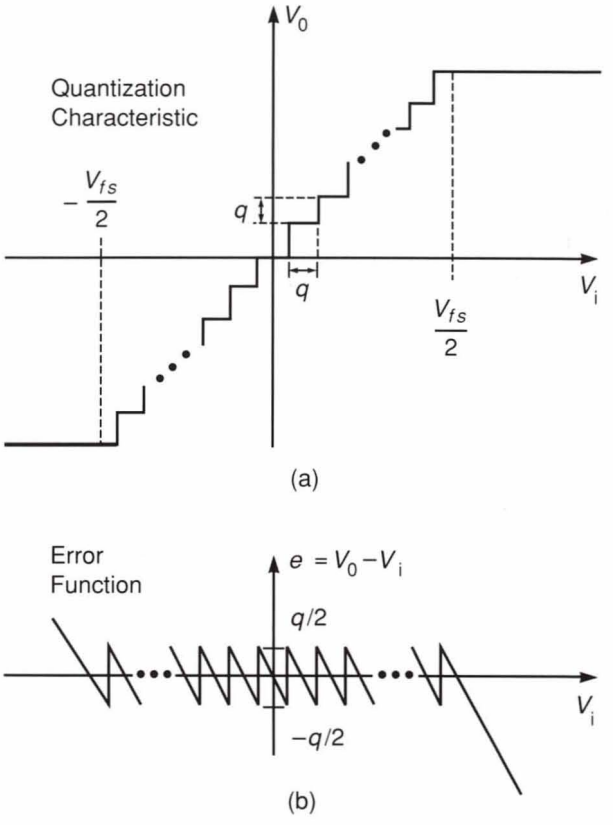

Fig. 1—Illustration of the waveform quantization process. (a) Transfer characteristic for an ideal quantizer. Inputs outside the range $[-V_{fs}/2, +V_{fs}/2]$ *are represented by minimum and maximum output codes. The quantization interval is denoted by* q. *(b) Quantization error function corresponding to the transfer characteristic given in (a).*

Properties of the Quantization Error

The error produced by the quantization process is a deterministic function of the input voltage, as illustrated in Figure 1(b). If the ADC input is a periodic time function, the resulting error waveform is also periodic, and it can be expanded into a Fourier series based on the period of the input signal. The power present in the error waveform is concentrated in a line spectrum harmonically related to the input frequency, so that the quantization appears as harmonic distortion. Harmonic frequencies above the Nyquist frequency are aliased back into the Nyquist band, which causes all of the error-waveform power to be present in the fre-

quency interval $[-f_s/2, +f_s/2]$.

While the ADC quantization error is a deterministic function of the input signal, it is convenient to describe the error statistically. The input signal is viewed as a random process with specific statistical properties, which allows a statistical determination of the quantization error. For a wide class of inputs the error waveform is assumed to be uniformly distributed on the interval $[-q/2, +q/2]$. With uniform distribution, the root-mean-squared (rms) quantization noise power, P_q, is given by

$$P_q = \frac{q^2}{12R} , \qquad (1)$$

where R is the input resistance of the ADC. Equation 1 is derived from random input signals, but it accurately describes many deterministic waveforms as well, including sinusoids. Equation 1 depends strongly on the condition that clipping effects are negligible. For random input signals, this condition on input signal power level is equivalent to requiring that the probability of clipping be small enough to be neglected. Equation 1 leads to the well-known expression for output signal-to-noise ratio (SNR) produced by an ideal N-bit converter with a full-scale sinusoidal input

$$\text{SNR} = 6.02N + 1.76 \text{ dB}. \qquad (2)$$

While the assumption of random signal behavior allows the derivation of Eq. 2 as an accurate expression for SNR, the ADC output error for a sinusoidal excitation is still composed of harmonics of the input frequency. If the power of all these harmonics is summed, the resulting quantity will then be the rms quantization power given by Eq. 1.

Effects of Additive Input Noise

Figure 1(b) shows that the error waveform produced by the quantization of a periodic waveform is itself periodic. The error waveform, which appears as harmonic distortion at the ADC output, is undesirable in spectral analysis applications. For high dynamic-range systems, the quantization noise power should be spread uniformly over the Nyquist band rather than concentrated in a discrete set of spectral lines. The harmonics present in the error waveform are coherently related to the ADC input signal; this presence complicates the task of processing the signal spectrum to differentiate between actual input signals and system-induced spurious components.

To randomize the quantization error, the ADC input signal must contain a random noise component with rms power $\geq P_q$ and an essentially flat power spectrum over the Nyquist bandwidth [1]. While setting the additive noise power equal to P_q degrades the SNR by 3 dB, it also virtually destroys any statistical correlation between sample errors. The composite noise spectrum at the ADC output is thus flat over the Nyquist bandwidth with spectral density $(P_n + P_q)/f_s$, where P_n is the power of the additive noise and f_s is the sampling frequency.

In receiver applications, the randomization discussed above can be accomplished by adjusting the receiver front-end gain, so that the noise present at the ADC input is at the correct level to randomize the quantization error. This design procedure, given the individual performance limitations of the receiver front end and the ADC, results in a configuration that optimizes system dynamic range. The actual noise level that provides the best performance will vary, depending on the particular choice of ADC and the specifics of the intended application. Typically, P_n is chosen to be several dB above P_q.

Non-Ideal ADC Behavior

The above description of quantization effects is based on an ideal ADC. Real-world devices, however, deviate significantly from ideal behavior. Errors in ADC performance can be divided into two categories—static and dynamic. Static errors affect the ADC performance for DC and low-frequency inputs, while dynamic errors affect performance at higher input frequencies. While dynamic errors can theoretically cancel static errors at a given input frequency, dynamic errors almost always further degrade performance from the level imposed by the static errors. For any performance measure chosen, ADC performance is best at low input frequencies,

but it degrades as the input frequency approaches the Nyquist frequency.

Static errors significantly affect the spectral purity of the ADC output, particularly those that result in nonuniform spacing of the quantization thresholds. Dynamic error sources include harmonic distortion in the sample/hold buffer amplifier, signal-dependent sample jitter, dynamic settling of internal amplifier stages, and frequency-dependent variations in quantization thresholds.

Given the sources of both static and dynamic errors in an ADC, optimum dynamic range in the frequency domain is usually obtained at operating power levels below P_{fs}. An ideal ADC yields the highest dynamic range when operated at full scale. Real ADC performance usually diminishes rapidly as the input power approaches full scale. As a result, the optimum dynamic range is found at an empirically determined input power level below full scale. The tests described in the next section quantify the dynamic range performance for the ADC under test, and determine the optimum operating power level.

ADC Performance Tests

Because of the wide range of static and dynamic error sources for ADCs, performance testing on a given converter should mimic the intended application as closely as possible [2–5]. Unfortunately, no direct quantitative method exists for translating performance measured by one test into performance measured by another. The following section describes tests particularly suited to evaluation of ADC performance in spectral analysis applications requiring maximum instantaneous dynamic range. Examples of such applications are communication receivers, adaptive array processors, radars, and digital spectrum analyzers.

Histogram Tests

The histogram is a graphical representation of the measured likelihood of occurrence of each ADC state (or code). The x-axis represents the code values, while the y-axis represents the probability of occurrence of the given code, based on the measured data. The histogram of an ideal ADC coincides exactly with the probability density function of the input waveform after appropriate scaling. Differences between measured and ideal histograms provide an indication of the actual device linearity. Nonlinearity in an ADC transfer characteristic manifests itself as nonuniform threshold spacing. A quantization interval that is larger than the ideal value exhibits a higher than theoretical probability of occurrence, and is often referred to as a *long code*. Similarly, a quantization interval that is less than the ideal value and exhibits a lower than theoretical probability of occurrence is an example of a *short code*. A quantization interval of zero is called a *missing code*.

Frequency Domain Testing

Frequency domain tests characterize the spectral purity of an ADC output. All frequency domain testing performed by Lincoln Laboratory is based on computation of the DFT, which is implemented by an FFT algorithm. For narrowband tests, the input signal is a single tone or pair of tones. For the test to be valid, the tones must have spectral purity that exceeds the level to be measured by approximately 10 dB. When this condition is met, spurious signals identified in the DFT plot can be assumed to be due entirely to ADC nonlinearities.

In applying the DFT to processing signals of finite time duration, a window function must be applied to the collected data. This procedure minimizes the effects of spectral leakage that otherwise invalidates spectral measurements. Detailed treatments of window functions and their properties are found in Refs. 6 and 7.

The specific choice of input frequency (relative to the sampling rate of the ADC) is an important consideration for dynamic testing. A particular test is representative of the ADC performance only if all of the ADC codes between the maximum and minimum signal values are exercised during the course of a test. An example of a bad input frequency would be $f_s/8$, where f_s is the sample frequency. In this case, assuming that the input and ADC clock signals are phase coherent (this is always the case for testing

performed on the ADC test bed), the user observes only eight distinct samples. Thus, at most eight ADC codes are exercised. To exercise all codes, the user chooses the frequencies of all single-tone inputs to coincide exactly with DFT cells corresponding to an odd number. This procedure guarantees that the DFT length (which is a power of two) and the input signal cell number are mutually prime, which avoids degenerate cases such as that described above.

Single-Tone Spurious-Free Dynamic Range Tests

Spurious-free dynamic range (SFDR) is a parameter that characterizes the dynamic performance of an ADC in the frequency domain. It measures how well an ADC allows simultaneous detection of small amplitude signals in the presence of large amplitude signals. While many possible methods exist to quantify the SFDR parameter, the method described here evaluates the output power spectrum of the ADC with a single tone as the signal input. The SFDR is then defined as the difference in dB between the fundamental power level and the power level of the highest amplitude spur as seen at the output of the ADC. Figure 2 illustrates the determination of SFDR in a DFT plot obtained from measured ADC output data. The quantity P_{qf} denotes the average quantization noise power present in each DFT frequency cell. The behavior of the SFDR as a function of input power level is important. The user can set the signal level at the ADC input and predict the maximum expected spur level. Figure 3 illustrates a typical plot of SFDR. In Fig. 3(a), the smooth sloping curve, labeled P_f, denotes the power of the fundamental as a function of input power. As expected, this curve exhibits a slope of 1 dB per 1 dB over a wide range of input power levels. The jagged curve, labeled P_s, denotes the maximum spurious power as a function of input power. The difference between these two curves is the SFDR as a function of input power.

The SFDR curve shown in Fig. 3(b) illustrates an important difference between real and ideal ADCs. An ideal ADC always yields maximum SFDR when operated at full scale. In contrast,

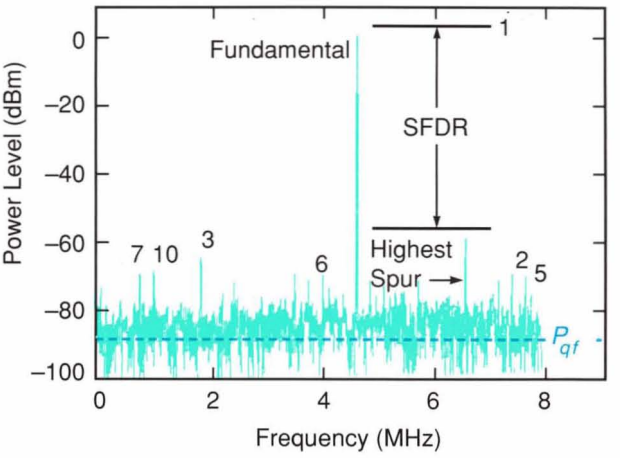

Fig. 2—Determination of spurious-free dynamic range (SFDR) from a discrete Fourier transform (DFT) magnitude plot. The SFDR is the difference in dB between the power of the fundamental signal and the largest spurious signal component. The quantity P_{qf} (blue line) is the expected quantization noise power present in each DFT frequency cell.

real ADCs yield maximum SFDR at an input power at least several dB below full scale. This discrepancy is caused by distortion mechanisms that become increasingly significant as the signal level approaches full scale.

It is tempting to conclude that the ADC can be operated at an input power level corresponding to the peak of the SFDR curve, and that the operating SFDR will then equal the peak value. In practice, however, this assumption is invalid. An operating point for an ADC is desired such that all spurs will always be less than a known threshold level. Since an actual signal environment will contain a superposition of signals of various power levels, an operating point should be selected so that spurs resulting from all signals present will be below the chosen threshold.

An operating point based on the considerations given above can be found in the plot of maximum spur power. Figure 3(a) illustrates an example for the following discussion. Begin by drawing a horizontal line that coincides with the highest spur level, P_{smax}, occurring in the linear operating range of the ADC (point *A* in the figure). The term *linear* denotes the ADC power input range for which a 1-dB change at the input

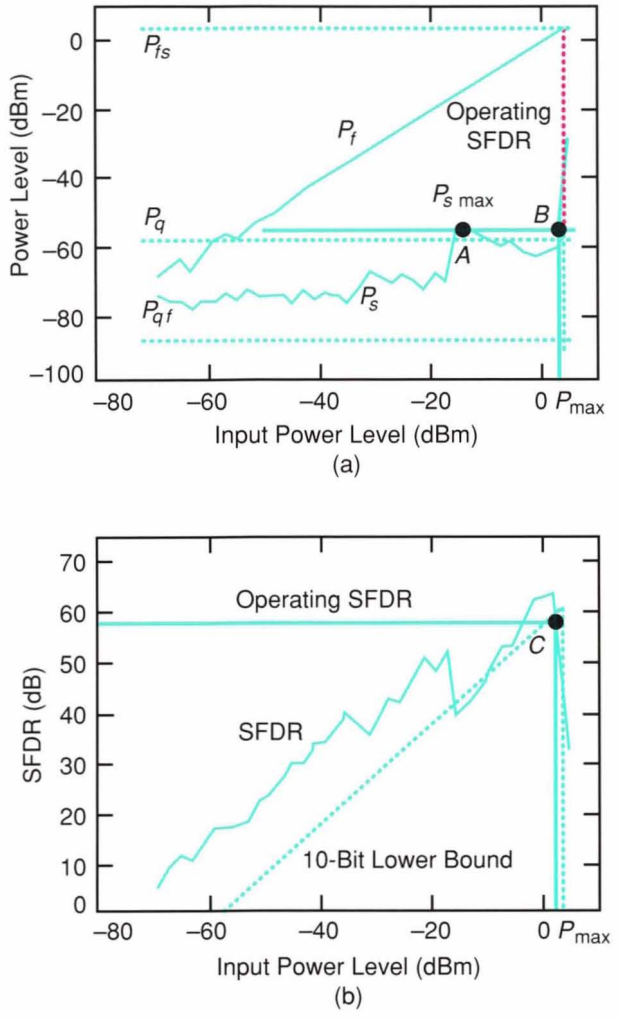

Fig. 3—Sample plots of SFDR data. (a) A plot of fundamental power P_f and maximum spurious power P_s. The operating SFDR (red line) is the difference between the P_f and P_s curves at the input power value P_{max}. The quantity P_q is the root-mean-squared (rms) quantization noise power. (b) A plot of SFDR, or the difference in dB between the P_f and P_s curves in (a). Note that the operating SFDR (point C) is less than the maximum value of the SFDR.

produces a 1-dB change in fundamental power at the output. The point where the horizontal line touches the maximum spurious-power curve P_s identifies the input power that produces the maximum spur level, P_{smax}. Next, follow the horizontal line to the right to the point where the spurious-power curve P_s crosses the horizontal line (point B in the figure). The input power corresponding to this crossover point, P_{max}, is the maximum operating power for the ADC. By obtaining P_{max}, the value of the operating SFDR can be obtained from either the maximum spurious-power plot (the difference between P_f and P_s at P_{max}) or the SFDR plot (point C in Fig. 3[b]).

Occasionally the above method fails to yield the maximum operating SFDR. This case occurs when the crossover point B is in a region where the maximum spurious power is increasing at a rate less than 1 dB per dB of input power. Since the power of the fundamental in this region increases at the 1-dB-per-dB rate, better performance is obtained for an input power level higher than indicated by the crossover point. For this case the peak of the curve describing SFDR versus power determines the maximum operating SFDR.

Residual Error Tests

Residual error is the aggregate of all quantization products, such as nonlinear distortion, quantization error, and random noise, that are not part of the desired signal. *Bits lost* is the parameter used to quantify residual error. To motivate the definition of the bits-lost parameter, recall that the mean-squared quantization error, for an ideal N-bit ADC, is given by

$$\sigma^2 = \frac{q^2}{12}\,, \tag{3}$$

where q is the quantization interval of the ADC. For a real ADC, the mean-squared residual noise can be related to an effective quantization interval, q_{eff}, of an ideal ADC; this interval would yield the same error variance. The difference between the bit resolution of this fictitious ADC and the stated resolution of the unit under test is the bits lost, b_l. The value of b_l is computed from the relation

$$b_l = \log_2\left(\frac{q_{eff}}{q}\right). \tag{4}$$

Note that one bit of resolution is lost for each doubling of q_{eff}.

To compute b_l, the actual signal must be estimated and removed from the data samples. The estimation is done by performing a least-

squares fit of a sine wave to the data. The estimate is then removed from the data to obtain the residual error waveform. The mean-squared value of the residual error is then computed, which in turn allows b_l to be calculated with Eq. 4.

Noise Power Ratio Test

The dynamic performance of an ADC with broadband inputs is characterized by a quantity known as noise power ratio (NPR). In ADC applications in which the input signal contains a large number of noncoherent tones or narrowband signals, spurs resulting from combinations of strong signal components should not interfere with detection of weaker signal components. A signal that contains a large number of frequency-multiplexed voice channels is an example. In this case, spurs resulting from strong voice signals must not excessively spill over into frequency bands occupied by other signals; this spillover limits achievable dynamic range. Since it is impossible to design a test that embodies the specific features of all possible applications, the NPR serves as a measure that characterizes ADC performance under the above operating conditions.

The measurement of NPR proceeds as follows. A random noise process is generated with an essentially flat spectrum up to a cutoff frequency, f_c, of less than half the sampling frequency. A notch filter then removes a narrow band of frequencies from the noise. To obtain a meaningful measurement, the depth of the notch in dB must be significantly greater than the maximum NPR value being measured. Also, the width of the notch should be small compared to the overall noise bandwidth. When the notched noise is applied to the ADC input, the frequency spectrum of the resulting code sequence is computed. The NPR, usually expressed in dB, is then calculated as the ratio of the power spectral density inside the notched frequency band to that outside the notch. Figure 4 shows a reconstructed sample power spectrum obtained for an 8-bit ADC. For this spectrum, the cutoff frequency is approximately 20 MHz, and the notch is centered at 2 MHz. An ideal ADC with infinite resolution would yield an output power spectrum with an infinitely deep notch at the output, resulting in an infinite NPR. For a real ADC, the notch depth is bounded by the resolution of the ADC and the input power level relative to full scale. Further degradations in NPR arise from the non-ideal behavior of the ADC.

Typically, NPR is measured as a function of the total rms input noise power. Consideration of an ideal N-bit ADC reveals insight into ideal NPR behavior. Equation 3 gives an expression for the mean-squared quantization error in terms of the quantization interval, q, where it is assumed that the effects of clipping are negligible. If V_{fs} is the full-scale peak-to-peak voltage input range of the ADC, then $q = V_{fs}/(2^N)$, where N is the ADC resolution in bits. As long as clipping effects are negligible, Eq. 3 is valid, and the quantization noise power is independent of input power level. The quantization noise has a flat power spectrum, with the noise power distributed evenly in the interval between DC and half the sampling frequency. Thus, the noise-power spectral density in the notch is specified. As the input noise power is increased, each 1-dB increase in input power yields a 1-dB increase in NPR (since the quantization-power spectral

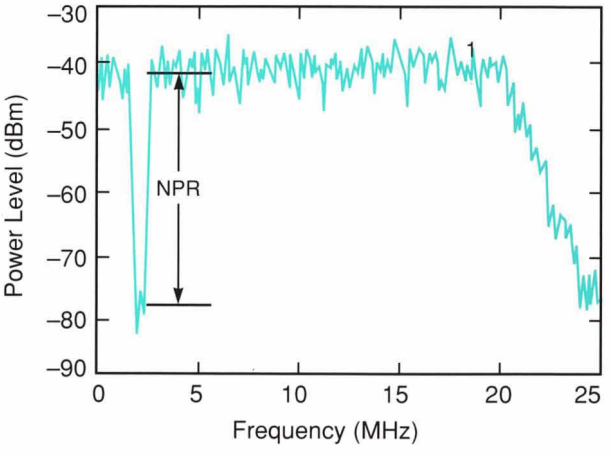

Fig. 4—Reconstructed power spectrum from an NPR test on an 8-bit ADC. The NPR is the depth in dB of the spectral notch. An ideal system with infinite resolution would produce a notch of infinite depth.

density in the notch remains constant). The NPR increases 1 dB per 1 dB until ADC clipping becomes significant, usually at an input level several dB below full scale. Once clipping occurs, the noise power in the notch is no longer constant, and it increases more rapidly than the input power level. Thus, the NPR curve rapidly approaches 0 dB for large input power levels.

The ADC Test Bed

The following section describes the test bed that was assembled for this program. It is a unique instrument not currently matched in performance by commercial instruments. The primary feature of the test bed is the ability to store 128K contiguous 16-bit samples (the current commercial maximum is 4K 16-bit words) at clock rates up to 250 MHz. Alternately, 256K contiguous 8-bit samples can be stored at clock rates up to 500 MHz. The long contiguous sample acquisition makes the test-bed instrument particularly useful in determining valid histograms and statistical parameters for ADCs of 12 or more bits.

The test-bed designers made the decision to perform only dynamic tests, thus omitting or skipping the many static tests that manufacturers use to specify resolution, linearity, and other parameters [8]. W.A. Kester and others [2–4] reported the importance of dynamic testing for video ADCs, while J.R. Naylor [5] concluded that ADCs should be tested extensively for each intended application. Our decision was to test all ADCs in an environment that emulated the intended application of acquiring digitized data for signal processing. To obtain wide dynamic signal range, all static properties must be valid, but the converse is not necessarily true. Thus, proper static performance is necessary but not sufficient to obtain proper dynamic performance.

Overall, the test system represents a substantial commitment in cost to device testing, especially when combined with the cost of software development. Currently the test system is run by menu-driven programs that allow ADC testing to be performed by any operator with minimum training and supervision.

Test-Equipment Configuration

Figure 5 shows a photograph of the ADC test bed. The test equipment is mounted in a double equipment rack with a tape recorder and computer positioned to the left of the rack. The tape recorder system has a separate general-purpose interface bus to access and store data during a test. The device under test is placed on the shelf in front of the equipment rack, along with a universal interface board customized for each unit. The left half of the equipment rack contains analog signal sources along with a dual-channel signal conditioner (filter bank) for providing inputs to the device under test. The right half of the equipment rack contains power supplies, the data-acquisition unit, and a spectrum analyzer.

A diagram of the test-equipment setup appears in Fig. 6. The figure shows that, while the computer controls the test-equipment setup, it does not control switches to route or combine signals in the system. The computer also controls the buffer or data-acquisition unit over a dedicated input-output path, which gives it direct access to the buffer memory as an extension of its own dynamic memory. The custom-designed buffer was fabricated at Lincoln Laboratory specifically for this project. It collects and stores 128K contiguous 16-bit words at clock rates up to 250 MHz. It can also collect 256K of 8-bit bytes at a clock rate up to 500 MHz. As a result, the buffer will allow testing of state-of-the-art ADCs for several years to come. If devices in excess of 500 MHz need testing, reproductions of these memory units could be multiplexed, with the additional development of a suitable high-speed controller.

Proper testing of ADCs requires precautions that should be noted. First, the ADCs are sensitive to ground loops and power-supply noise. Grounding must be handled in a consistent and careful fashion. All power-supply leads connected to the device under test are shielded twisted pairs and are cabled inside a shielded jacket. Second, test signals must be properly filtered and isolated, so that the input to the ADC is well defined and free from distortion. Figure 6 shows that the filters in the signal conditioner

Fig. 5—Photograph of the ADC test-equipment configuration.

have cutoffs at octave values with 90-dB stopbands. In addition, the test-signal noise floor should not be allowed to dominate the noise of the ADC under test. Third, test frequencies and sample clock rates must be stable, accurately defined, and free of phase noise. To achieve this result, all synthesizers and clock sources use an oven-stabilized, filtered, 10-MHz reference. The reference has a stability of 3×10^{-10} parts per day, which yields excellent stability for the collection of each buffer of data. Finally, since the FFT is used extensively to view ADC performance in the frequency domain, the data must be properly windowed. To acquire an understanding of the windowing requirements for wide dynamic range responses, see Refs. 6 and 7.

Available Tests and Procedures

The ADC test-bed structure shown in Fig. 6 determines the various tests that can be performed. By setting the signal conditioner and configuration switches, the device under test can have as input a single sinusoid, two sinusoids, a sinusoid plus noise, or noise alone. The noise source, a Marconi 2090C automated noise test set that tests telephone carrier channels, is used either for the standard noise power ratio (NPR) test [3] or for dithering the ADC input [1].

The computer is a 32-bit machine with a 12.8-MHz floating-point CPU and 10 MB of dynamic RAM. Thus a weighted 4,096-point FFT is computed in approximately 14 seconds. The menu-selected tests listed in Table 1 are available in a single program.

ADC Compensation: Previous Work

A significant difference exists between ideal and real ADC performance. To minimize this difference, error correction or compensation of the ADC output samples can yield improved performance. Compensating the ADC errors is equivalent to linearizing the transfer characteristic illustrated in Fig. 7. This model treats the

ADC as a cascade of two systems—a nonlinear system followed by an ideal quantizer. For the purpose of the present discussion, the nonlinear system is memoryless, so that the system output depends only on the current input value. The simplest improvement is to determine the actual operating characteristic of the ADC and use this knowledge to add a correction term to each ADC output sample. Several authors have implemented such a correction based on the static and memoryless operating characteristic of the ADC [9–11]. Their goal was to derive an efficient way to represent the nonlinear error of the ADC as a function of the input voltage.

Once the nonlinear transfer characteristic of the ADC is established, compensation to correct the nonlinearity can be derived. Several ADC manufacturers attempt compensation by pro-

Table 1. Menu-Selectable ADC Tests

- Waveform plotting
- FFT power spectra
- Probability density estimates
  - Sinewave
  - Noise
  - Other
- Dynamic range responses
  - Residual error
  - Noise power ratio
  - Distortion analysis
    - Fundamental
    - Harmonic power levels (first 99)
    - Maximum spur
    - Power calibration
    - SFDR

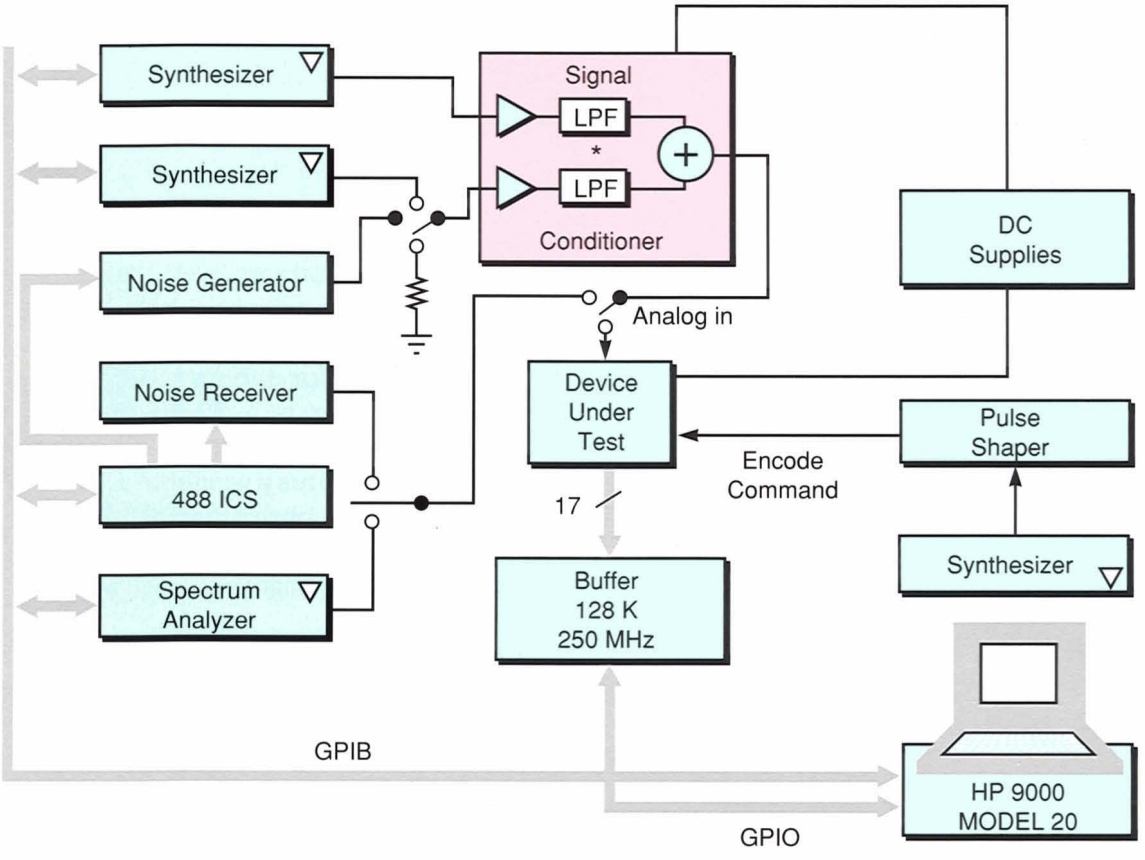

Fig. 6—ADC test-bed block diagram. Cutoff frequencies in the signal conditioner are at octave values with 90-dB stopbands. The symbol ∇ refers to a 10-MHz reference signal common to the synthesizers and the spectrum analyzer.

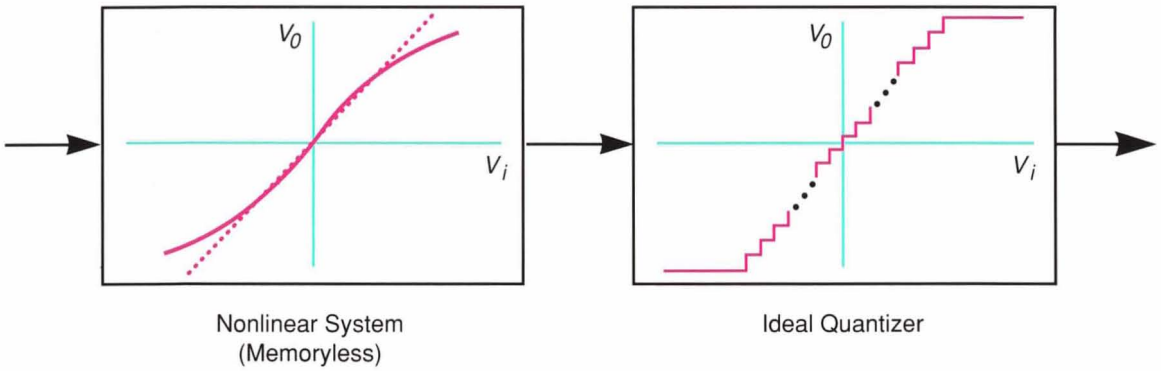

Fig. 7—Static nonlinear model for a real ADC. The model is a cascade of a memoryless nonlinear transfer function and an ideal quantizer.

gramming a ROM lookup table with values of the ADC error that correspond to each of the ADC states. The output samples of the converter point to locations in the lookup table. The error value is then added to the ADC output code to obtain the corrected sample.

Unfortunately, the technique described above is limited. The static error-modeling approach is effective only for input signal frequencies that are very low relative to the ADC input signal bandwidth. The limitation is linked to the assumption that the nonlinearity can be modeled by a static nonlinear transfer characteristic. In actual converters, the nonlinear error corresponding to a given output state depends not only on the current instantaneous input voltage but on the past history of the applied signal as well. In other words, the nonlinear system preceding the ideal quantizer is dynamic rather than static and memoryless. To obtain effective error compensation over a broad range of input frequencies, the error model must be generalized to account for dynamic effects.

Dynamic ADC Compensation

Efforts in ADC compensation at Lincoln Laboratory have focused on generalizing the ADC nonlinear error model to include dynamic effects. The work was begun by the authors [8] and developed by T.A. Rebold [12]. A further refinement, which provides a different perspective on ADC error modeling, was suggested by N.W. Spencer [13].

Figure 8 shows an extension of the ADC model described in Fig. 7. The output of the nonlinear system preceding the quantizer is a function not only of the present input, but also of the first derivative of the input signal. This model, while not a completely general nonlinear dynamic model, is more general than the static model.

A causal nonlinear dynamic system produces an output that is a function of all past input. In practical systems (ADCs included), however, the output is strongly dependent on the history of the signal for a relatively brief time just prior to and including the present instant. Thus, if a description of the input signal during this brief interval is available, the nonlinear system output can be determined to arbitrary accuracy. By considering only input signals that are well behaved (continuous and differentiable), the input signal, $v_i(t)$, can be accurately approximated by a Taylor series expansion in the vicinity of time t_0:

$$v_i(t) = v_i(t_0) + \sum_{k=1}^{\infty} a_k \, (t - t_0)^k \left. \frac{d^k v_i}{dt^k} \right|_{t = t_0} .$$

By further restricting attention to bandlimited signals (such as the input to an ADC), the Taylor series can be truncated to a finite number of terms and still retain high accuracy. The signal value and some finite number of signal derivatives therefore characterize the

ADC input signal.

To assess the validity of the model shown in Fig. 8, the user must determine when the first derivative of the input signal is sufficient to characterize the nonlinear system behavior. In ADC applications, the input signal bandwidth must be limited to less than half the sampling frequency. The input signal can be accurately modeled near the kth sampling instant, kT, as a linear ramp as long as

$$\left| t - kT_s \right| << T_s,$$

where T_s is the sampling interval. If the dynamic behavior of the nonlinear system depends primarily on the signal for a period that is also much less than the sample interval, then the instantaneous signal value and the first derivative determine the nonlinear system output.

Given the above restrictions on sampling interval, the ADC output state and the first derivative of the input signal at the time the sample was taken describe the ADC nonlinear dynamic system. This model leads naturally to the concept of a two-dimensional correction table, in contrast to the one-dimensional table used in previous work. The correction table represents the ADC error as a function of the converter output state and the slope of the analog input signal at the time the sample was taken. The two-dimensional plane on which the error surface is defined is referred to as a phase plane.

Phase-plane error modeling is central to the dynamic error-compensation technique described in this article. Some difficulties arise, however, in implementation of the compensation model. The first difficulty is in evaluating and representing the derivative of the ADC input signal. Different methods of estimating the slope give rise to different hardware implementations of the compensation algorithm. The second difficulty is the accuracy with which the error surface is represented. Since the surface is represented as digital data, it is by necessity stored as a finite number of discrete samples. Each sample represents the value of the error surface in a small rectangular region or cell of the phase plane. To be an accurate representation, the error surface must be essentially constant everywhere within each cell.

Compensated ADC Architectures

Estimating the slope of the ADC input signal is a significant step in representing the error surface. Several implementations are available to perform the estimation.

Figure 9 shows a straightforward implementation of the error correction scheme, which we call compensation structure 1. The output samples of the ADC are used to estimate the analog slope of the input signal by using the central difference estimator (CDE):

$$\hat{\dot{x}}_k = \frac{x_{k+1} - x_{k-1}}{2T_s}.$$

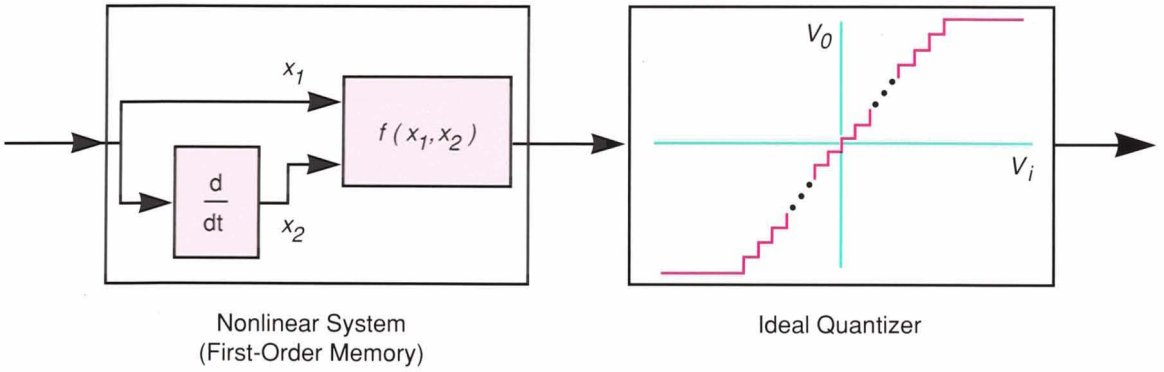

Fig. 8—Dynamic nonlinear model for a real ADC. The model is again a cascade of a nonlinear transfer function and an ideal quantizer, but the model now includes a first-order memory characteristic.

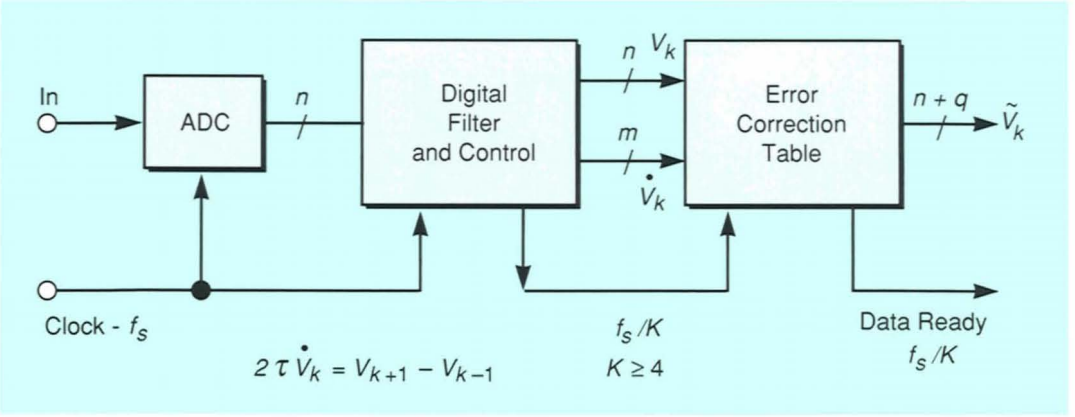

Fig. 9—Compensation structure 1 employs a single ADC and the central difference estimator (CDE). The ADC output code and slope estimate are used to address memory containing an error-correction lookup table.

The circumflex indicates that the quantity in question is an estimate of the derivative (as opposed to the true value), while T_s is the sample period. In general, the slope axis can be quantized to a different resolution than that of the ADC samples themselves. In Fig. 9, m bits represent the slope estimate while n bits represent the code value. The error table is stored in memory, and the code and slope values address the relevant memory location containing the error correction value. The correction value is then summed with the ADC code output to produce a corrected sample. In general, the corrected value will have a larger word length than the raw samples. Hence, in the figure the final output has a word length of $n + q$ bits. Figure 9 also shows that the output data rate, f_s/K, from the converter is a submultiple of the true ADC sample rate. The reduction in output sample rate accounts for the limitations of the CDE in estimating the signal derivative.

The major limitation in compensation structure 1 is in the use of the CDE to estimate the analog slope. The CDE can accurately represent only the slope of signals that are at frequencies less than one-fourth the sample rate. Usually signals are quantized with frequencies approaching the Nyquist limit of half the sample rate.

Another candidate for estimation of the time derivative of the converter input signal is the backward difference estimator (BDE), which estimates the slope according to the expression

$$\hat{\dot{x}}_k = \frac{x_k - x_{k-1}}{T_s}.$$

The BDE estimates the signal slope more accurately at frequencies above one-fourth the sample rate. In actual experiment the BDE does in fact outperform the CDE in compensating for ADC errors over a broader range of input frequencies.

Figure 10 shows a proposed ADC compensation structure 2. Two ADCs are used—ADC 1 is the primary converter (which is compensated) while ADC 2 operates at a much higher sample rate. In the figure, the ratio of the two sample rates is 10. By using a number of samples from ADC 2, a digital filter provides an estimate of the slope. The filter scheme illustrated in Fig. 10 is a finite impulse response (FIR) digital filter that uses samples located symmetrically in time relative to the sample to be compensated. The fact that the samples from ADC 2 occur at a much faster sample rate allows a more accurate slope estimate to be obtained for input frequencies approaching the Nyquist limit of ADC 1. If a causal estimate of the slope is desired, the FIR filter can be designed to employ only samples that occur simultaneous to or prior to the sample under compensation.

The primary limitation of compensation structure 2 is the requirement of the high-speed auxiliary ADC. For example, if ADC 1 operates at 25 Msps, then ADC 2 would operate at 250 Msps to maintain the 10:1 ratio of sample rates. While the ratio can be reduced, the high sample-rate

requirement for ADC 2 could limit the utility of structure 2. Since ADC 2 only estimates slope, it does not need to be as accurate as ADC 1. An inexpensive high-speed flash converter with less resolution than ADC 1 should perform adequately.

Figure 11 shows a proposed compensation structure 3. A continuous-time linear network generates the time derivative of the input signal. ADC 2 then samples the derivative signal. In structure 3 the auxiliary ADC does not need to operate at a higher sample rate than the primary ADC. Hence, ADC 2 does not have stringent speed or resolution requirements, and an inexpensive converter should suffice.

State Space Compensation

In the discussion above, causal (BDE) and noncausal (CDE) estimators of the slope of the ADC input signal are compared. The results imply that causal estimators are superior in the sense that the ADC being compensated is a causal system, and hence is more accurately modeled with causal estimators. The BDE was proposed as one such causal estimator. However, the BDE is a function of x_k and x_{k-1}, and $(x, \dot{\hat{x}})_k$ and x_k are related to x_k and x_{k-1} by a linear transformation. An error surface constructed with the independent variables x_k and x_{k-1} contains the same information as an error surface constructed with the BDE. The error surface with x_k and x_{k-1} as independent variables is called a state space representation (SSR). Because of identical error surfaces, the performance obtained with the SSR is identical to that obtained with the BDE. The advantage of the SSR is that no arithmetic computation of slope is necessary, since the independent variables are merely the present and most recent past ADC samples. Because of speed considerations, this advantage is important in applications in which real-time ADC compensation is performed. A disadvantage is that this estimator is also limited to calibration frequencies near one-fourth the sample frequency.

Error-Surface Construction

The ADC compensation described in this article relies on a valid error surface to perform the correction of the ADC output samples. To implement any of the ADC compensation schemes, the error surface must first be constructed. Efforts to date have not been directed at constructing the error surface at low cost or in a minimal amount of time; rather, emphasis has been placed on high integrity in the mea-

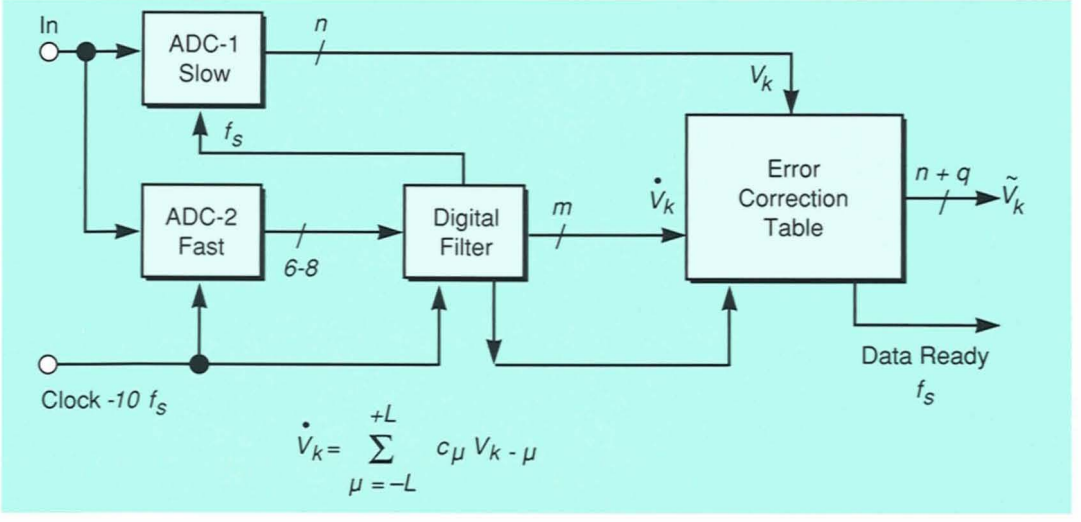

Fig. 10—Compensation structure 2 employs two ADCs and a symmetrical finite impulse response (FIR) derivative estimator. The second high-speed ADC yields an improved slope estimate over compensation structure 1.

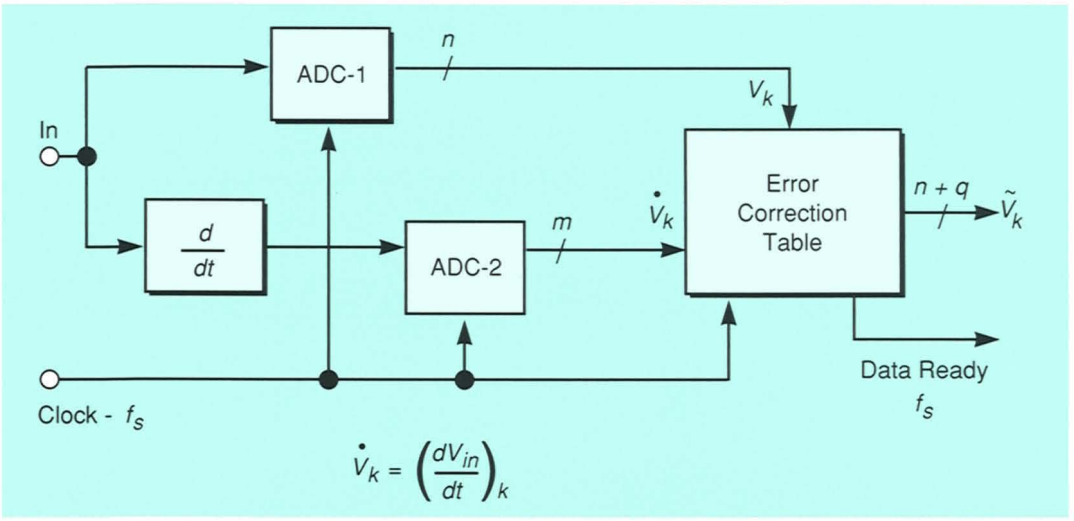

Fig. 11—Compensation structure 3 employs two ADCs and a sampled analog derivative estimator. This structure offers the potential for highest sampling rates.

surements and repeatability of the results. A brief description of the construction procedure follows.

The error surface is implemented by a lookup table in digital memory, and is defined on a rectangular region of the two-dimensional real plane (x, $\dot{x}$). The region is divided into a uniform grid of rectangular subregions called cells. The value of the error surface, determined by statistically averaging the error values obtained for each cell over a period of time, is assumed to be constant in each cell.

To obtain the error value in each cell, a known ADC input signal is necessary. If the signal is well defined, it can be separated from the error in the resulting ADC output data stream. Sinusoids were selected as the calibration waveforms, due to the relative ease of generating high-purity tones of variable amplitude. A sinusoidal waveform with a fixed frequency and amplitude sweeps out an elliptical path in the phase plane. In order to fill all the cells, a number of test tones are applied at varying amplitudes. The test-tone frequency corresponds to the desired maximum signal frequency, while the number of applied tones required is tied to the cell resolution selected for the table construction.

As each test tone is applied to the ADC, a record of data is stored in memory. The DFT estimates the signal component present in the stored data, then the estimated signal component is subtracted from the data to yield the error values. Next, the error values are assigned to cells addressed by the current and previous sample values for which the error value was obtained (in the case of BDE or SSR). Finally, error values accumulate in cells until the test is complete, at which time the accumulated values in each cell are averaged to derive a statistical estimate of the error value for each cell.

Compensation Examples

Three ADCs were selected to illustrate compensation results. These are the Analog Devices MOD-1205, the Burr-Brown ADC600, and the TRW TDC1025. The MOD-1205 lends itself well to compensation because it is stable over long periods of time. Error tables several weeks old consistently provide improved performance for the device. The ADC600 does not compensate effectively because the error tables appear random, and because high resolution and extensive averaging are required to obtain consistently valid corrections. The TDC1025 is a flash converter without a front-end sample-hold circuit. Near 11-bit performance is obtained for a signal bandwidth of approximately 4 MHz. This performance result dramatically illustrates the

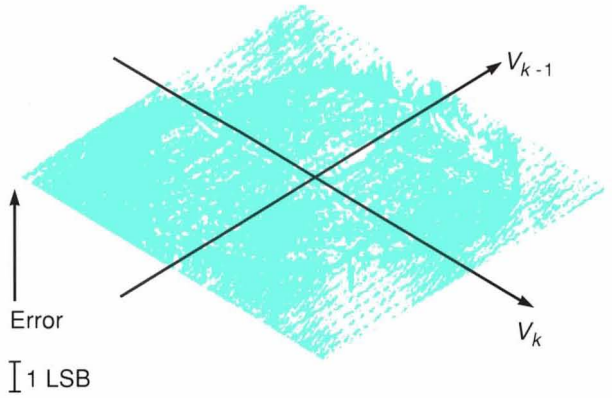

Fig. 12—Compensation table for Analog Devices Mod-1205 converter (SSR). The table is plotted as a function of the current ADC output state, V_k, and the previous output state, V_{k-1}. The cliff in the rear portion of the table is a prominent feature.

potential of the proposed compensation process. The following paragraphs summarize additional details for all three converters.

Analog Devices MOD-1205

The MOD-1205 is a 12-bit, 5-Msps converter. The SFDR achieved without compensation is approximately 65 dB. Figure 12 shows a representative compensation table generated with the SSR. The cliff in the rear portion is a prominent feature of the surface; this local functional dependence means that the error depends on the previous state as well as the current state.

Figure 13 shows a plot of SFDR versus input frequency for the uncompensated ADC, for the ADC compensated with the CDE, and for the ADC compensated with the SSR. The plot indicates that both the SSR and CDE yield dramatic improvement near the calibration frequency (approximately 0.5 MHz). The performance enhancement with the CDE nearly vanishes at 0.25 MHz; in contrast, the performance enhancement with the SSR is more broadband in nature. The gaps in the curves in Fig. 13 represent frequencies that could not be tested because the proper filters needed to purify the tone were unavailable.

Figure 14 shows plots of fundamental and highest spur level as functions of input power and frequency as described earlier with respect to Figs. 2 and 3. On the basis of these figures, an overall SFDR enhancement of 8 dB is obtained in the MOD-1205 over a wide range of operating frequencies.

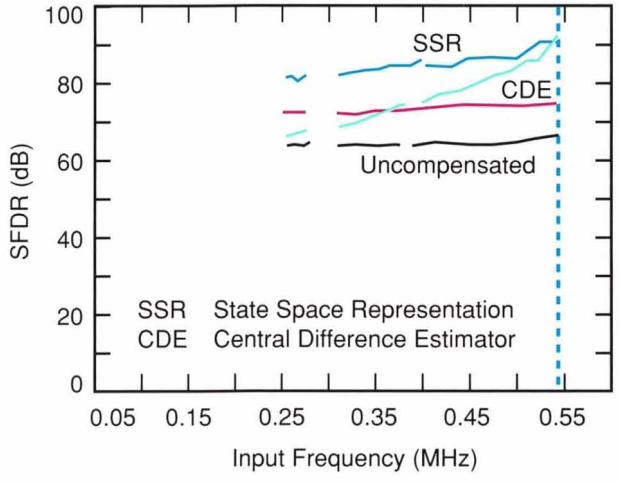

Fig. 13—Comparison of compensation techniques applied to the Analog Devices MOD-1205 converter. The CDE performs well only near the calibration frequency (blue dashed line), while the SSR performs well over a wide range of frequencies.

Burr-Brown ADC600

The Burr-Brown ADC600 is a 12-bit, 10-Msps converter. Without compensation, it exhibits SFDR that rivals results obtained from internally compensated 14-bit converters. Figure 15 shows the error surface obtained for the ADC600 with the SSR. Also shown is a plot of fundamental and maximum spurious power as a function of input power. The error compensation improves slightly at higher input power levels and not at all at low input power levels. Inspection of the error surface indicates that much of the ADC error is concentrated along a ridge near the origin of the current state axis. The shape of the ridge depends somewhat on the previous state. The table in this figure was constructed with 256 cells along the current state axis; additional resolution could model the behavior of the ridge more accurately. Since the converter can produce 4,096 distinct output states, 256 resolution cells in the current state

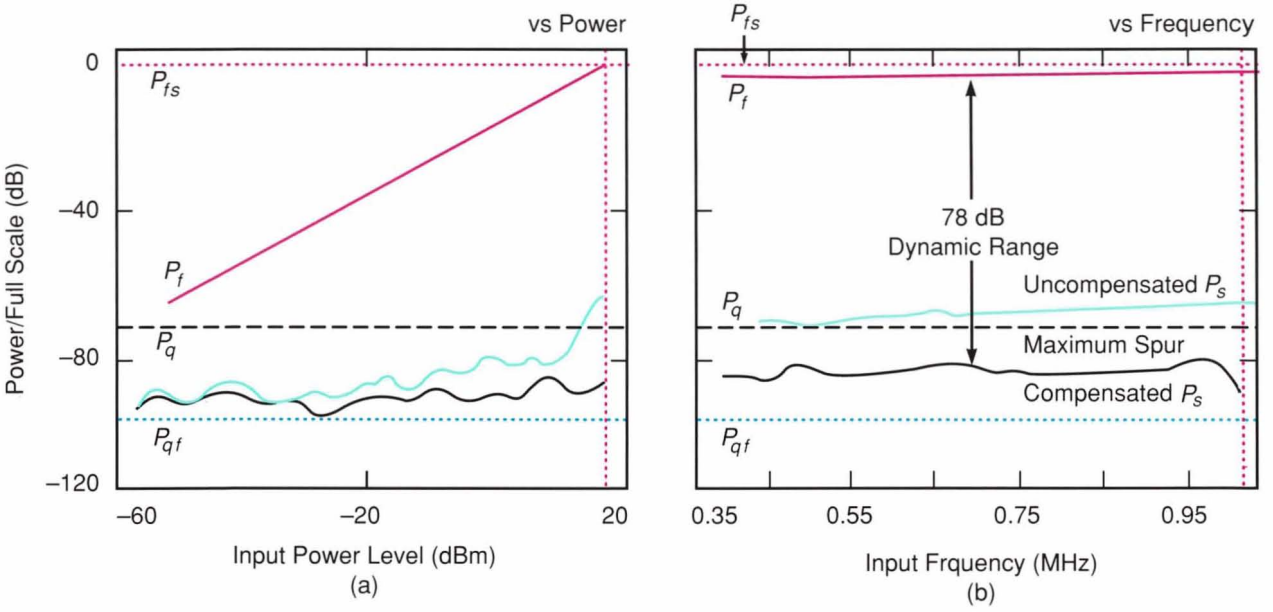

Fig. 14—Compensated SFDR performance of the Analog Devices MOD1205 converter. The quantity P_{qf} (blue dashed line) is the expected quantization noise power present in each DFT frequency cell. (a) The plot of fundamental power P_f and maximum spurious power P_s versus input power shows increasing SFDR enhancement with increasing input power. (b) The plot of fundamental power and maximum spurious power versus frequency shows at least 8-dB improvement over a wide frequency range.

axis is insufficient. At present, additional tests with improved table resolution have not been performed. This is seen as an area for future work.

TRW TDC1025

The TRW TDC1025 is an 8-bit, 50-Msps flash converter without a sample/hold circuit. Figure 16 shows the error table obtained with CDE compensation. The table size is 64×32 on the code and slope axes, respectively. The sample frequency is 25.6 MHz, while the tone frequency used to construct the table is 3.9875 MHz. The high degree of structure in the table (in contrast to that shown in Fig. 15 for the Burr-Brown ADC600) indicates that compensation will provide significant performance improvement. Figure 17 shows plots of the SFDR obtained at two widely different signal frequencies—3.5 MHz and 0.95 MHz, respectively. At both frequencies the usable SFDR is in excess of 60 dB, an astonishing result for an 8-bit converter, which represents an improvement over the uncompensated converter of more than 12 dB.

The promising results shown for the TDC1025 flash converter indicate that the use of flash converters in combination with dynamic error-compensation techniques can achieve high SFDR at high sample rates. Figure 18 shows an illustration of this fact from the work of Rebold [14]. The ADC input consists of a low-level FM signal near 8 MHz in the presence of a large narrowband interferer near 4 MHz. Figure 18(a) shows the uncompensated response. The desired signal is approximately 15 dB below the ADC quantizing power level P_q, and approximately 3 dB below the maximum spur P_s. Figure 18(b) shows the compensated response. All harmonic distortion is removed and the desired signal is clearly detectable above the filtered quantizing level P_{qf}. In addition, high SFDR does not fundamentally require high bit resolution in the ADC. Typically, however, ADC system designers specify resolution that is higher than necessary. The extra resolution is necessary to overcome the differences in real-

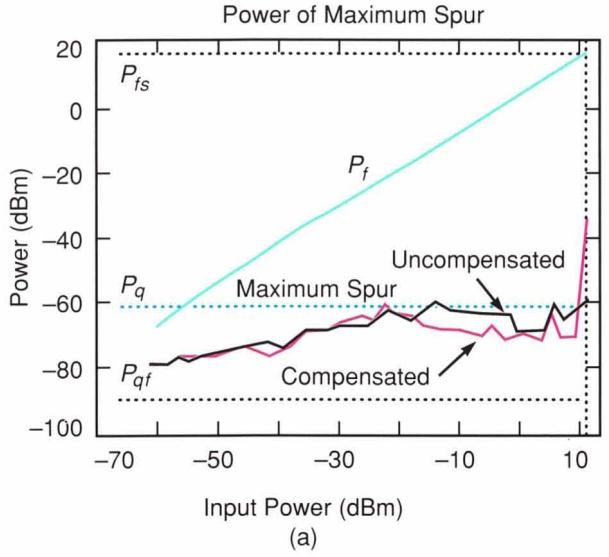

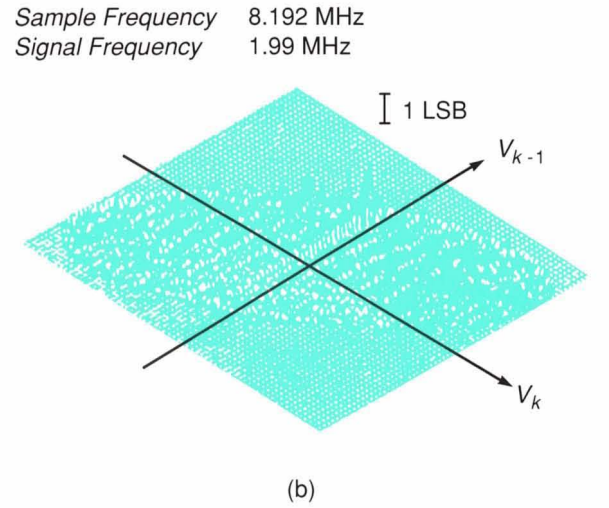

Fig. 15—SSR compensation applied to the Burr-Brown ADC600 converter. (a) The improvement due to compensation is modest, and occurs only over a limited input power range. (b).The compensation table has little structure except for a small ridge along the $V_k = 0$ *line.*

versus-ideal ADC performance.

Areas for Future Work

Dynamic error compensation is a promising technique that enhances the dynamic range performance of ADCs in spectral analysis applications. While preliminary results are promising, many topics remain to be investigated, such as the following:

(a) phase-plane validity for all ADCs,

(b) determination of cell size versus converter resolution,

(c) optimum derivative estimator, and

(d) performance limitations of structures 2 and 3.

This article proposes three compensation structures. Present work at Lincoln Laboratory employs compensation structure 1, while the effectiveness of the other structures requires further examination. For at least one converter, a very high-resolution table is required to obtain a worthwhile performance improvement. Higher-resolution tables and nonuniform cell sizes should be explored.

Another unexplored topic is the sensitivity of ADC nonlinearity to time and temperature changes. The dynamic correction techniques in this article are only practical if the ADC characteristics are stable and repeatable. ADC behavior will ultimately relate directly to the precise and stable representation of the error surface.

Summary

Dynamic error-compensation techniques

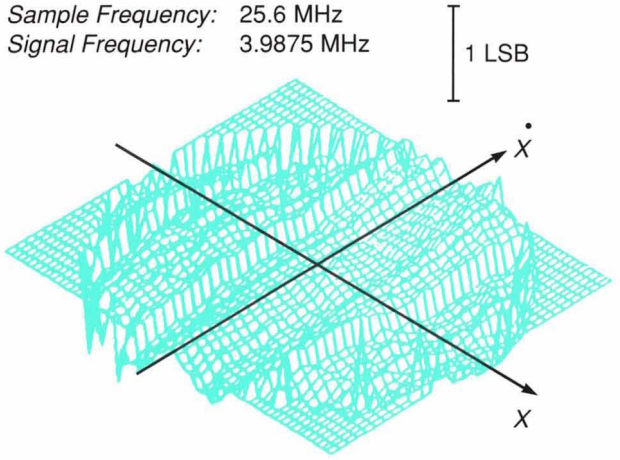

Fig. 16—CDE compensation table for the TRW TDC1025 converter. The table shows a high degree of structure, which indicates that compensation will be effective.

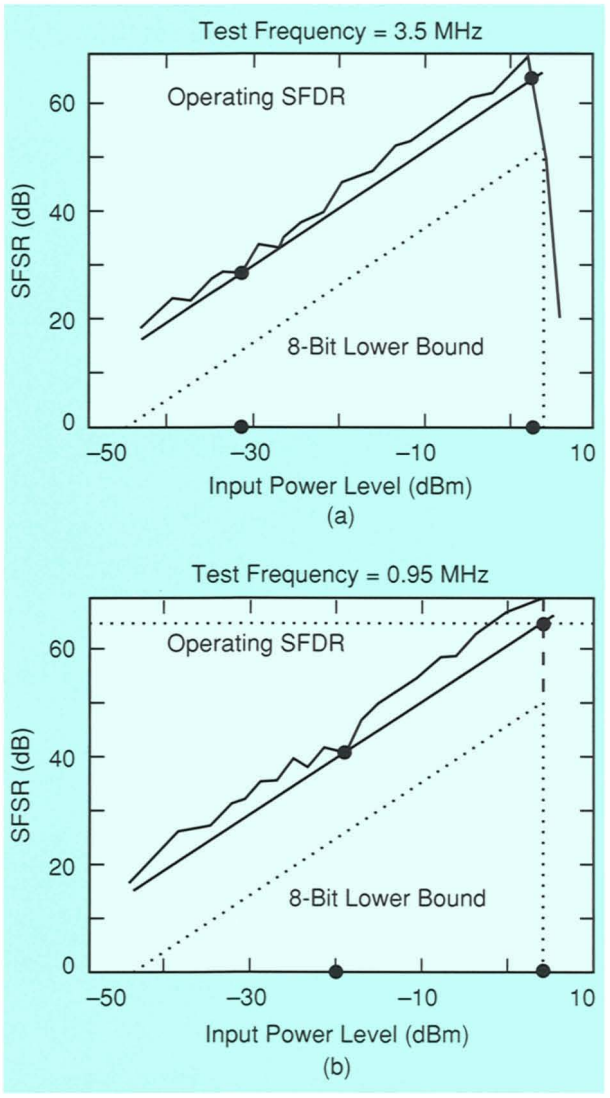

Fig. 17—Compensated SFDR for TRW TDC1025 converter using the compensation table in Figure 16. The two plots demonstrate that a 65-dB compensated SFDR can be obtained over a wide bandwidth. (a) Compensated SFDR at a test frequency of 3.5 MHz, near the calibration frequency of 3.9875 MHz. (b) Compensated SFDR at a test frequency of 0.95 MHz, much lower than the calibration frequency.

enhance the linear range performance of ADCs in spectral analysis applications. The important feature distinguishing these techniques from previous attempts at compensation is the use of an error surface defined on a two-dimensional real plane. The multidimensional error description allows dynamic error effects to be modeled subject to certain restrictions on ADC nonlinearity.

Three examples of error compensation were described. In two of the three examples, significant dynamic range enhancement was obtained over a wide range of operating frequencies, while in the third example only modest improvement was obtained. Inspection of the error surface indicates that increased table resolution may be required to obtain significant performance improvement.

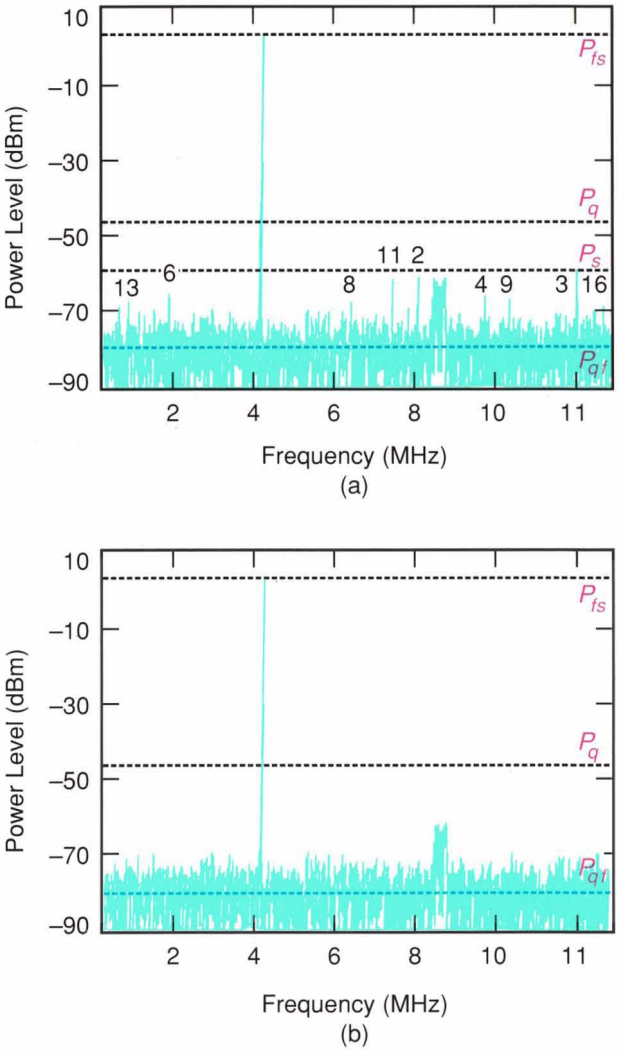

Fig. 18—Compensation example for the TRW TDC1025 converter, illustrating elimination of spurious components that mask a small signal. (a) The uncompensated response shows a near full-scale sinewave at 4 MHz that masks a low-level FM signal near 8 MHz. (b) The compensated response, which uses the correction table shown in Fig. 16, eliminates the spurious components and reveals the desired signal.

Acknowledgments

The success of any program depends upon many contributions, and this project is no exception. The authors would like to thank all contributors and to acknowledge the following special efforts.

Al Horst was responsible for the design and construction of the high-speed data buffer, and for supervision of the layout and assembly of the final test bed. He was ably assisted by Mike DiMare and Carl Reppucci, who deserve special thanks for their very careful attention to detail in providing a low-noise, high-performance test instrument.

Alan Chaiken and Tom Rebold developed several programs for this project, and contributed considerable input in simulation studies to help develop an understanding for both ideal and non-ideal ADC effects. Special thanks are due to Tom for his persistence and insight in the development of useful compensation procedures. Bill Rogers organized many of these programs to work together through a common parameter list and a single interactive menu, thus providing a professional performance for the final test bed.

The ideas and emphasis developed in this work are due in part to David Goldfein's and Jeff Friedhoffer's guidance. Their support and knowledge of all aspects of a communication receiver and its special requirements for spectral analysis were especially useful in shaping the goals for this study.

This project was supported by the Air Force.

References

1. N.S. Jayant and L.R. Rabiner, "The Application of Dither to the Quantization of Speech Signals," *Bell Syst. Tech. J.* **51**, 1293 (1972).
2. W.A. Kester, "Test Video A/D Converters under Dynamic Conditions," *EDN* **27**, 103 (1982).
3. W.A. Kester, "Characterizing and Testing A/D and D/A converters for Color Video Applications," *IEEE Trans. Circuits Syst.* **CAS-25**, 539 (1978).
4. "Dynamic Performance Testing of A-to-D Converters," *Product Note 5180A-2*, Hewlett-Packard Co., Santa Clara, CA (Dec. 1981).
5. J.R. Naylor, "Testing DA and AD Converters," *IEEE Trans. Circuits Systems CAS-***2,** 526 (1978).
6. R.W. Hamming, *Digital Filters* (Prentice-Hall Inc., Englewood Cliffs, NJ, 1977).
7. F.J. Harris, "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform," *Proc. IEEE* **66**, 51 (1978).
8. F.H. Irons, "Dynamic Characterization and Compensation of Analog-to-Digital Converters," *IEEE Int. Symp. on Circuits and Systems, San Jose, 5–7 May 1986* **IEEE Cat 86Ch2255-8,** 1273.
9. J.T. Millman, "Error Minimization of Data Converting Systems via Generalized Spectrum Analysis," *Ph.D. Thesis* (MIT, Cambridge, MA, Sept. 1975).
10. E.A. Slome, "A System of Converter Testing Using Walsh Transform Techniques," *Intl. Test Conf., Philadelphia, 27–29 Oct. 1981* **IEEE Cat. 81CH1693-1,** 304.
11. M. Vanden Bossche, J. Schoukens, and J. Renneboog, "Dynamic Testing and Diagnostics of A/D Converters," *IEEE Trans. on Circuits and Systems* **CAS-33**, 775 (Aug. 1986).
12. T.A. Rebold and F.H. Irons, "A Phase-Plane Approach to the Compensation of High-Speed Analog-to-Digital Converters," *Proc. of IEEE Intl. Symp. on Circuits and Systems, Philadelphia, 4–7 May 1987* **IEEE Cat. 87CH2394-5,** 455.
13. N.W. Spencer, private communication.
14. T.A. Rebold, "Dynamic Error Correction Method for High-Speed Analog-to-Digital Converters," *S.M. Thesis*, (MIT, Cambridge, MA, 1987).

DANIEL ASTA is a staff member in the Advanced Techniques Group. His research speciality is in the field of adaptive array processing. Dan received his B.S. and M.S. degrees in electrical engineering from the University of Illinois, and his Ph.D. in electrical engineering from UCLA. At Illinois, he received the Bronze Tablet, the university's highest award for academic achievement. Before coming to Lincoln Laboratory in 1987, he worked at Hughes Aircraft Company in El Segundo, Calif.

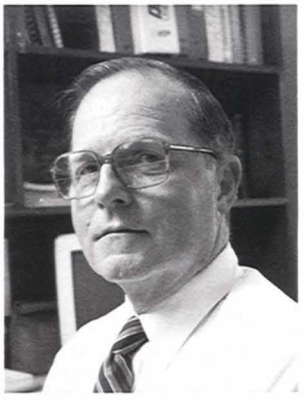

FRED H. IRONS is a staff member in the Advanced Techniques Group. He has been at Lincoln Laboratory for 12 years, and his interests are in the area of signal processing and system analysis. He received his B.S. from Ohio State University, his S.M. from MIT, and his Ph.D. from Lehigh University, all in electrical engineering. Before coming to Lincoln Laboratory, Fred was an associate professor of electrical engineering at the University of Maine at Orono.