

An 18-Bit Digital-to-Analog Converter for High-Performance Digital Audio Applications*

JOEL M. HALBERT AND MARK A. SHILL

Burr-Brown Corporation, Tucson, AZ 85706, USA

Design and implementation details of a new 18-bit digital-to-analog converter are presented. The device is suitable for digital audio applications requiring an unusually large dynamic range, low wide-band noise, high speed, and low distortion. Key design considerations include careful management of thermal effects, a low-glitch switching architecture, and isolation of parasitic noise sources.

1 THE CASE FOR CONVERTERS BEYOND 16 BITS

1.1 Professional Applications

Since the introduction of consumer digital audio equipment at the beginning of the decade, professional audio engineers have found themselves in a new and difficult position. Through the 1970s professional recording and mixing electronics could be designed to have measurably better performance than most consumer playback machines. Even with several generations of mixing, dubbing, equalization, and other processing, the final master tapes were usually cleaner than all but the most expensive home reproduction systems. Furthermore, the sound from top-quality playback systems was still limited by the fidelity of mass-produced LPs and cassettes.

This comfortable situation was seriously challenged by the introduction of Compact Disc (CD) technology. First, CDs largely eliminate the quality variables associated with vinyl and analog tape media. Second, the accuracy of the analog output signal is limited only by the digital-to-analog converter (DAC) deglitcher, filter, and small-signal postamplifier. Although not perfect, the 16-bit DACs available to CD player manufacturers have extremely good performance—in many cases better than the total signal chain used in the recording and mastering process. Designers of professional studio equipment can buy converters that are specially selected and guaranteed for linearity, but these

provide only partial improvement because the fundamental limitation of the 16-bit quantization error remains.

In order to realize the full potential of audio fidelity for the end-user, the signal digitization and processing steps must have a greater dynamic range than the final recording. The studio engineer is then free to equalize, expand, filter, mix, and amplify without accumulating errors higher than the 16-bit level. An ideal signal processing system would therefore be based on digital words of more than 16 bits.

Modern studios use both analog and digital processing equipment. Connecting them together requires digital-to-analog and analog-to-digital conversion elements which should also provide greater than 16-bit resolution. The basic building block of these elements is a reliable, easy-to-use, and compact DAC.

1.2 18 Bits in Consumer Equipment

High-quality consumer playback equipment is also moving toward the use of 18-bit DACs, even though the source material (CD and DAT) contains only 16-bit data words. At this time the authors feel that the relative merits of such applications have not been well established—but this does not imply that there is no justification for such use. Certainly, it is advantageous to improve the linearity and reduce the thermal noise of the playback DACs, improving distortion and signal-to-noise ratio (SNR). Beyond that, the addition of 2 bits has several potential advantages, which require further investigation.

It is certainly possible to recover real information below the 16-bit quantization level during periods when

* Presented at the 83rd Convention of the Audio Engineering Society, New York, 1987 October 16–19.

the signal bandwidth is considerably less than the Nyquist frequency (20–22 kHz for CDs). This is because a limited-bandwidth signal is effectively being oversampled by the 44.1-kHz clock, and a properly designed digital signal processor can use extra DAC resolution to best advantage. (Note that this is not the same thing as the “oversampling” digital filters currently used in many CD players, which cannot generally extract any real information beyond 16 bits. Though the interpolation hardware used in these filters may have extra roundoff bits, the significance of the resulting signal is actually less than 16 bits, leaving the total residual error roughly constant.)

For the widely used total harmonic distortion test frequencies below 5 kHz, a reasonably linear 18-bit DAC can produce total harmonic distortion figures in excess of -100 dB. Low-level low-frequency gain tracking will also benefit from the extra dynamic range. Again, results which at first glance seem to violate the 16-bit “limit” are indeed feasible, assuming specific input signal characteristics. To take full advantage of 18-bit converters in a 16-bit CD or DAT player, more sophisticated signal-adaptive DSP hardware will be required. Although particular specifications can be improved in this way, the subjective benefits with real musical signals are yet to be documented. It is also quite possible that the extra bits can be used to produce effects which subjectively improve the sound quality, even if the measured total error power is not reduced. The increasing market presence of players with “18-bit technology” will stimulate manufacturers to maximize both measured and perceived performance.

The 18-bit DAC described here is designed to provide the “building block” solution for demanding applications in studio equipment. The converter can also be applied to consumer equipment, providing new levels of performance along with flexibility for manual or automatic calibration.

2 OPTIONS FOR HIGH-RESOLUTION DAC DESIGN

In order to understand the design decisions that were made in the development of the 18-bit DAC, it is useful to review various basic approaches to DAC design.

Although there are many possible architectures for high-resolution DACs, most general-purpose designs are based on the summing of weighted current or voltage sources. Certain performance criteria [for example, differential linearity (DL)] can be optimized using different approaches, but these have usually been reserved for applications with special requirements, which are beyond the scope of this discussion.

2.1 Binary-Weighted Designs

The binary weighting for N bits is often achieved using N equal-valued sources summed into an $R-2R$ “ladder” resistor network (Fig. 1). This approach is very economical, but it requires each individual source to be precisely trimmed. The required precision of the

most significant bit (MSB), bit 1, source is equal to the resolution of the DAC itself; the next MSB, bit 2, must be adjusted with at least half of the overall DAC precision, and so on. In general, the source (or ladder element) corresponding to bit k must be adjusted to a relative accuracy of 1 part in 2^{N-k+1} . This is not difficult for the lower (less significant) bits, but can become a problem for the upper bits of a high-resolution device.

Alternatively, $2^N - 1$ equal-valued sources can be used, where one source is assigned to the least significant bit (LSB), two sources are assigned to the next bit, four to the next, and so on in binary progression (Fig. 2). The total weight of each bit k is then determined by the summation of 2^{N-k} sources. Thus any individual source is allowed to deviate randomly from its ideal value as long as the statistical average of many such random deviations is close to zero.

This “multiple-source” approach has gained favor in monolithic DACs because it exploits the advantages of monolithic devices (good statistical matching and relatively small parametric variations across the chip), while minimizing the cost associated with the large transistor count. Even so, it does become impractical to have $2^N - 1$ total sources for high-resolution converters. Therefore the architecture will usually be split between an upper M -bit DAC, which employs $2^M - 1$ sources for highest accuracy, and a lower $(N - M)$ -bit DAC, which uses an $R-2R$ ladder network (using $N - M$ sources) for the less significant bits.

2.2 Superposition Errors

One problem that can occur with binary-weighted DACs is called superposition error, in which the various sources interact to produce nonlinearity. If the weighting of the bit 2 source is slightly different depending on whether bit 1 is on or off, then it will be impossible to achieve perfect converter linearity over the entire code range. These errors have a number of causes, including parasitic resistance in ground lines, thermal interaction of various devices on the chip, and self-heating of critical circuit elements. For this reason, designers have looked for methods that will reduce the transfer function’s sensitivity to layout imperfections.

2.3 Segmentation versus Straight Binary Weighting

An extension of the multiple-source approach, which has been used in monolithic DACs for several years, is known as segmentation. This is similar to the previously described method in that many equal-valued sources are summed to achieve the total output value. The difference is that each source is associated with a particular input code, rather than being tied directly to a particular bit. As the input code increases monotonically, sources are switched on and add 1 LSB to the total output signal (Fig. 3). Unlike a binary-weighted DAC, a completely segmented design will have guaranteed monotonicity (that is, the output signal will never “reverse,” but will always increase as the input code increases). For the reasons described, however,

segmentation is usually applied only to M critical upper bits. Still, segmentation provides another advantage—the M segments of the DAC transfer function can be adjusted individually to compensate roughly for superposition errors, gain compression (S -curve nonlin-

earity), and other integral linearity problems.

Despite the attractiveness of segmenting the upper bits, there are still a few key advantages to straight binary weighting. No decoding logic is required, which helps maintain equal propagation delays from the TTL

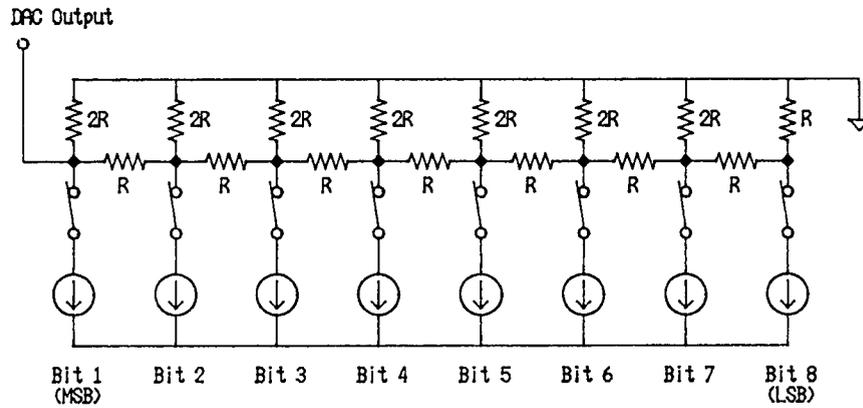


Fig. 1. 8-bit DAC using conventional $R-2R$ ladder network.

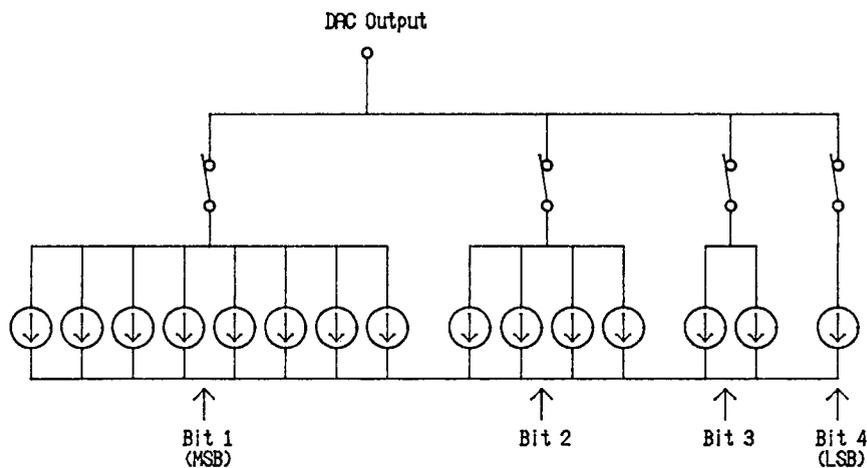


Fig. 2. 4-bit DAC using binary weighting with $2^N - 1$ sources.

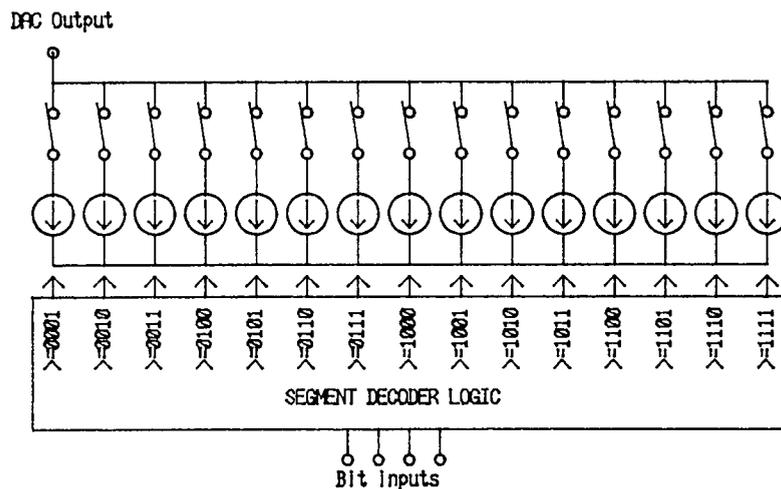


Fig. 3. 4-bit segmented DAC using 15 equal-valued sources.

bit inputs to the actual bit switch. This is important for reducing the nonlinear glitch energy at the DAC output. Another major advantage is the ability to adjust the DAC linearity using a strictly differential approach. Assuming that superposition errors have been held to negligible values, the linearity trim can be accomplished by simply adjusting each code "carry" transition for a 1-LSB change. In other words, the DAC is calibrated to itself. No external reference DAC (or ADC) is required, except for purposes of final verification. This feature makes it much easier to trim very high resolution converters (16 bits and up) during the manufacturing phase—particularly at the wafer probe operation, where it is difficult to maintain microvolt accuracy between the device and the test equipment. Furthermore, differential trim provides a very simple and reliable calibration procedure for the end-user. Calibration is covered in more detail following a discussion of the 18-bit DAC circuit.

2.4 Dynamic Element Matching

Yet another twist, which is familiar to audio engineers, is called dynamic element matching. This is a feature that may be applied to either binary-weighted or segmented designs to achieve close and stable matching of the current sources, while reducing the required precision of the trims. In dynamic element matching, the weight of each bit or segment is determined by the average of many primary current sources. The switching algorithm is arranged so that over a long period of time, each bit or segment receives the correct proportional contribution from the primary sources. (Extrapolating to the simplest case, one primary source can be switched to create all the bit sources by time-averaged weighting.) The dynamic approach does have appeal from a conceptual standpoint, although practical implementations do not eliminate all of the possible bit mismatches. To achieve the lowest possible errors, some provision for factory, user, or automatically controlled adjustment would still be required to "close the loop." For extremely accurate designs, then, the disadvantages of the dynamic switching, filtering, and logic circuits may not be acceptable if a secondary correction method is required anyway.

2.5 Other Architectures

Finally it must also be noted that weighted summing of current or voltage sources is not the only design approach that has been used for audio converters. Weighted capacitor DACs, dual and triple slope integrating designs, and various types of "oversampling" DACs are among the possible methods. While all of these methods can be made to work in a theoretical sense, it is the opinion of the authors (supported by the results in the marketplace) that these other approaches are currently best suited for low-to-medium-performance audio products. Capacitor-based designs suffer from the lack of an acceptable trim method, which tends to limit the ultimate linearity to the 10- or 12-bit level (although automatic correction circuitry can

be added to improve that situation greatly). Integrating converters have been used successfully in CD players. Unfortunately, every 1-bit increase in resolution puts an enormous strain on the design; it involves a doubling of the internal clock rate or additional subbranching circuitry, which complicates the architecture.

Recent advances in oversampled or "noise-shaping" converters have generated a great deal of interest. The simplified models that are used to present the case for these converters require further improvement. At present the limits of real-world converter performance versus oversampling ratio are not clearly defined, and better objective tests should be developed to evaluate results. Nonetheless, autocalibration and oversampling techniques hold great promise, and further refinement is certain.

3 SPECIFIC ARCHITECTURE USED FOR THE 18-BIT DESIGN

3.1 General Description

Following the logic of the preceding discussion, it was decided that a fairly conventional binary-weighted architecture would serve as the best approach for the new design. The limitations and potential problems of this topology are well understood, based on extensive experience with monolithic 16-bit converters. Therefore achieving a meaningful performance improvement depends on the elimination of known superposition, noise, and dynamic errors. The monolithic process used is also well characterized, possessing fast, high-accuracy NPN transistors with good matching characteristics. Nichrome laser-trimmable thin-film resistors have a low temperature coefficient (TCR) and, more important, sub-ppm matching of TCR across the chip.

Fig. 4 shows a simplified diagram of the 18-bit DAC. The 3 MSBs are made up of seven individual current sources, which are used in a 4-2-1 arrangement described. This not only relaxes the matching requirements on each source, but it also permits them to be physically interdigitated on the chip to reduce code-dependent thermal errors. Bits 4-16 are made up of unit-valued current sources which feed the $R-2R$ ladder network. The currents for bits 17 and 18, the least critical, are split off from a single unit-valued source. The relative gain of the 3-bit "upper" DAC can be adjusted against the total weight of the 15-bit $R-2R$ "lower" DAC by trimming the scale-down network.

3.2 Bit Switch Cell—Switching and Noise Characteristics

The individual current source/bit switch cells (Fig. 5) have several features which improve the noise and settling characteristics of the converter. The current source transistor Q_6 is isolated from the switch itself by Q_5 , the cascode transistor. Q_5 absorbs the switching waveform on the differential bit switch pair Q_1-Q_2 and prevents it from coupling to V_{ref} through the collector-base capacitance of Q_6 . Since the reference generator circuit (discussed below) is a low-bandwidth

servo loop, keeping the reference line as clean as possible is important to minimize DAC settling time.

The bit switch itself is driven differentially by the outer switch pair Q_3-Q_4 , which further reduces the glitch energy at the emitters of Q_1-Q_2 . However, the primary reason for including the Q_3-Q_4 circuit is to decouple the wide-band noise from the zener diode level translator. In a standard monolithic process, high-speed level translation of several volts is difficult to

achieve without using zener diodes. (The zener is actually an avalanche diode which uses the reverse breakdown of an NPN base-emitter junction.) Unfortunately the avalanche mechanism is inherently quite noisy. Previous bit switch designs drove the base of Q_2 directly from the level shifter, which coupled noise through the base-emitter capacitance of Q_2 and into the emitter of the output device Q_2 . This noise was a limiting factor in the SNR performance, but has es-

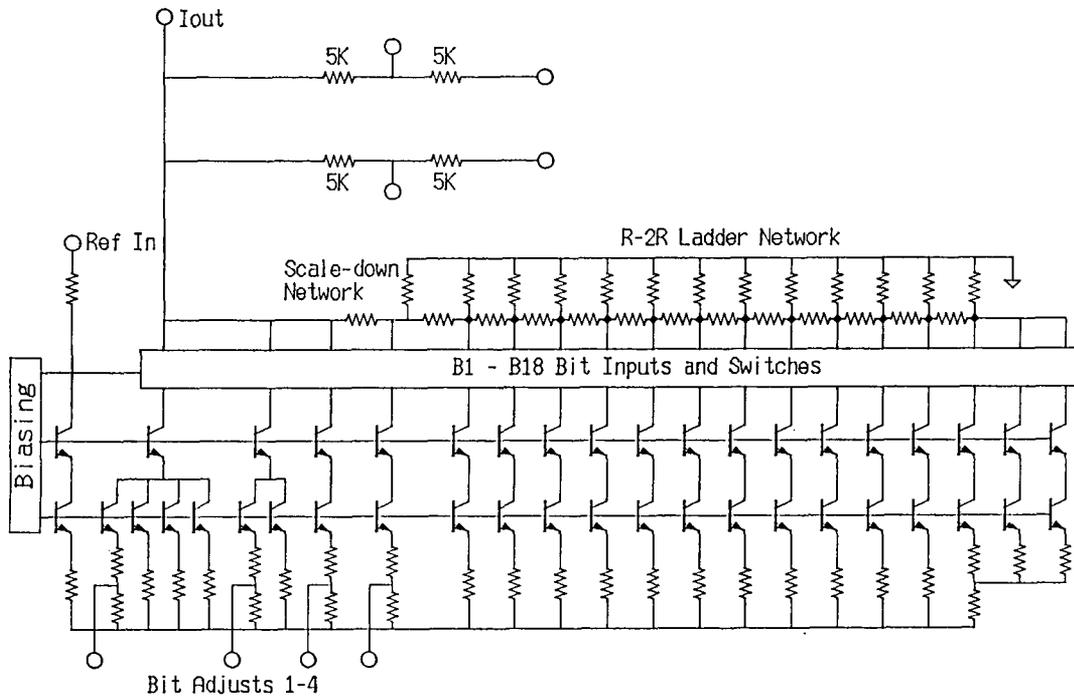


Fig. 4. Monolithic 18-bit DAC.

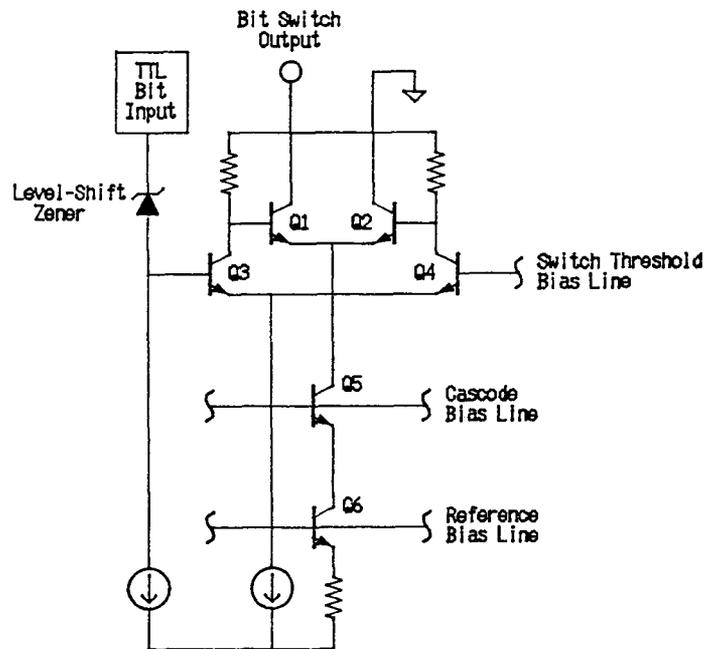


Fig. 5. Bit switch cell.

entially been eliminated from the present design.

3.3 Reference and Total DAC Noise

Another noise source is in the reference generator itself. Accurate and stable references are usually created using specially designed and compensated zener diodes, which are inherently noisy, as noted. Band-gap reference generators are usually even noisier than zener references, because a high-gain amplifier is necessary to boost the band-gap output voltage to a usable reference level. In the 18-bit design, the reference noise is decoupled using an off-chip 0.01- μ F capacitor.

The resulting noise from the DAC has been measured at various output codes, and at high frequencies it is clearly dominated by the thermal noise of the 4-k Ω output impedance (the impedance of the $R-2R$ ladder and the upper/lower DAC "scale-down" attenuator). This is a fundamental limitation, and it is not possible to reduce the spectral noise density of the DAC any further without increasing the current levels. In any case, the random noise is below the quantization limit, even for highly oversampled (wide-band) audio systems. Usually the system noise will be dominated by the output amplifier and the following circuitry.

3.4 Thermal Considerations and Off-Chip Components

The 18-bit DAC was designed to maximize performance and to avoid the tradeoffs that were encountered previously in single-chip implementations. This implies a willingness to partition the design, taking peripheral elements off the chip if necessary.

As mentioned, the straight binary-weighted approach has important advantages, assuming that superposition error is kept under control. One LSB value at 18 bits is approximately equal to 4 parts per million (ppm). In order for the binary approach to work well, all code-dependent shifts in the bit weights must be held to no more than 1 or 2 ppm. This requires great care in the routing of signal and ground traces, good isolation of analog and digital signals, and, most important, careful thermal design.

As the differential pairs change state in each bit cell, the local heat sources move from one device to another, and may also increase or decrease in magnitude. These local temperature changes must be kept far away from the current source transistors—even a small fraction of 1°C at the current source will cause unacceptable shifts in the bit current. Besides physical distance, one way to minimize the temperature change is to balance the layout by placing the differential pair transistors symmetrically with respect to the current source device. Then, even when the location of nearby heat producers shifts, the average distance does not shift, and constant heating on the current source transistor is maintained.

Another major source of differential power is found in the feedback resistor network. The feedback network is left on chip for maximum gain stability and TCR tracking, but this raises the problem of unequal and code-dependent heat flow across the chip. The power

in the feedback network varies from 0 to 20 mW (10 V/k Ω) in the worst case. To reduce and balance the effects of this power change, the feedback network is placed along the top edge of the chip, quite far away (relative to monolithic integrated-circuit dimensions) from the critical bit switch cells. Also, the network is centered over the MSB cells (which themselves are symmetrically interdigitated), further reducing any differential effects among the upper few bits. Self-heating of the feedback network is minimized by making the resistors large, thus maintaining a low power per unit area. The chip layout in Fig. 6 illustrates the points discussed.

Existing 16-bit digital audio DACs fabricated on this process are designed to be completely monolithic devices. They include both a complete reference circuit and a fast-settling output operational amplifier (op amp) on the chip. These provide very economical, convenient, and high-performance solutions and are widely used in both consumer and professional equipment of all types. However, the completeness of these chip designs also leads to subtle imperfections, which limit their ultimate linearity and noise performance.

First the output op amp (current-to-voltage converter) is designed to supply at least 5-mA source or sink current to an external load. Under loaded conditions, the output transistors will dissipate between 0 and 40 mW, varying nonlinearly with the signal level. This is a larger effect than the feedback resistor problem discussed and greatly compounds the problem of differential heating on the current source transistors. In the 16-bit devices (0.0015–0.003% total linearity error) the problem is held to an acceptable level by employing a partially segmented architecture. However, the design goals of the 18-bit DAC require much lower superposition error. The new device provides the capability for linearity adjustment up to true 18-bit performance (0.0002%), which means that the variable power from the op amp becomes an almost intractable problem. The solution in this case is to remove the op-amp circuit from the chip and provide the function in the DACs hybrid integrated-circuit package. For versions of the product which package the chip separately (and as an option in the hybrid version), the end-user may choose the most suitable amplifier for the application.

Removing the amplifier from the DAC chip does not eliminate all of the problems associated with the output op amp circuit. Any monolithic op amp is itself subject to thermal feedback errors, due to input stage heating from the output stage. Driving a real-world load causes small variations in the input offset voltage, an effect that limits the linearity of precision op amps. Several candidates for the op amp function were evaluated. Unfortunately no amplifier was found that had both 18-bit linearity and adequate speed (settling time) to meet design goals.

The solution, again, is to partition the function between two chips. In the hybrid (complete) version of the 18-bit DAC product, a high-speed FET-input amplifier chip is combined with a monolithic buffer am-

plifier to create a high-speed and highly linear DAC output amplifier. The buffer amplifier itself is an extremely wide-band device, and can be placed inside the feedback loop of the fast-settling FET amplifier without significantly affecting total phase margin. The combination works because the FET amplifier drives only the high-impedance buffer input, and therefore it suffers no linearity degradation from thermal feedback.

Another feature of the existing 16-bit designs is an on-chip reference circuit. Critical elements of the reference circuit must remain on the chip for good matching with the DAC itself, but large portions of the reference circuit (reference op amp, reference zener, and mis-

cellaneous bias circuitry) have no important matching requirements. These noncritical reference components dissipate well over 100 mW, which create a troublesome hot spot if left on the chip. Although the power dissipated by the reference does not change versus the input code, it is still a potential source of error. The reference op amp may experience small offset shifts due to changing power elsewhere on the DAC chip, causing the gain to shift nonlinearly with code changes. Also, heat from the reference causes a temperature gradient across various parts of the DAC circuit. Although transistor mismatches due to the gradient can be trimmed out at any given ambient temperature, sec-

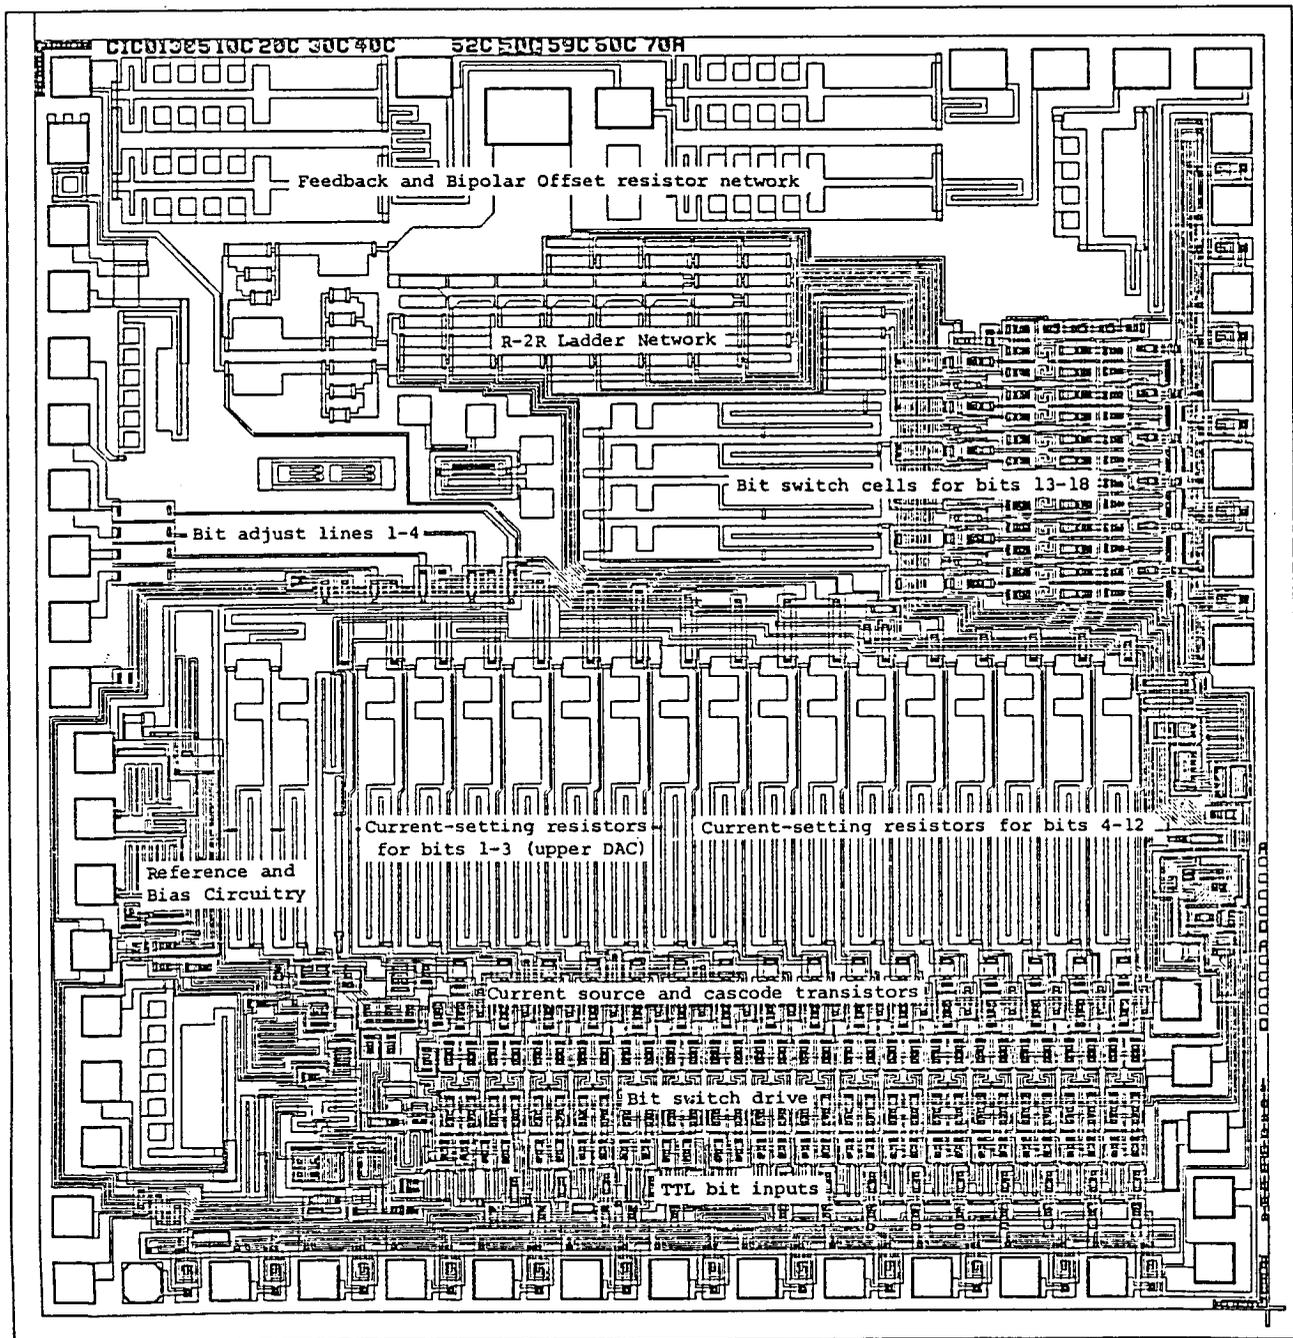


Fig. 6. Monolithic 18-bit DAC. Die size is 3.9 by 4.0 mm (0.154 by 0.159).

ond-order terms will cause the mismatches to reappear at other temperatures. Finally, some applications for the 18-bit DAC require a gain reference that is extremely stable over temperature and time. The best results and the greatest versatility are achieved by building the reference as a separate circuit—again, included in the hybrid package.

Fig. 7 shows the reference control loop (again in simplified form), which is largely a conventional design.

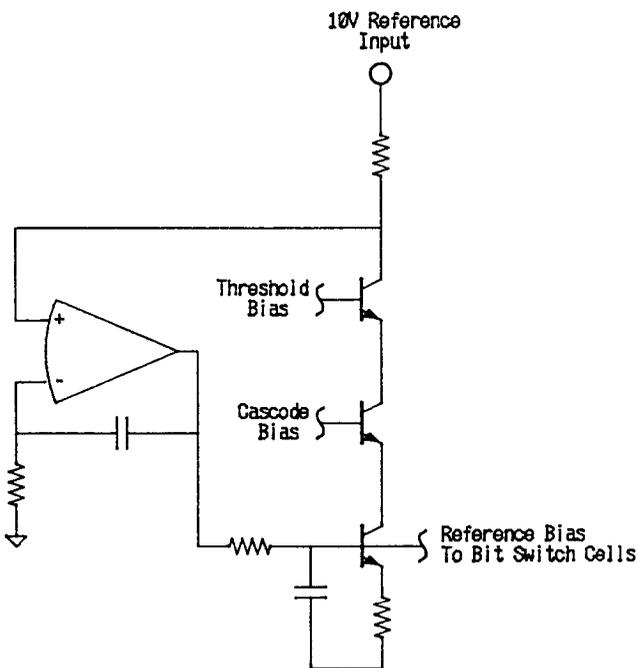


Fig. 7. Reference servo control loop.

A large (0.01- μ F) capacitor is used for noise decoupling, and also for complete suppression of any switching transients that manage to get through the cascode isolation. This creates the need for a similarly large loop-stabilizing capacitor around the servo op amp. The 10-V reference generator is temperature compensated in manufacturing and achieves typical drift performance of 1–3 ppm/ $^{\circ}$ C.

The complete 18-bit DAC circuit is shown in simplified form in Fig. 8. The hybrid device contains six semiconductor chips (DAC, output op amp, buffer, reference servo amplifier, reference output amplifier, and reference zener diode) and several passive devices (capacitors for noise reduction and bypassing, thin-film reference resistor network). The package itself is a 40-pin side-brazed integrated circuit, 1.2 in² (7.7 cm²) in area, fabricated with an integral multilayer substrate. Power dissipation is approximately 1 W.

4 PERFORMANCE CONSIDERATIONS AND SUMMARY OF MEASURED RESULTS

4.1 Linearity, with and without User Adjustment

Linearity performance for the 18-bit design has several aspects. First, the DAC should be as linear as possible when shipped to the customer. Second, although the linearity is expected to shift to some degree over the operating life of the product, it is always advantageous to minimize such shifts and to extend the period between calibrations for very demanding applications. Third, and ultimately most critical, is the assurance that the converter can in fact be calibrated to deliver the required linearity over the entire transfer function.

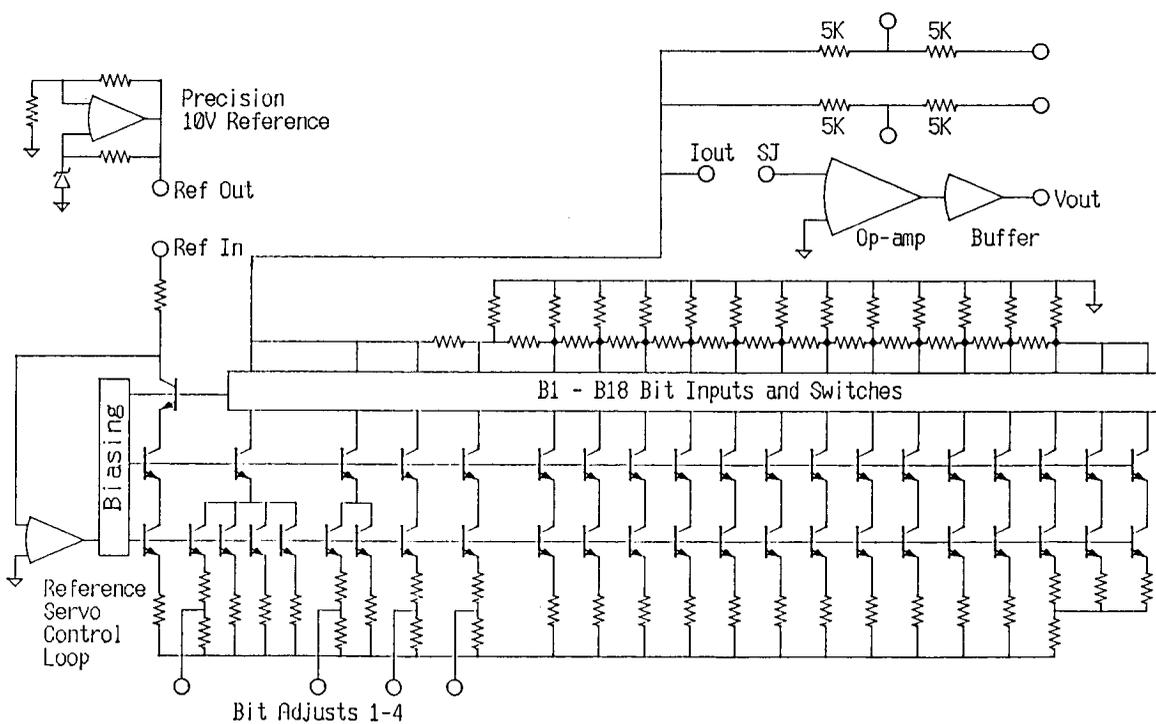


Fig. 8. Complete hybrid DAC.

These three considerations will be emphasized differently depending on the application. Stability over time and temperature is of course limited by the characteristics of the monolithic process—the transistor junctions and the thin-film resistors will shift slightly. Given these limitations, however, the most important factor governing long-term stability is the quality of the environment. Hermetic sealing in an inert atmosphere is the best way to minimize shifts, though it will not completely eliminate them. Hermeticity is achieved by seam welding the cover of the hybrid version. On the other hand, some applications demand a lower cost than the complete hybrid DAC can provide. For those instances, customers will purchase the monolithic version in an injection-molded plastic package and add external components to achieve the output amplifier and reference functions themselves, thereby minimizing cost and tailoring the performance to their needs. The device's stability in a plastic package is somewhat inferior to that of the hermetic version, and the packaging process itself can induce shifts from various sources. Therefore an initial in-circuit calibration is expected to be helpful for the hybrid version and required for the monolithic version.

Actual results from manufacturing lots of the hybrid DAC show excellent yield to 0.00075% unadjusted linearity (1 LSB of 17 bits) at the time of shipment. After 500 h life testing at 85°C ambient, which is roughly equivalent to 32 000 hours or 3.7 years at 25°C ambient (room temperature), some of these devices will exhibit shifts that bring them to the range of 0.0015% (1 LSB at 16 bits) unadjusted linearity. Linearity drift over temperature is well below 0.5 ppm/°C and does not increase significantly over life.

These results, though not "perfect," represent the highest unadjusted linearity performance yet achieved in high-speed integrated-circuit DACs. Furthermore, the key advantage of the new design lies in its capability for true 18-bit performance after a simple and straightforward calibration procedure. In the typical environment of the studio, the DAC (once factory adjusted by the equipment manufacturer) will stay within the limits of 18-bit performance for several months. Because the DAC is calibrated by measuring against itself, no sophisticated equipment is required, and the recalibration procedure is simple enough to be performed reliably by on-site personnel.

If true 18-bit performance is required and recalibration in the field is not a viable option, then a straight-

forward autocalibration circuit can be built which replaces the adjustment potentiometers with low-cost adjustment DACs operating under software control.

4.2 Superposition Error Results

As described earlier, low superposition error is the key requirement for a (binary-weighted) DAC to achieve true 18-bit performance. Measured results on the new DAC confirm a superposition error of typically 2 ppm or 20 μV on a 10-V full-scale range (FSR). These results are derived as follows. The largest sources of superposition error are found in the upper bits, or MSBs, of the DAC. This is because the upper bits have the most weight and any errors in the current sources of the lower bits will be scaled down binarily. Any superposition error introduced by self-heating of the DAC's feedback resistor will also be restricted to the upper MSB codes since it is only for these codes that the feedback resistor will have significant power fluctuation. With these points in mind, superposition testing was done for the first 4 MSBs. To determine superposition error, all possible combinational codes of the first 4 bits were tested. If the DAC has zero superposition error, the total linearity error at any given code (regardless of its deviation from ideal) should equal the algebraic sum of the individually measured errors for each individual bit. For example, the error for the code 11 0100 0000 0000 0000 (the "electronic sum" with bits 1, 2, and 4 turned on at the same time) should equal the algebraic sum of the individual errors for bits 1, 2, and 4. Any deviation from this sum is defined to be superposition error or bit interaction. A list of measured (unadjusted) bit errors for a typical DAC as well as the difference between algebraic and electronic sums for bits 1–4 are given in Tables 1 and 2. The data are corrected for gain and offset. The measurements are subject to errors from the digital voltmeter used, which has a claimed linearity of 10 μV over this range. The results show superposition errors at or below half an LSB at 18 bits, which is low enough to permit 18-bit total linearity.

Figs. 9–11 are oscillograms showing differential linearity (DL) results from a portable demonstration fixture. The DAC input code is exercised at ± 4 LSBs around the critical "major carry" transition (011111 . . . to 100000 . . . transition). In each figure (a) shows unadjusted 18-bit DL for the device under test (each vertical division represents one nominal LSB value); (b) shows 18-bit DL after the bit calibration procedure

Table 1.

	Individual bit weights (V)		Error (μV)	18-bit error (LSB)
	Nominal	Measured		
Bit 1	5.000000	4.999971	-29	-0.76
Bit 2	2.500000	2.500010	+10	+0.26
Bit 3	1.250000	1.249990	-10	-0.26
Bit 4	0.625000	0.625030	+30	+0.77

(c) shows unadjusted DL for a 16-bit resolution (vertical sensitivity has been corrected to compensate for the larger LSB value at 16 bits); and (d) shows the 16-bit trace after adjustment. The lower trace in each oscillogram shows the switching of the bit 1 (MSB) input line, which is inverted due to the complementary input coding.

Fig. 9(a) shows a device that is monotonic at 18 bits without adjustment. The adjusted 18-bit performance [Fig. 9(b)] is virtually indistinguishable for this DAC.

Fig. 10(a) shows an unadjusted device with approximately 1 LSB of positive DL error at the major carry. Even at the 16-bit level [Fig. 10(c)] the DL error is perceptible. After adjustment [Fig. 10(b) and (d)] the device achieves true 18-bit performance.

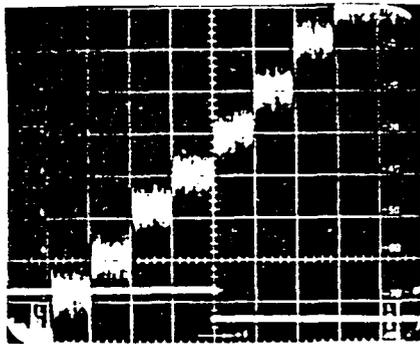
Fig. 11(a) is an example of a DAC with a slightly negative DL, causing nonmonotonic behavior at the major carry. Again, this part can be adjusted to 18-bit accuracy, as shown in Fig. 11(b) and (d).

4.3 Settling Time

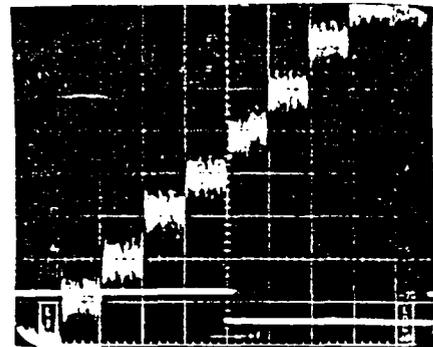
The primary measure of dynamic performance is the settling time. The 18-bit DAC was designed to minimize settling imperfections and reduce glitch energy compared to previous devices. Although excellent techniques have been developed for measuring high-speed settling at the 12–16-bit level, the direct determination of current-output settling time to 18-bit accuracy is a serious challenge, and a direct measurement has not yet been implemented. However, settling time can be

Table 2.

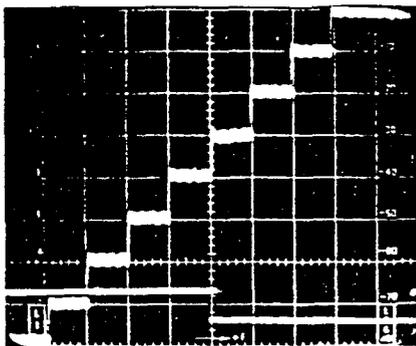
Bit combinations	Algebraic error sum (μV)	Electronic error sum (μV)	Superposition error (μV)
Bits 3, 4	+20	+16 μV	-4 (-0.10 LSB)
Bits 2, 4	+40	+40 μV	0 (0 LSB)
Bits 2, 3	0	+2 μV	+2 (0.05 LSB)
Bits 2, 3, 4	+30	+34 μV	+4 (0.10 LSB)
Bits 1, 4	+1	+18 μV	+17 (0.45 LSB)
Bits 1, 3	-39	-20 μV	+19 (0.50 LSB)
Bits 1, 3, 4	-9	+2 μV	+11 (0.29 LSB)
Bits 1, 2	-19	-6 μV	+13 (0.34 LSB)
Bits 1, 2, 4	+11	+16 μV	+5 (0.13 LSB)
Bits 1, 2, 3	-29	-22 μV	+7 (0.18 LSB)
Bits 1, 2, 3, 4	+1	0 μV	-1 (0.03 LSB)



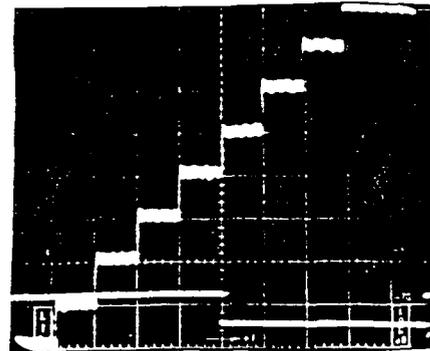
(a)



(b)

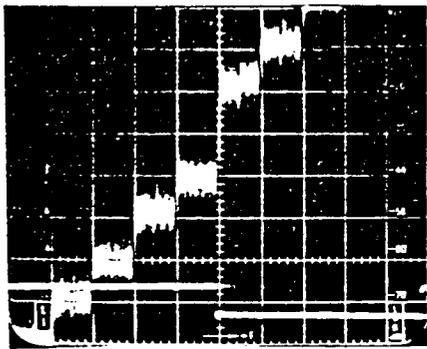


(c)

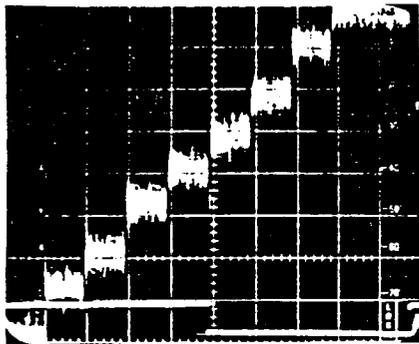


(d)

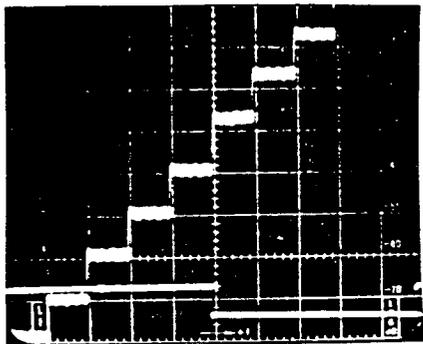
Fig. 9. Differential linearity performance at major carry, device A. (a) Unadjusted transfer function, 18-bit resolution. (b) Adjusted transfer function, 18-bit resolution. (c) Unadjusted transfer function, 16-bit resolution. (d) Adjusted transfer function, 16-bit resolution.



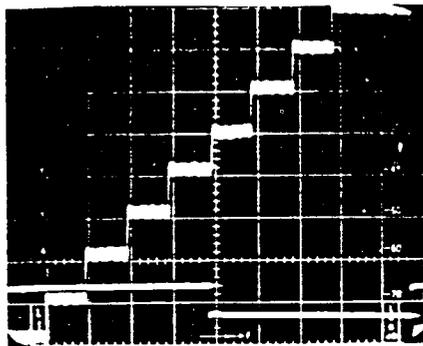
(a)



(b)

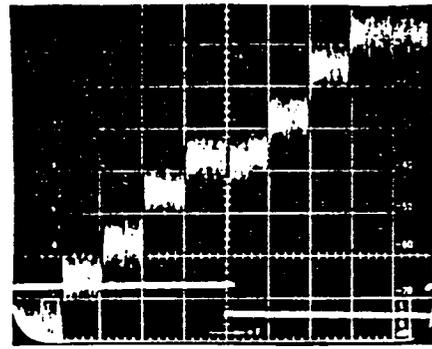


(c)

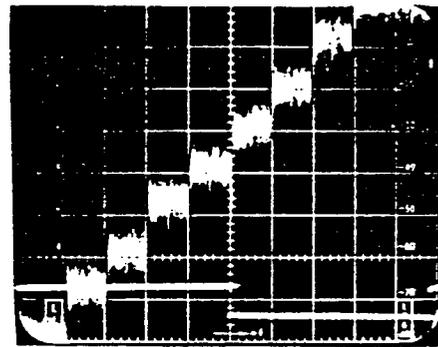


(d)

Fig. 10. Differential linearity performance at major carry, device B. (a) Unadjusted transfer function, 18-bit resolution. +1 LSB error at major carry. (b) Adjusted transfer function, 18-bit resolution. (c) Unadjusted transfer function, 16-bit resolution. (d) Adjusted transfer function, 16-bit resolution.



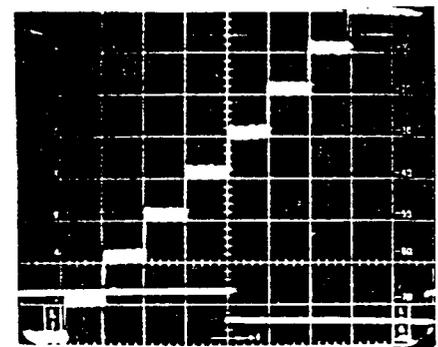
(a)



(b)



(c)



(d)

Fig. 11. Differential linearity performance at major carry, device C. (a) Unadjusted transfer function, 18-bit resolution. Nonmonotonic response at major carry. (b) Adjusted transfer function, 18-bit resolution. (c) Unadjusted transfer function, 16-bit resolution. (d) Adjusted transfer function, 16-bit resolution.

inferred with high confidence based on the DAC's performance in critical circuits. The DAC chip has already been incorporated into a high-performance ADC, in which the DAC must settle accurately for acceptable performance. In addition, a special test amplifier was built into the 18-bit chip to measure settling of the primary current-source bias line (the most likely cause of poor DAC settling). Based on results from these tests, the DAC's current output settles in approximately 300 ns to better than 0.001%. In most DAC applications, the settling time will be completely dominated by the characteristics of the output amplifier. Again, 18-bit results are not available, but lower accuracy measurements on the FET op amp (used in the hybrid) indicate

0.0015% settling in 2–4 μ s with no measurable settling "tails."

4.4 Distortion Performance

Sine-wave distortion tests at approximately 1 kHz have yielded results of around 0.0008% total harmonic distortion, which again challenge the limits of the available test equipment. Since transfer function linearity can be demonstrated at the 0.0002% level, the limitation on total harmonic distortion is primarily due to dynamic considerations. Great care must be used in the deglitcher design to minimize nonlinear switching products and feedthrough.

THE AUTHORS



J. M. Halbert

Joel M. Halbert was born in 1959 in Oak Ridge, Tennessee. He attended Yale University, where he received a B.S. degree in engineering and applied science in 1981. He then joined Burr-Brown Corporation in Tucson, Arizona, where he has been employed as a designer of high-speed and high-resolution data conversion products.

He served on the AES R-3 Subcommittee of the Digital Audio Technical Committee. In addition to the

AES, Mr. Halbert also is a member of the IEEE and Sigma XI.

•
Mark A. Shill received a B.S.E.E. degree in 1984 from the University of Arizona. He joined Burr-Brown Corporation as a design engineer and has been involved in the development of precision monotonic data converters and associated circuit functions. (Mr. Shill's photograph was not available at press time.)