


20 Bit "Colinear" DAC, a Solution to Low Level Problems

Preprint# 2963
Session-Paper# F-1-3

Fred J. Highton and Toshio Murota
Burr Brown Corporation
Tucson, AZ USA

**Presented at
the 89th Convention
1990 September 21-25
Los Angeles**



This preprint has been reproduced from the author's advance manuscript, without editing, corrections or consideration by the Review Board. The AES takes no responsibility for the contents.

Additional preprints may be obtained by sending request and remittance to the Audio Engineering Society, 60 East 42nd Street, New York, New York 10165 USA.

All rights reserved. Reproduction of this preprint, or any portion thereof, is not permitted without direct permission from the Journal of the Audio Engineering Society.

AES

AN AUDIO ENGINEERING SOCIETY PREPRINT

20 BIT "COLINEAR" DAC, A SOLUTION TO LOW LEVEL PROBLEMS

Fred J. Highton and Toshio Murota
Burr Brown Corporation
Tucson AZ 85734-1400
Atsugi Japan

ABSTRACT

Digital audio systems have traditionally used digital to analog converters (DACs) that employ laser trimmed, current sources in order to achieve sufficient accuracy. However, even the best of these suffer from potential low level nonlinearity due to errors at the major carry transition. More recently, DACs that utilize noise shaping techniques and very high oversampling frequencies have been introduced. These DACs do overcome the low level nonlinearity problem, but at the expense of signal to noise performance and often at the expense of channel separation and intermodulation distortion if the succeeding circuitry is not carefully designed.

A new solution is required, one that combines all the advantages of a conventional DAC: excellent Full Scale performance, high signal to noise ratio and ease of use with excellent low level performance. Two DACs used in a complementary arrangement provide an ideal solution.

The circuit employs two DACs that share a common reference and even a common R-2R ladder network to ensure perfect tracking under all conditions. By physically interleaving the individual bits of each DAC and employing laser trimming the highly accurate match required between DACs is achieved.

This new complementary linear, "Colinear", approach, which steps away from zero with small steps in both directions, avoids any glitching or "large" linearity error and also provides an accurate current output. The low level performance of this part is such that a 20 bit converter actually has validity.

BACKGROUND

In the search for high fidelity sound reproduction the limitations of analog systems became apparent. Digital signals were easier to store and recall without degradation and sixteen bit pulse code modulated (PCM) data was chosen as the standard. This allowed a system dynamic range of approximately 98 dB which is an entirely adequate level if judged on tests performed by the British Broadcasting Corporation. These tests demonstrated that a dynamic range of 86 dB was sufficient to ensure that most listeners would not detect any artifacts of digitization even on the most sensitive material [1]. Of the components required to build an economical system, the 16 bit digital to analog converter (DAC) presented a formidable challenge. In 1980 such a converter, which was only guaranteed to 14 bit accuracy, cost \$40.00. A decade later a dual 18 bit converter, guaranteed to 15 bits, is available in large quantities at under \$4.00. This level of cost reduction has enabled millions of people to enjoy the benefits of high fidelity reproduction but it is significant to note that although the cost dropped by a factor of ten during this period, the accuracy improved by only a factor of two.

After the first generation of 16 bit compact disc (CD) players were introduced the manufacturers sought to increase the features and reduce the cost of later models. Unfortunately, in many instances performance standards fell. This situation was allowed to develop because many reviewers and most consumers were not sophisticated enough to detect the difference. The reproduced sounds they were hearing were far superior to those of the analog systems to which they were accustomed. However, the situation was rightly brought to light by Lipshitz and Vanderkooy [2], who presented evidence in 1988, which showed that many players were not meeting the minimum 14 bit (86 dB) requirement. This revelation shocked the CD player manufacturers and their suppliers into action. Most of the ensuing attention was focused on the DAC despite the fact that the digital filters were equally criticized.

The DAC suppliers learned that while it was possible to correct for linearity shifts with an external adjustment, the Original Equipment Manufacturer (OEM) in most cases was unwilling to make it. Other solutions were sought. Many OEM's looked to new architectures such as oversampling converters. The DAC suppliers began to understand the reasons for the larger assembly shifts, associated with consumer packaging and took steps to overcome them. But the gauntlet was down, low level linearity was the goal and 14 bit performance was never going to be sufficient again.

SEARCH FOR A SOLUTION

It has been shown [3] that conventional current source DACs can meet the most demanding specifications if sufficient care is taken. But the particular product referenced above and its successor have been available to OEMs for quite some time and have only been used in the very highest performance CD players. The major

reasons are the relatively high price and the requirement for large numbers of external components. It is obvious that a true solution to this problem must be easy to use and low cost.

In response to the initial criticism [2], several OEMs sought exoneration by touting a new architecture. The situation gave impetus to a movement, already underway, to produce high resolution converters without the need for extreme analog accuracy. These oversampling converters seek to replace resolution in amplitude with resolution in time. They make use of oversampling and noise shaping techniques which rely on feedback and therefore are prone to instabilities. In fact, a host of new problems emerge. One is due to the unpredictable nature of the gain associated with a one bit quantizer, another is the possible generation of idle tones caused by partial correlation of the quantization noise. Although these problems may be reduced by adding complexity to the basic architectures, certain signal-to-noise limitations result. Since noise shaping creates large out-of-band noise this must be dealt with carefully not to introduce distortion which may be folded back into the base band. To make use of oversampling and especially multi order noise shaping techniques highly complex digital filters are required. These devices are themselves imperfect. They are prone to mathematical and round-off errors [2] and introduce large amounts of digital noise.

These new converters may have succeeded in overcoming low level linearity errors but it is at the expense of reduced signal-to-noise ratio (SNR). Although the specification of these converters show SNR well in excess of 100 dB at infinite zero (a pattern of repeated bipolar zero code), actual dynamic range measures well below these levels. Some of the limitations of one such converter are evident in Figures 16 and 18.

The "Achilles' heel" of conventional DACs, and that responsible for the problems highlighted in 1988 [2], is the error due to the most significant bit (MSB) toggling at bipolar zero (BPZ). In a 16 bit converter the MSB should be accurate to 16 ppm and maintain this over time and temperature which is a very difficult task. Any small MSB error together with the glitch caused by every bit changing state, may create an error that is audible at very low signal levels. If this error is shifted away from BPZ the situation becomes more tolerable. Some DACs now seek to solve this problem by employing a digital offset (accompanied by a compensating analog offset) to move the transition away from BPZ. In this way the transition occurring at BPZ will be that caused by some lower bit. For example, if the code is offset by the weight of bit 9 the transition error of bit 9 is substituted for that of bit 1 and a much smaller error results. Even though the original bit 1 error is still activated by signals as small as 50 dB below full scale (FS), the error is much less noticeable. There is therefore, a possible tradeoff between the size of error to be tolerated at BPZ and the smallest signal level at which the MSB transition has an effect. This technique may be audibly superior to current one bit solutions but is not the ultimate solution for the highly discerning listener or for more demanding professional applications.

One seemingly obvious answer is the true signed magnitude configuration which employs switching of the output amplifier between positive and negative

excursions of the input code. With this arrangement, it is only the least significant bit (LSB) which toggles on and off around BPZ and thus the smallest possible error and glitch occur at the lowest signal levels. This seems such a perfect solution and one which has been available to OEM's since the very beginning. Why is it not employed? One reason may be that the data is not coded in this way and needs to be converted, but considering some of the other complexities in the system, this would seem to be a trivial task. Another more plausible reason is associated with switching the operational amplifier. This must be accomplished in such a manner, that the switching glitch, offset and gain error combine to produce a total noise of less than -100 dB. This not only requires fast and accurate switches but a very expensive amplifier. This expensive circuitry could be added to the DAC integrated circuit but there are several arguments against this. Firstly, the non trivial switching difficulties have to be solved, secondly the varying power associated with the amplifier can cause DAC errors [3] and thirdly high quality OEM's prefer to select their own amplifier.

After careful consideration of all the options, the preferred solution is one in which two separate DACs are combined in a complementary arrangement to provide a current output which changes minimally around BPZ.¹ The user is then free to follow this device with an amplifier of choice and not be concerned with switching or deglitching networks. The only disadvantage of this approach, compared with a signed magnitude approach, is the added thermal noise of one DAC which is always on at BPZ. However, with careful design, this noise can be reduced to insignificance when compared to other noise sources. Figure 1 is a block diagram of this "Colinear" DAC and Figure 2 shows the code conversion between the input, binary twos complement (BTC) and that required for "Colinear" operation. This conversion is carried out on the chip and is transparent to the user.

From the coding it can be seen that only small bit changes occur either side of bipolar zero and since these bits can be matched to within one percent without any trimming, the linearity at BPZ will be accurate to 1/100 of one LSB or essentially perfect. At higher signal levels large bits are activated and the linearity degrades. But since the highest linearity is required for the lowest signal levels an increase in the resolution of the converter is appropriate. Noise, thermal and systematic will obviously place a limit on the achievable dynamic range and in the audio bandwidth, 20 bit resolution appears to be a practical limit.

DEFINING THE PRODUCT

Although this part will be attractive to professional audio users, it is primarily targeted at the high end CD player market. The goal is then to produce a 20 bit DAC that will overcome all the problems previously described and still be convenient and

¹Prior to publication of this paper, the authors became aware of a previous application of this technique based on a discrete board level implementation [6].

economical to use. Purists may well argue the need for 20 bits in this application but many OEMs are already employing 18 or 20 bits. They are using gain ranging architectures which suffer from noise modulation problems but they do demonstrate a desire for higher resolution. Moreover, if digital filters make 20 bits available, truncation errors are at least, moved from the 16 to the 20 bit level. The extra dynamic range may, in addition, facilitate the inclusion in the system of some digital signal processing (DSP). For example it would permit a 24 dB boost for digital tone controls or the provision of digital loudspeaker equalization [4].

From experience in the targeted marketplace, the importance of making the correct design tradeoffs is well understood. Three factors must be carefully considered and balanced accordingly. These are:

- 1) Performance Low level linearity problems must be eliminated while maintaining FS total harmonic distortion (THD) and SNR at least equivalent to the best digital audio DACs currently available.
- 2) Economy Pricing must be equivalent to currently available high end converters.
- 3) Ease of Use External adjustment and the use of external components must be largely eliminated. Operation from widely available power supplies must be ensured.

Consideration of the third factor implies the use of an "on chip" reference, serial to parallel converter and bipolar offset current generator. The inclusion of these circuits can have deleterious effects on performance [3] but with careful consideration to layout and bus routing an acceptable level can be achieved [5]. In 1986 a digital audio DAC was introduced which did not need an external deglitching circuit and as a result a glitchless output current is now mandatory. In addition, many customers wish to use an external resistor in their current to voltage converter and therefore an accurate current output is required. Furthermore, the part must operate from +/-5 Volt supplies to maximize its acceptability to the widest range of customers.

The second factor demands the adherence to very strict cost targets. The foremost of these is die cost which is dependent on the monolithic process chosen and the number of available dice per wafer. The selected 20 Volt bipolar process, which has been used for high precision audio converters since 1980, contains fast well matched NPN transistors with both high current gain and high output impedance. These transistors make excellent current sources and switches. The process provides stable nichrome resistors which have an extremely low temperature coefficient of resistance (TCR) and extremely accurate TCR matching. Besides its technical merits the process is running a large volume of product which makes it consistent and high yielding. These factors help to provide economic silicon. To maximize this advantage the chip area must be minimized and this leads to performance tradeoffs from such effects as matching, trimming and thermal and electrical feedthrough.