RECENT DEVELOPMENTS IN DIGITAL AUDIO TECHNIQUES

PREPRINT NO 1956 (B5)

By
K.J. Gundry, D.P. Robinson, and C.C. Todd
Dolby Laboratories
San Francisco

Presented at the 73rd Convention 1983 March 15-18 Eindhoven, The Netherlands

AN AUDIO ENGINEERING SOCIETY PREPRINT

Recent Developments in Digital Audio Techniques

K.J. Gundry, D.P. Robinson, and C.C. Todd

Dolby Laboratories, San Francisco and London

Abstract

An idea from the 1950's is re-introduced, re-evaluated and improved to produce a digital system of high quality with a bit rate of only about 250 kbits/second. Where channel widths or bandwidth considerations are such to prevent the use of higher quality 16-bit pcm systems, the new approach can be used giving a quality significantly superior to today's broadcast standards, at a greatly reduced cost to pcm, and possessing a fundamentally more robust resistance to many types of errors than pcm.

0. Introduction

The development of pcm systems has led to the requirement for rapid world-wide standardization of parameters. In this process, several simpler techniques for digital encoding and decoding of audio signals have been neglected. There are many applications where the complexity, performance, and expense of 16-bit pulse-code modulation are inappropriate. For several of these applications, the form of encoding known as delta modulation, and its various improvements, has significant merit. This paper will describe one such improvement, referred to as digitally controlled adaptive delta modulation, which the authors have been investigating, and will discuss the effects of digital data errors.

The output of a digital transmission or recording system differs from the original analog signal by the quantizing error. It is normal to consider this error as noise, since in general it subjectively resembles random noise. In this paper, the terms quantizing error and noise will be used interchangeably except where the distinction is made specifically.

1. Pulse-code Modulation

The use of pulse-code modulation (pcm) for high quality audio is well established; an excellent review of the subject appears in reference (1). Assuming a sampling frequency little more than twice the highest audio frequency, the signal-to-noise ratio of a pcm system with uniform quantizing steps (linear pcm) can be expressed as $\sqrt{1.5} \times 2^n$, where n is the number of (binary) bits per sample.

If 80 dB is taken as the minimum requirement for the signal-to-noise ratio of a high quality system, then the quantizing resolution must be at least 13 bits per sample (2); 14 and 16 bit systems are common (for example EIAJ format for recording pcm on video cassettes, compact disc (3)).

For many practical applications, linear pcm has major disadvantages.

a) Cost. Converters normally need components with a precision better than 1 part in 2^n, where n is the number of bits per sample (e.g., about 0.006% for a 14 bit system). Such precision is inherently expensive, and even with the use of more sophisticated techniques and higher production levels, the cost of converters will remain high.

b) Bandwidth. The minimum bit-rate for linear pcm with the required performance is in the region of 500 kbit/s per channel, and many current professional systems require bit rates in the range of 700-800 kbit/s. Some recording and transmission systems are only practical with a markedly lower rate.

To quote reference (1), "one might speculate that the optimum conversion system would be one that was perceptually matched to the human auditory system since quantization error would be designed for minimum perceptual degradation. Such a system would make extensive use of spectral masking, as well as forward and backward temporal masking." Our paper discusses approaches to optimizing conversion schemes, considering both performance and cost.

2. Companded pcm

In order to reduce the required bit-rates, companding techniques may be used.

a) Analog companders. Inadequate linear pcm systems (e.g., 10 bits/sample) can be combined with analog companders. If the noise reduction system is well designed, this approach can work satisfactorily. For some years there seemed to be reluctance to acknowledge the economic and technical virtues of such an approach (2), (4), but more recently the advantages are beginning to be acknowledged (5), (6).

b) Instantaneous digital companding. If the analog signal has been quantized to a high precision (14 or 16 bit), the transmitted bit-rate can be reduced by sending the samples with full resolution only when they are small in magnitude (low level input signals). Larger samples are sent with coarser resolution, and quantizing error is therefore increased (7), (8), (9).

c) Block digital companding. This is similar to b) above, but the reduction in resolution occurs over blocks of data (e.g., 16 or 32 samples) in accordance with the magnitude of the biggest sample within each block. This method gives a greater saving in bit-rate, but requires storage of blocks of data (4), (10).

Digital companders resemble analog wide-band companders in that the level of the quantizing error rises as the signal level rises, and the spectrum of the quantizing error is independent of the nature of the signal; that is, digital companders are fundamentally susceptible to noise modulation. Since the spectrum of the quantizing error from a well-designed pcm decoder resembles white noise, which to the human ear is predominantly high frequency, noise modulation is most likely to be audible in the presence of low frequency signals. Our experiments have shown that if the noise above 2 kHz rises to worse than 65 to 70 dB below full-scale in the presence of say a 200 Hz full-scale signal, then noise modulation will be audible on critical musical material. Similar figures can be deduced from (10) and (4). (Our work and these references indicate that (12) is overly optimistic, and should not be used as a criterion to judge the audibility of noise modulation.)

With moderate amounts of pre- and de-emphasis, (13), this requirement can be met if the minimum resolution in the presence of full-scale signals corresponds to a 10 bit system.

The British Broadcasting Corporation's NICAM 3 (11), a block range-switching scheme with initial quantization to 14 bits, digitally companded to 10 bits with pre-emphasis, is an example of an optimized system. Digitally companded systems whose resolution drops below 10 bits can be expected to give audible noise modulation under critical conditions.

Digital companding schemes are unattractive for consumer applications because they retain the cost disadvantages of linear pcm, resulting from the need for 14 or 16 bit converters, and the required bit rates remain moderately high (a minimum of about 320 kbits/s per channel if noise modulation is to be avoided).

3. Delta Modulation

Delta modulation (dm) can be considered as a special case of differential pcm in which n = 1; that is, there are only two quantizing levels. The ratio N of maximum signal to quantizing noise ratio can be expressed (derived from (8) pp. 200 on):

$$N \geq \frac{\sqrt{3}}{4\pi} \cdot \frac{F^{3/2}}{f_m f_s^{1/2}}$$

where f_m is the frequency of a sine-wave at full scale,

f_s is the audio bandwidth,

and F is the sampling frequency*.

The equality applies when there is no correlation between successive samples; in practice the sampling frequency F is normally much greater than $2f_s$, so that there is always a high degree of correlation. The resulting value of N is typically 2-3 dB better than that calculated from this formula.

If the quantizing step-size is fixed (linear dm), the full-scale signal level falls at 6 dB/octave with increasing signal frequency; the system is therefore slope limited. A linear dm system designed to accept high-quality audio (full-scale signals up to at least several kHz) with a bit-rate in the range of 200-500 kbit/s, will have a signal-to-noise ratio considerably inferior to pcm systems operating at the same bit rates. This has in the past led to the rejection of dm for high quality applications (15).

However, dm has several virtues which will be discussed in detail later in the paper.

a) Its circuitry does not require any expensive or high tolerance components; the cost is very small compared with any conceivable pcm system, and therefore dm is particularly attractive for consumer applications.

b) Delta-modulation is inherently much more tolerant than pcm of isolated bit errors (16) and therefore for some applications elementary error correction schemes are adequate.

c) The small number of components lead to smaller size and much lower power consumption compared with pcm.

---

* Other sources give formulas with slightly different constant terms; for example (14) gives 0.2 instead of $\frac{\sqrt{3}}{4\pi}$.

d) Since the sampling frequency in dm is much higher than twice the highest audio frequency, only very simple and low cost anti-aliasing filters are needed.

4. Companded dm

The slope-limited overload characteristic of linear dm, apparently disadvantageous, proves useful when companding is applied. Figure 1 shows the quantizing error of a linear dm system relative to the level of a full-scale signal as a function of the signal frequency (in this case, the sampling frequency is 250 kbit/s). The dotted line shows the equivalent information for a linear pcm system at about the same bit rate (8-bit quantizing). At low and middle frequencies, dm is markedly better. When the input signal frequency is low and would not therefore mask high frequency noise, a dm system optimized for that low frequency would have a good signal-to-noise ratio. When the input signal frequency is high, a system optimized for that high frequency would have a degraded signal-to-noise ratio, but the increased noise level would be masked by the signal.

An ideal compander would therefore adjust the parameters of the system so that at all times the signal just reaches full-scale, whatever its frequency content. This objective cannot be achieved by an analog compression-expansion system controlled by the analog signal itself, since it implies an infinite compression ratio, and therefore an infinite expansion ratio. A digital, or quantized, compander in which the step-size adopts one of a discrete number of values, just sufficient to accommodate the slope of the input signal at that instant, can provide an approximation to the ideal characteristic. Unfortunately this requires very high precision, comparable with that of pcm systems, because errors in step-size lead to discontinuities in the slope of the reproduced audio, and hence to additional wide-band quantization error.

However, by using a hybrid analog-digital approach, an almost ideal characteristic is achieved. The step-size is made continuously variable at a syllabic rate under the control of the bit-stream, building on techniques described in (17) and (18). When the data contains strings of 1's or 0's, indicating that the coder is lagging behind the slope of the input signal, the circuit tends to increase the step-size until the strings are broken up. The step-size therefore varies in accordance with the peak value of the slope, adapting so that the modulator is always close to full-scale. The method does not require high precision because the step-size changes continuously and at a syllabic rate, and hence any discrepancies between the ADC and DAC lead only to a low frequency amplitude modulation of the signal.

One of the unusual properties of this digitally controlled adm, or dcadm, is that there is no inherent maximum signal level at which overload is reached. In theory, as input signal levels are raised, the adaption continues to increase the step-size so as to accommodate the signal.

Real circuits resemble amplifiers whose limitations are introduced by power supplies or imperfections in the circu.try. Since the adaption is slope-dependent, these real limitations become apparent at high frequencies first.

5. Error Feedback

A well-known flaw of dm is that the quantizing error resulting from sine-wave inputs does not resemble random white noise, but contains discrete frequencies, commonly known as birdies.

Error feedback is a very old but not well known technique whereby the output of an imperfect transmission path is compared with its input and the discrepancy due to the imperfection is added to the input signal in such a sense as to reduce the discrepancy, in theory to zero. Note that this is different from negative feedback in which the whole output is fed back and imperfections are reduced by finite amounts.

The application of this technique to digital systems was first proposed in the 1950's (19) and it is ideally suited to over-sampled digital systems (20), (21). Delta modulation is a grossly over-sampled digital coding system, that is, the sampling frequency is typically much greater than twice the highest audio frequency. Hence much of the power in the quantizing error lies above the audio band. By applying error feedback within the audio band only, the quantizing error at audio frequencies, including the birdies, is reduced at the expense of increasing it above the audio band where it is of no significance (22). In practice it is not possible with economic circuits to reduce the error to zero, but birdies are reduced by 20 or more dB (effectively to inaudibility), and broad-band noise is reduced by 10 dB or so (see figure 1). The appendix contains a simple analysis of the technique as applied to dm (derived from (22)).

6. Program-adapting pre-emphasis

The performance of digitally controlled adm with error feedback can be further enhanced by a program-adapting response shaping system. This can be understood by considering the following signal conditions.

a) When signal amplitudes are small, high frequency emphasis in the coder and de-emphasis in the decoder can improve the subjective noise level by reducing the noise power in the part of the spectrum where the human ear is most sensitive.

b) When the input signal contains large amplitudes at low frequences alone, a degree of low frequency attenuation in the coder, and corresponding boost in the decoder reduces the extent of step-size adaptation and therefore reduces the change in high frequency noise due to the increases in step-size.

c) When the signal contains high amplitudes at high frequencies, the effect of both the high frequency emphasis and low frequency attenuation in the coder is to boost the low frequency noise in the reproduced signal; it is therefore desirable to reduce these response changes, or even to invert their directions.

Hence a family of emphasis characteristics of the form shown in figure 2 gives good results. Since the significant factor in determining the required characteristics is the slope of the input signal, which in turn is directly proportional to the step-size, it is possible to control the emphasis with the same information (derived from the bit-stream) that controls the step-size.

As with any signal-controlled syllabic compander, the adapting pre-emphasis of digitally controlled adm might be expected to display a lag between the onset of a large amplitude signal and the adaptation of the circuit, with a resultant need to compromise between speed of response and modulation distortion (23). In the adaptive pre-emphasis this compromise is greatly reduced by using a two-path configuration, in which the pre-emphasized components of the signal present in both encoder and decoder are subjected to non-linear complementary overshoot suppression, similar to that used in the Dolby analog noise reduction systems (24), (25).

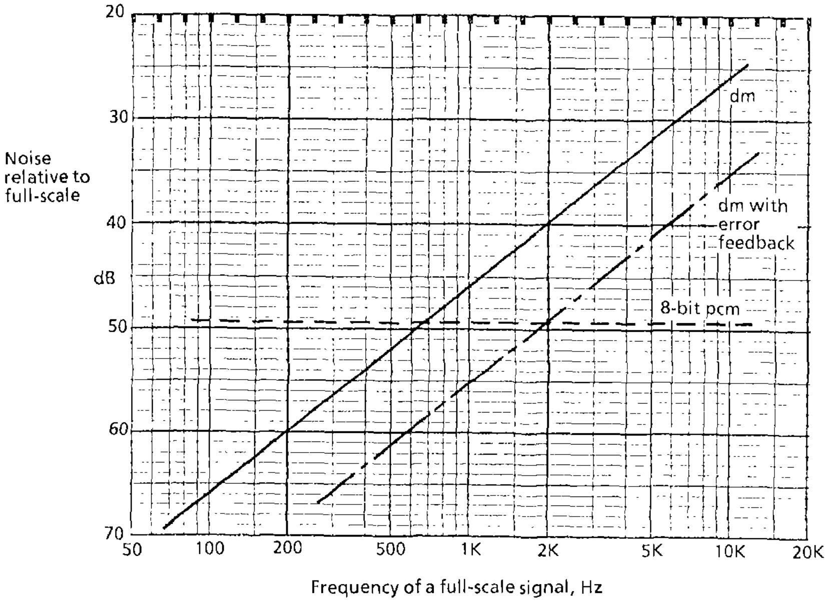

Figure 1: **Ratio of maximum signal to noise as function of frequency of a full-scale (sine-wave) signal.**

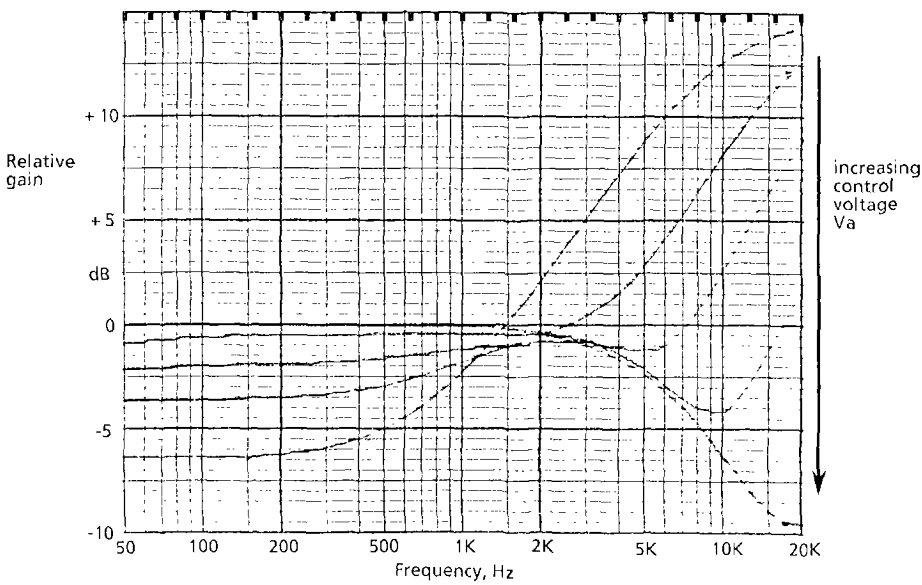

Figure 2: **Typical set of responses of adaptive pre-emphasis**
(measured results)

7. Practical system

Figures 3 and 4 show block diagrams of an example of a practical encoder and decoder. This has been built with general purposes components (linear bipolar operational amplifiers and CMOS logic) and performs as described in section 8 below.

For economical mass production a dedicated integrated circuit is desirable. A study conducted for us by Silicon Systems Inc., the California IC company known for its expertise in switched capacitor techniques, has concluded that the whole converter can be condensed to one chip with an area of 20 mm^2, requiring only a few non-critical external coupling and decoupling capacitors and a feed of clocking frequency.

7.1 Input filter (block 1)

Unlike pcm systems, dm uses a sampling frequency much higher than twice the maximum audio frequency. There is no need for high-order anti-aliasing filters since there is plenty of space between the top of the audio spectrum and the Nyquist limit (half the sampling frequency). Two- or three-pole low-pass filters are therefore sufficient, with obvious economic and performance advantages compared with pcm.

7.2 Program-adapting pre-emphasis

The band-limited signal passes via a response shaping network 2 and summation amplifier 3 to the delta modulator (blocks 8-14).

When the signal slope is small, high frequency components pass via blocks 4, 5 and 6 to summation amplifier 3, providing substantial high frequency boost (e.g., 12 dB at 10 kHz). At the same time low and middle frequencies are fed back negatively from the output of 3 via low-pass filter 7 and blocks 4, 5 and 6, giving an overall reduction in gain (e.g., 6 dB at 200 Hz and below).

As the slope of the processed audio increases, the control signal V_a derived from the delta modulator increases. This raises the frequency of the high-pass filter and attenuates the low frequencies, both in block 4, and thus reduces the high frequency pre-emphasis and increases the level of low frequencies.

A representative family of responses is shown in figure 2; each curve corresponds to a value of control signal, which in turn is proportional to the peak slope of the processed audio. At high slopes, the high-pass filter moves up beyond the audio band, and therefore the overall response tends to that of the response shaping network 2 in the main path. This is a low-pass filter at about 6 kHz, and reduces the degree of adaption for signals with predominant spectral components at the top of the audio spectrum.

7.3 Delta modulator

Blocks 10, 11, 13 and 14 constitute a conventional delta modulator in which the output data controls the polarity with which the control signal V_a is applied to the integrator 14. If for example a "1" in the output data indicates that the reconstructed audio is more positive than the input signal at the time of sampling, the up-down switch 13 feeds the appropriate polarity to cause the integrator output to move negatively, reducing the error. Similarly a "0" indicates a negative error and feeds the opposite polarity to the integrator.

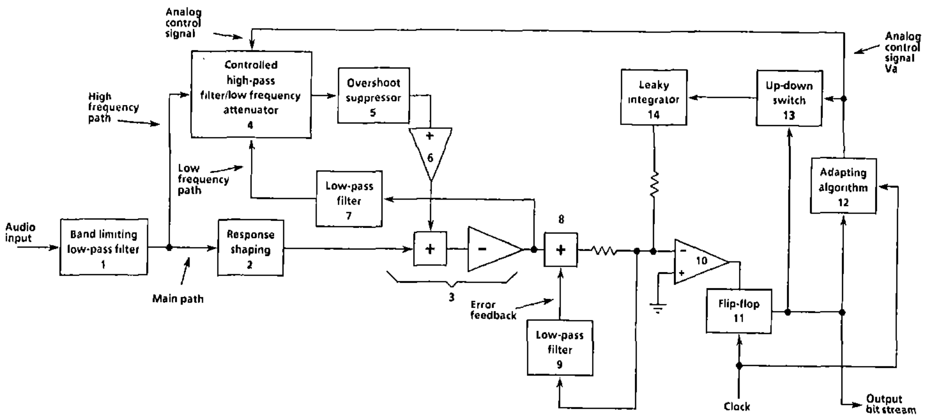

Figure 3: Practical Encoder

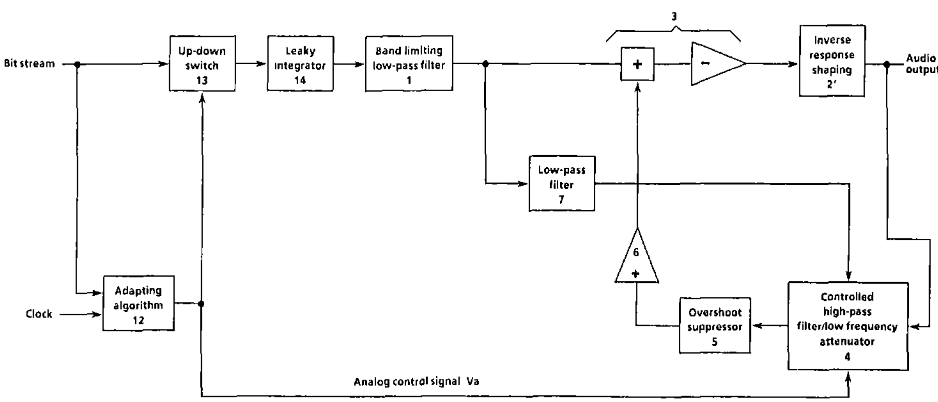

Figure 4: Practical Decoder

7.4 Adapting algorithm (block 12)

The rate at which the reconstructed audio can follow the input audio depends on the integrator characteristics and on the control signal V_a. If the output data of the encoder contains a series of "1's" or "0's" without transitions, the reconstructed audio is failing to follow the input, and it is necessary to increase V_a.

A particularly simple and effective adaption algorithm (by deFreitas) looks at the length of a current string of 1's or 0's. If the string is sufficiently short (say 1, 2, or 3 bits long) then the control signal V_a is allowed to decay at a syllabic time constant in the order of 50 ms. If a string of 1's or 0's exceeds some length, then V_a is increased. Longer strings cause the increase in V_a to become exponentially faster. V_a will ride up and down and stay slightly above the value required to prevent slope overload. The step-size is directly proportional to V_a.

Since the variation of V_a is syllabic, minor differences between the control signals generated in the ADC and DAC do not add significant noise; in practical circuits, no components need a tolerance tighter than ±1%, and audibly satisfactory results can be obtained with even looser tolerances.

7.5 Error feedback

The quantizing error at the input of the comparator 10 is fed to a low-pass filter 9 to select its audio frequency components, and is then added to the input signal. See the appendix for an elementary analysis.

7.6 Decoder (figure 4)

The decoder uses identical blocks rearranged and with the omission of the comparator 10, flip-flop 11, and error feedback components 8 and 9. In many applications the encoder may be switched to perform as a decoder, with a resultant economy.

8. Performance

8.1 Frequency response

The frequency response of the system described above is determined almost entirely by the input and output filters, which can be simple two- or three-pole low-pass filters. Assuming a sampling frequency of 250 kHz, a response up to 15 or 16 kHz (or even 20 kHz) can be achieved to any desired degree of flatness; with active filters using 5% tolerance components, ±1 dB to 15 kHz is a practical specification for consumer equipment.

8.2 Dynamic range

Figure 5 shows the overload point as a function of frequency; below 3 kHz this is determined by the clipping level of the amplifiers on the power supplies used.

The noise level is proportional to the value of the control voltage V_a. In principle, this can decay to zero, resulting in an output noise level of zero or an infinite dynamic range. In practical circuits, however, there will be some lower limit of control voltage. The adaptor will only allow V_a to drop to zero in the absence of any input signal. With real circuits there will always be an input signal consisting of amplifier noise, and there are also noise sources in the integrator and

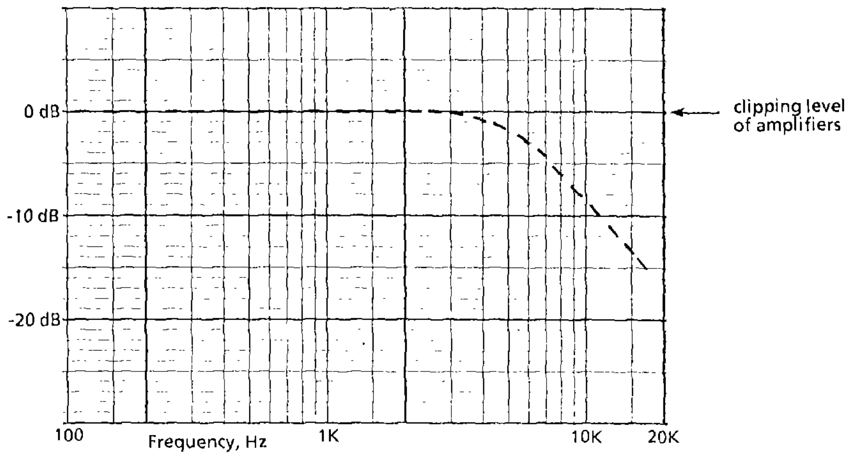

Figure 5: **Overload level as a function of frequency**

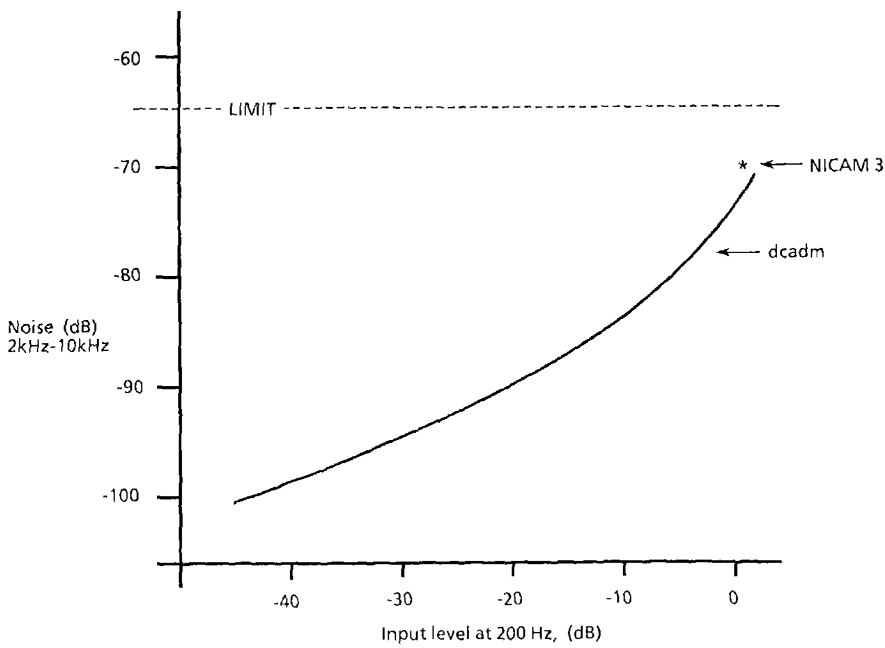

Figure 6: **High frequency noise v level of low frequency signal**
(measured results, 0dB = full-scale)

comparator. These noise sources combine and are converted into the bit stream. The adapter will always make a control signal large enough to encode this noise. The dcadm decoder will reproduce this noise with its own amplifier and integrator noise added. The combination of all these noise sources limits the bottom end of the dynamic range. Using typical integrated circuit operational amplifiers (e.g., TL072), the dynamic range is in the order of 100 dB.

8.3 Signal to noise ratio

Since the noise is a function of signal (as with any companded system, analog or digital), the signal-to-noise ratio is not the same as the dynamic range. Realistic assessment of noise in the presence of signals requires consideration of masking, and normally in practice employs extensive subjective testing. However, it is possible to make measurements which give a good indication of the probability of audible noise modulation, by examining high frequency noise in the presence of low frequency signals, and low frequency noise in the presence of high frequency signals, and comparing the results with limits arrived at by listening tests.

Figures 6 and 7 give examples of such measurements, together with the limits, and show that noise modulation is unlikely to be audible. To place this information in perspective, calculated points are given for the NICAM 3 (11), using CCITT pre-emphasis (+6 dB at 15 kHz).

Listening tests on critical musical material and artificially generated signals confirm that noise modulation is not a problem.

8.4 Dynamic performance

Figure 8 shows the typical response to a high amplitude high frequency transient. The overshoot suppression eliminates the normal transient distortion resulting from the finite response time of a syllabically adapting system; indeed the minor rounding at the beginning of the tone-burst is largely the result of band-limiting in the input and output filters. This test signal is more critical than any real musical or speech waveform.

9. Digital Data Errors

The reason for converting audio signals from an analog into a digital representation is often to assure more reliable storage or transmission of the signal, since digits may be stored or transmitted with essentially no error. Error free transmission or storage has a price however, in that excessive amounts of bandwidth or power or storage media may be required. Since practical systems are usually designed for efficient utilization of the storage media or transmission channel, there will typically be some finite error rate.

Digital errors may be dealt with in several ways. If the error rate is low enough, the errors may simply be ignored, and the resulting sound quality degradation tolerated. One method to deal with the errors is to detect which bits are in error by means of a parity or algebraic check code. The unreliable bits can then be discarded (turning the bad data into erasures instead of errors) and some type of error concealment can be employed in the digital to audio converter. The ideal method of handling errors is to transmit extra redundant bits along with the audio bits so that errors may be detected and actually corrected in the digital domain. However, an Error Detection And Correction System (EDAC) entails considerable complexity, and will itself fail when the raw error rate is sufficiently high, thus requiring error concealment in the DAC. In the design of a digital transmission or storage system the effect of digital errors on audio quality should be considered. The nature of the degradation of various types of A-D-A systems when subjected to digital errors can have a significant impact on overall system design.

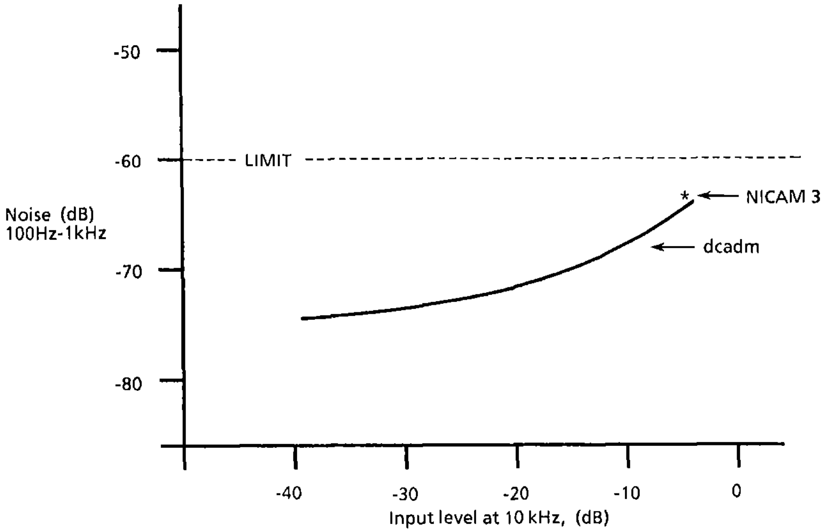

Figure 7: Low frequency noise v level of high frequency signal
(measured results, 0dB = full-scale)

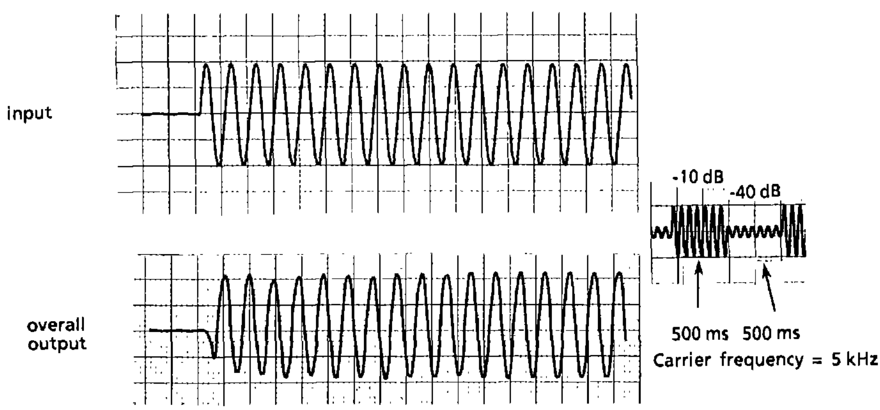

Figure 8: Typical response to a high level transient

The effect of digital data errors or erasures on pcm systems is well known. An isolated one bit error can result in an impulse of amplitude ranging from negligible (for an LSB error) to half of full scale (for an MSB error) as shown in fig. 9. Even a very low error rate can obviously have a severe impact on a pcm system. Isolated erasures, however, can be handled quite comfortably by a pcm system since straight forward linear interpolation can replace an erased value. Provided these erasures do not occur too frequently, the effect of this error concealment is quite innocuous. Linear interpolation can only conceal isolated erasures and not burst erasures; if burst errors are likely to occur, interleaving is usually employed so that a burst error is transformed into a repetitive error.

The effects of digital errors or erasures on a delta modulation system are quite different from on a pcm system. An isolated one bit error will cause the decoder to integrate in the wrong direction for one clock cycle, as shown in fig. 10. This has the effect of adding a small step to the audio waveform. The amplitude of this step is never large because every bit has equal weight--one quantum of resolution. Additionally, the spectrum of this disturbance is subjectively less annoying than that with pcm since it is shifted towards lower audio frequencies.

The innocuous effect on reproduced audio of a delta modulation bit error can yield some very great system advantages. Consider a digital transmission system with a bit error rate of 10^{-5}. This would lead to very audible impulses occurring a couple of times per second with pcm coding, but would be barely audible with dm. The pcm would require some form of error control in the transmission system; the dm would not.

If the error rate into the adm decoder becomes significant (i.e., greater than about 10^{-4}) the performance of the circuit described in section 7.4 will begin to suffer because the adaption circuit will begin to mistrack. Consider a single bit error which occurs in the middle of a string of 1's or 0's. The string which is broken in half would have produced a very large charging current for the syllabic filter. Instead, the error will cause the circuit to reset. The result is that the charging current will not reach as high a value and the decoder control signal will mistrack the encoder's control signal.

An alternative adaption algorithm has been developed which in the presence of errors helps to maintain a more accurate control signal. This scheme works by examining the density of bit stream transistions averaged over some number of bit periods (for example 12). Since the average transition density does not change dramatically with single bit errors, this algorithm maintains an essentially correct control signal value in the presence of a moderate number of errors. The decoder control signal will mistrack dynamically by an amount proportional to the error rate. The audible effect of this is that of amplitude modulation by a noise source. This alternative algorithm allows the adm to tolerate error rates of 10^{-3} to 10^{-4} with only modest audible degradation.

10. Applications

The performance specifications of digitally controlled adm with program-adapting pre-emphasis as described are similar to those of an optimally designed digitally companded pcm system such as NICAM 3. However the adm operates at a somewhat lower bit-rate, 250 kbit/s instead of 322, and eliminates the need for the expensive 14-bit converters and high-order anti-aliasing filters. The adm tolerance for errors means that a perfect transmission channel or storage system is not required. ADM is therefore very attractive for consumer applications, which include low cost digital tape recorders and direct distribution of program to the home via cable or satellite.

In those applications where the available data transmission rate is sufficient to accommodate pcm systems, there may still be virtues in deadm as described. Since the signal-to-noise ratio of adm is proportional to $F^{3/2}$, an increase in sampling frequency from 250 kHz to say 750 kHz (approximately the rate for professional 16-bit pcm) will reduce the noise level by 14 dB, giving excellent performance by any standards, while retaining the cost benefits and insensitivity to bit errors.

11. Acknowledgements

In addition to the many contributions made to this work by our colleagues, we would particularly like to acknowledge those of G.M. Jacobs and M.R. Smith.

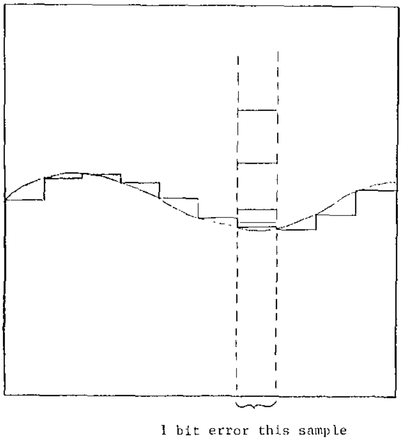

Figure 9a: Effect of a one bit error in a PCM system

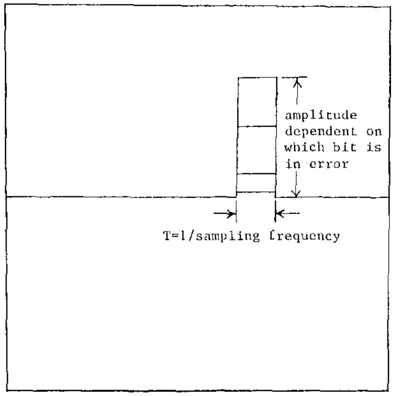

Figure 9b: Error generated by a one bit PCM error

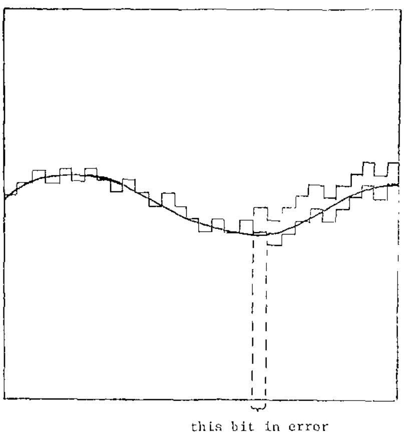

Figure 10a: Effect of a one bit error in a DM system

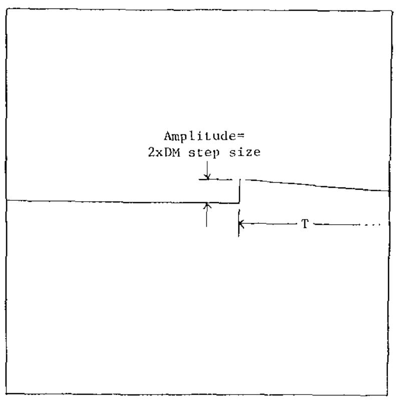

Figure 10b: Error generated by a one bit DM error
(The error decays at a rate set by the integrator leal

12. References

(1) Barry A. Blesser, "Digitization of audio", J. Audio Eng. Soc. vol. 26 pp. 739-771 (October 1978).

(2) "Pulse code modulation for high quality sound distribution: quantizing distortion at very low signal levels", BBC Research Dept. Report No. 1970/18.

(3) Compact disc, Philips Technical Review, vol. 40, 1982, No. 6.

(4) D.W. Osborne, "Digital sound signals: further investigation of instantaneous and other rapid companding systems", BBC Engineering No. 96, November 1973.

(5) K. Nakane et. al., "A new 8-bit PCM audio recording technique using an extension of the video track", IEEE Transactions on consumer electronics, August 1981, pp. 233-240.

(6) H.W.W. Korte, "Digitale Codierung von Tonsignalen unter Verwendung eines Dolby-A-Kompandes", Collected papers from Hörrundfunk 4, pp. 153-162, November 1976.

(7) P.F. Panter and W. Dite, "Quantization distortion in pulse-count modulation with non-uniform spacing of levels", Proc. I.R.E. vol. 39 pp. 44-48, Jan. 1951.

(8) K.W. Cattermole, "Principles of pulse code modulation", Iliffe Books Ltd. pp. 129 ff.

(9) A.H. Jones, "Pulse-code modulation for high quality sound signal distribution: a simulation of instantaneous companding", BBC Research Dept. Report No. 1969/26.

(10) M.G. Croll, D.W. Osborne, C.R. Spicer, "Digital sound signals: the present BBC distribution system and a proposal for bit-rate reduction by digital companding", International Broadcasting Convention 1974, IEE Conference Publication No. 119, pp. 90-95.

(11) C.R. Caine, A.R. English, J.W.H. O'Clarey, "NICAM 3: near-instantaneous companded digital transmission for high-quality sound programmes", Radio and Electronic Engineer vol. 50 No. 10, pp. 519-530, Oct. 1980.

(12) B. Blesser and F. Ives, "A reexamination of the S/N question for systems with time-varying gain or frequency response", J.Audio Eng. Soc. vol. 20 pp. 638-641 (Oct. 1972).

(13) CCITT Red Book Vol III p. 273.

(14) H.R. Schindler, "Delta modulation", IEEE Spectrum, vol. 7, pp 69-79 (Oct. 1970).

(15) C.J. Dalton, "Delta modulation for sound-signal distribution: a general survey", BBC Engineering No. 91 (July 1972).

(16) N.S. Jayant, "Digital coding of speech wave forms: pcm, dpcm and dm quantizers", Proc. IEEE vol. 62, pp. 611-632, May 1974.

(17) Wilkinson R.M., "A companded delta sigma modulation system", Report No. 68012, Signals Research and Development Establishment, Christchurch, Hampshire, U.K. (July 1968).

(18) R.E. DeFreitas, "Digital encoding circuitry", U.S. Patent No. 4254502 (1981).

(19) C.C. Cutler, "Transmission systems employing quantization", U.S. Patent No. 2927962 (1960).

(20) H.A. Spang, P.M. Schultheiss, "Reduction of quantizing noise by use of feedback", IRE Trans. Commun. Syst. vol. CS10 pp. 373-380 (Dec. 1962).

(21) D.J. Connor, R.C. Brainard, J.O. Limb, "Intraframe coding for picture transmission", Proc. IEEE vol. 60 pp. 779-791, July 1972.

(22) R.E. DeFreitas "Digital encoding circuitry with means to reduce quantization noise" U.S. Patent No. 4313204 (1982).

(23) D.E.L. Shorter, W.I. Manson, D.W. Stebbings, "The dynamic characteristics of limiters for sound programme circuits", B.B.C. Research Dept. Report No. EL-5 (1967).

(24) R. Berkovitz, K. Gundry, "Dolby B-type noise reduction system", The SERT Journal Vol. B, May/June 1974.

(25) R.M. Dolby, "An audio noise reduction system", J.Audio Eng. Soc. vol. 15, p. 383 (Oct. 1967).

Appendix Error Feedback

Figure 11a shows a simplified part of the delta modulator without error feedback. The input to the comparator is the total quantizing error e_t which can be considered the sum of two components, one containing only audio frequencies e_a and the other only frequencies above the audible range e_b.

$$e_t = e_a + e_b$$

The reconstructed audio, $-V_{in} + e_t$ therefore contains audio frequency components.

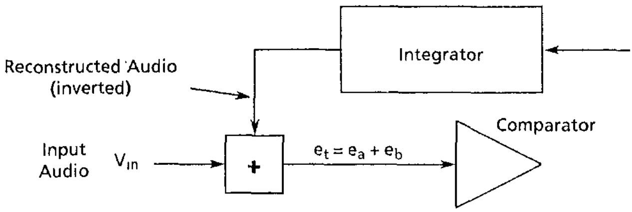

Figure 11a

Figure 11b shows the addition of error feedback. The audio component of e_t is selected by an audio frequency filter and is added to the input signal.

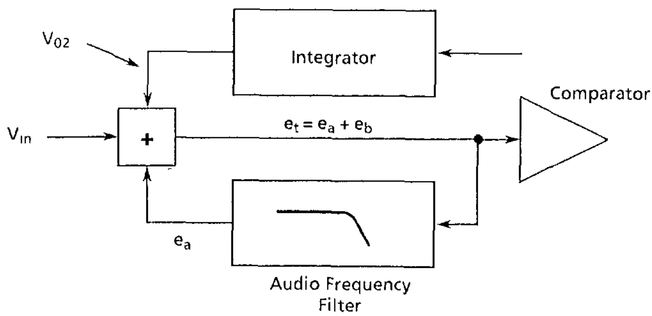

Figure 11b

Writing an equation for the summer, which is adding three input signals:

$$e_t = V_{in} + V_{02} + e_a$$

or, remembering that $e_t = e_a + e_b$

$$e_a + e_b = V_{in} + V_{02} + e_a$$

The e_a term cancels, and hence $V_{02} = -V_{in} + e_b$

Thus in theory V_{02} contains no audio frequency error.

This simple analysis does not take into account the limitations of practical filters, but with simple uncritical 2-pole filters it is easy to achieve about 10 dB of noise reduction.