Digital Audio Impairments and Measurements

N. H. C. GILCHRIST

British Broadcasting Corporation, Tadworth, Surrey, United Kingdom

In order to ensure the satisfactory performance of digital audio apparatus it is necessary to pay particular attention to providing a dynamic range which is sufficient for the audio signals, and to minimizing the errors and distortions which cause impairment. Coding range requirements for digital audio signals, and the audible impairments resulting from errors in the analog-digital conversion processes, from program-modulated noise and from bit errors have been studied by the BBC, and these are discussed in the paper.

0 INTRODUCTION

The recording or transmission of high-quality sound signals in digital form can offer a significant reduction in noise and distortion compared with analog recording and transmission, and digital audio processors can perform mixing, filtering and level-control operations with greater precision and versatility than their analog counterparts. The BBC has been using pulse-code modulation (PCM) for high-quality sound-program signals since 1972, when digital sound transmission equipment was introduced into the distribution network to replace analog circuits carrying the sound signals for radio and television from London (Broadcasting House) to the broadcast transmitters. Digital audio recording and processing equipment has also been developed in more recent years for applications in the studio.

Various impairments and distortions may be introduced into the audio signal during its conversion into digital form, during transmission or recording, or during reconversion into analog form. It is, of course, necessary to ensure that the levels of noise and distortion products contributed by analog circuits within the digital equipment are as low as possible, but there are certain types of distortion which are peculiar to digital processes. It is convenient to refer to these as errors, rather than distortions, as they can be attributed to fundamental errors resulting from the coding of analog signals in digital form, errors resulting from instrumental imperfections or bit errors introduced during transmission or recording.

1 IMPAIRMENTS INTRODUCED BY ANALOG-TO-DIGITAL AND DIGITAL-TO-ANALOG CONVERTERS

Conversion of an analog signal into digital form always introduces some distortion. This is because the digital codes which represent the samples of the original analog signal have only a finite number of values, depending upon the number of digits in the code, and after reconversion into analog form the instantaneous value of the signal may differ from that of the original analog signal by up to half of the quantizing step. Such quantizing errors introduce a form of distortion which normally sounds like random noise added to the signal, but when the signal is at a very low level, the effect is to impart a somewhat "gritty" sound to the signal, known as "granular distortion" [1]. As the number of digits (bits, for a binary code) is increased, so increasing the number of quantizing levels, the quantizing noise and granular distortion are reduced.

Because of practical limitations, the quantizing error always exceeds the theoretical figure. Consequently, the quantizing noise and distortion products are at levels somewhat higher than the theoretical in a practical converter, and the effective resolution is therefore less than the figure normally quoted as "resolution" by the designer or manufacturer of a converter.

Examples of analog-to-digital (A/D) converter errors which reduce the effective resolution are as follows:

1) Gain error, in which the full coding range of the A/D converter is not exploited by the audio signal at maximum level. The quantizing step therefore represents a more significant proportion of the useful coding range than is intended. This is illustrated in Fig. 1.

2) Offset error, in which there is a static shift of the transfer characteristic from the ideal, so that the overload characteristic is asymmetrical. Because of the asymmetry, the maximum signal applied to the A/D converter must be restricted, and so in this case too the full coding range is not used, with a consequent increase in the quantizing error (Fig. 2).

3) Irregularities in the clock signal which controls the sampling of the audio signal ("timing jitter"). Strictly speaking, perhaps one should not refer to this as "quantizing error," but the effect of the jitter is to increase the level of the random noise added to the signal by the A/D converter, thereby reducing the effective resolution.

4) Noise generated in analog circuits associated with the A/D converter or the digital-to-analog (D/A) converter. Unless the noise contributed by the analog circuits is at a very low level compared with the level of quantizing noise, this will reduce the effective resolution.

It is important to be able to determine the effective resolution of A/D and D/A converters for high-quality audio applications by making measurements during the development of conversion equipment, and during service use as a check on performance. The effective resolution of an A/D–D/A converter combination may be determined by measuring the level of quantizing noise, in the presence of an audio signal, and comparing this with the theoretical levels of quantizing noise for ideal A/D–D/A converter combinations with different resolutions.

For many years the BBC measured the level of quantizing noise introduced by the A/D and the D/A converters, using a 1-kHz tone applied to the A/D converter input, with a 1-kHz bandstop filter at the D/A converter output to remove the tone prior to the noise-measuring instrument. The arrangement is shown in Fig. 3. The distortion-measuring set incorporates a 1-kHz bandstop filter and an rms meter, and the 1-kHz tone source is an audio oscillator with very low distortion (in order to minimize the level of harmonic distortion products at the input to the A/D converter).

The 1-kHz tone method for measuring quantizing noise has been used successfully for measurements on converters with a resolution of up to 13 bits per sample. Attempts to use it to make measurements on converters with higher resolution have not been successful, and the reason for this appears to be the presence of harmonic distortion products in the output from the D/A converter. Although the harmonic products are at a very low level (and insignificant as far as their effect upon the sound quality is concerned), they prevent accurate measurement of quantizing noise at levels below the theoretical

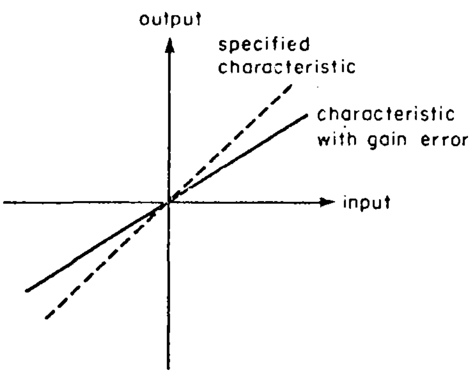

Fig. 1. Gain error.

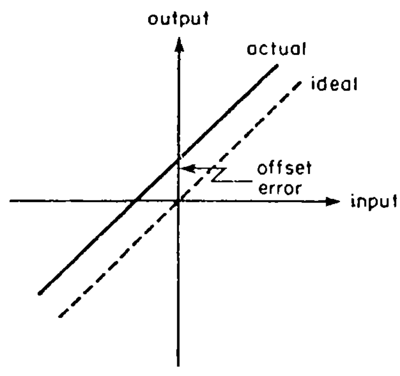

Fig. 2. Offset error.

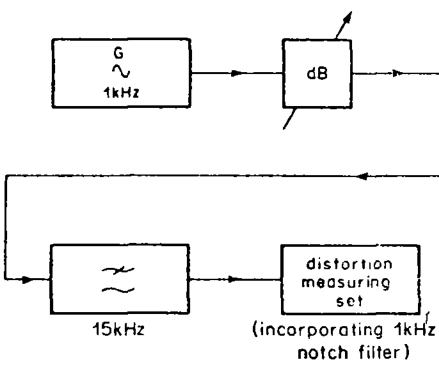

Fig. 3. Test arrangement using sinusoid at 1 kHz.

level for a 13-bit A/D and D/A converter.

If the 1-kHz tone is replaced by a tone at a much lower frequency, such as 30 Hz, a high-pass filter cutting off at about 300 Hz can be used at the D/A converter output to remove the fundamental tone and any low-order harmonics without significantly reducing the quantizing noise energy measured by the rms meter. The test arrangement is shown in Fig. 4. The distortion-measuring set of the previous arrangement has been replaced by two high-pass filters and an rms audio noise meter. A single 300-Hz high-pass filter was tried in the first instance, but gave insufficient rejection of the fundamental tone; a two-filter combination gave good results (distortion of the amplitude/frequency response, caused by reflections between the filters, was reduced by having different cut-off frequencies). The repeater coil which precedes the audio noise meter was found to be invaluable for reducing mains-induced low-frequency "hum," as was enclosing the filters in Mumetal boxes.

This test arrangement has been used to measure quantizing noise levels in audio digital converters with a resolution of up to 16 bits per sample.

Noise may also be detected at the output of a D/A converter when no input signal is applied to the A/D converter, or when the wanted signal at the input has an amplitude smaller than 1 quantum step. This "idle-channel noise" may be at a level higher or lower than the quantizing noise, and may change slowly in level and character over a period of time, depending upon instrumental imperfections, such as thermal drift and noise, in the A/D converter [2].

An ideal A/D–D/A converter combination would have a uniform transfer characteristic comprising a very large number of very small, equal quantizing steps. Departures from such an ideal transfer characteristic (which will be described as nonlinearity, for brevity) can be very important. It can occur in the analog circuits, taking the form illustrated in Fig. 5, in which case the degree of nonlinearity which can be tolerated is the same as for wholly analog sound equipment. However, another type of nonlinear distortion can occur in the digital conversion process when the quantum steps which form the transfer characteristic are not of uniform size.

This type of distortion, known as "differential linearity error," is illustrated in Fig. 6. If the differential linearity error exceeds the magnitude of one least significant bit, this can cause the slope of the transfer characteristic to reverse over part of the coding range;

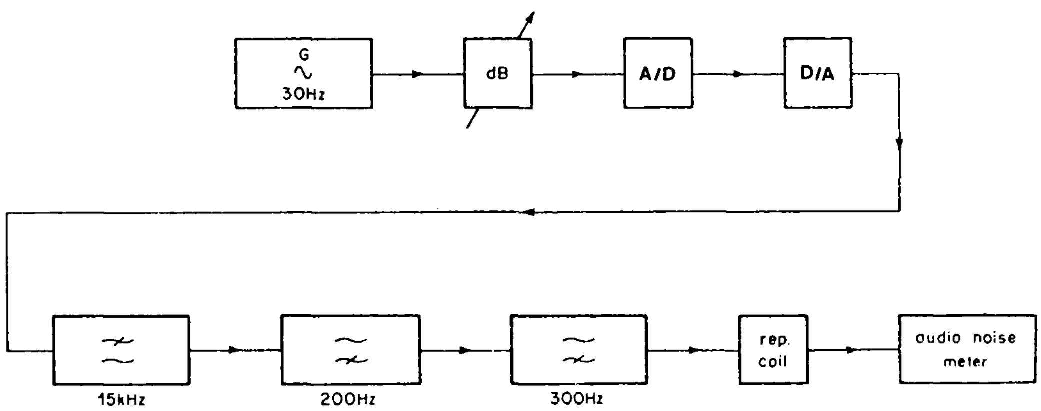

Fig. 4. Test arrangement using sinusoid at 30 Hz.

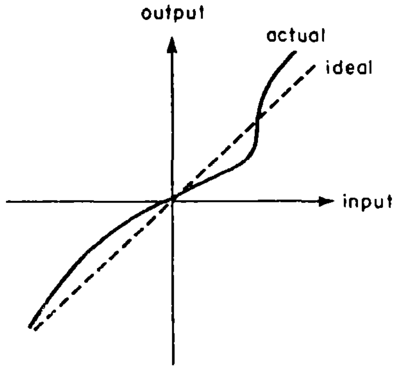

Fig. 5. Nonlinearity.

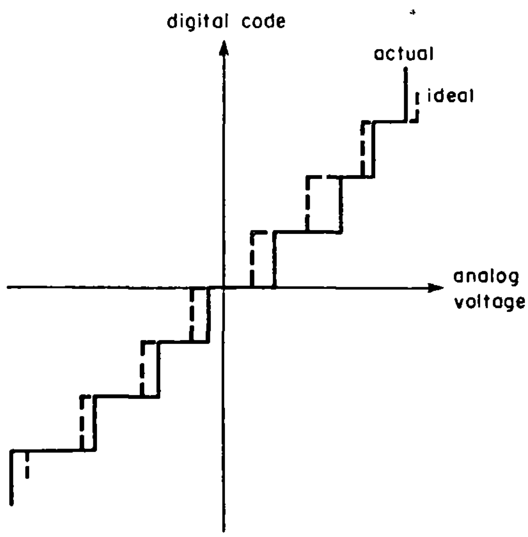

Fig. 6. Differential linearity error.

this nonmonotonicity would have serious consequences. An example of a nonmonotonic D/A converter transfer characteristic is shown in Fig. 7. Nonmonotonicity could be caused, for example, by an error in the value of one or more of the resistors forming the D/A converter ladder network. The magnitude of the error would need to be equivalent to at least one least significant bit for the characteristic to become nonmonotonic.

If a nonmonotonic D/A converter is incorporated in an A/D converter of the successive-approximation type, it causes one or more discontinuities in the transfer characteristic. The A/D converter is actually unable to form certain output codes, and the "steps" resulting from the missed codes do not in this case cause the slope to reverse. (In other words, the characteristic remains monotonic.)

The distortion caused by the differential linearity error may appear insignificant when measured by conventional analog methods, but the impairment resulting from even relatively small amounts of differential linearity error can be audible if discontinuities are created near the center of the transfer characteristic. This is because even a very small discontinuity can cause quite severe waveform distortion when the audio signal is at a low level. In a ladder-network D/A converter (Fig. 8) the greatest percentage accuracy is required of the resistor that determines the current switched by the most significant bit, if this kind of error is to be avoided. (In Fig. 8, this resistor is the one feeding the switch at the right-hand end.) If this resistor is sufficiently in error, the dicontinuity occurs at the center of the characteristic—in the worst possible place for an audio D/A converter.

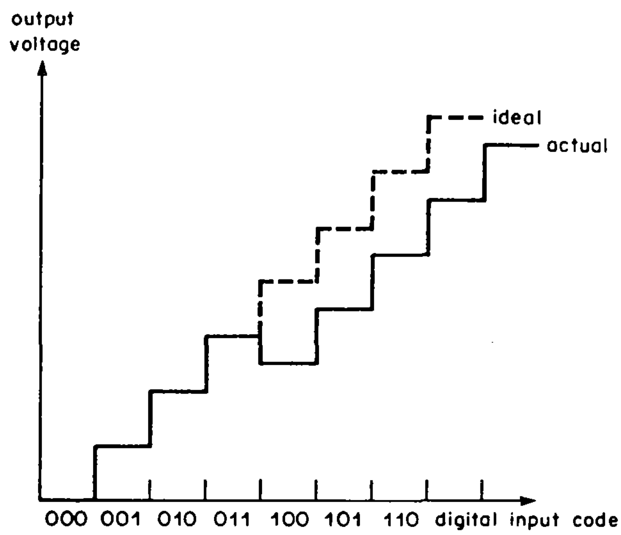

Fig. 7. Nonmonotonicity in a D/A converter.

There is a way of using the ladder D/A converter which shifts the most likely point for the discontinuity from the center to two positions midway between the center and the extremes of the transfer characteristic. This involves the arrangement shown in Fig. 9. The most significant bit determines the polarity of the voltage applied to the reference input of the ladder network, with the other bits controlling the network switches in the normal way. The input signal to this type of D/A converter must be in sign-magnitude coded form. The inverting reference voltage can be generated conveniently (and accurately) by another D/A converter. Slight inequality of the positive and negative reference voltages is relatively unimportant, as it cannot cause a "step" in the characteristic. Using this arrangement, an n-bit ladder network may be used to make a D/A converter with $n + 1$ bits per sample resolution. It is also fairly straightforward to make a floating-point D/A converter[1] by changing the magnitude of the reference voltage (once again, this can be effected very easily if the reference voltage is generated by a D/A converter) [3].

Because the conventional analog testing methods do not distinguish between differential linearity error and other types of distortion, it can be very difficult to

[1] Floating-point D/A converters (and A/D converters) use relatively low-resolution converters in an arrangement whereby the gain of the converter is changed, according to the amplitude of the audio signal, so that the resolution is increased when the signal level is low. By enabling a low-resolution A/D or D/A converter to operate effectively as a high-resolution audio digital converter at low signal level, while affording only a low resolution when the signal level is high (and when the higher level of quantizing noise can be masked by the signal), the floating-point technique provides the means to make audio digital converters at relatively low cost.

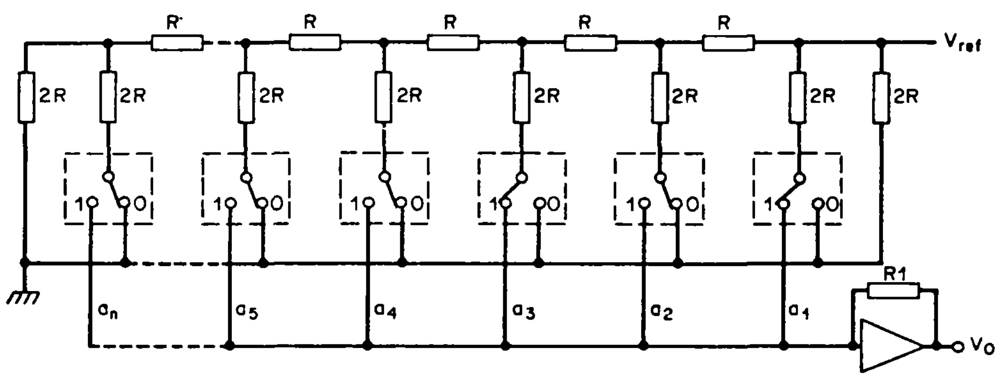

Fig. 8. Circuit diagram of ladder-network D/A converter.

detect the presence of this type of error in A/D and D/A converters, unless facilities are available for a listening test. The effect of the differential linearity error upon analog measurements and the degree of audible impairment caused are being studied by the BBC at the present time, with the intention of finding a suitable method for measuring the differential linearity error and of relating the measurement results to the impairment of the signal.

2 IMPAIRMENTS INTRODUCED DURING THE TRANSMISSION OF DIGITAL AUDIO SIGNALS

The BBC's first use of digital coding of sound-program signals in the broadcasting service was for point-to-point transmission. In 1972 a digital "sound in syncs" system was installed to carry the sound accompaniment to the BBC-1 and BBC-2 television signals over the major part of the television distribution network.

Pulse-code modulation was also introduced into the sound-program distribution network in the same year, when a 13-channel multiplex system was installed to carry signals from London (Broadcasting House) to the transmitter at Wrotham in Kent; an analog monochrome television super-high-frequency radio link was used to carry the PCM bit stream. The PCM system has been expanded over the years, and now it extends into Wales, Scotland, and Northern Ireland and forms a significant part of the BBC radio distribution network.

The analog television links carrying the PCM sound signals normally operate with a generous fading margin, and the digital signals received over these links are decoded without error for most of the time. When bit errors do occur, they normally do so infrequently and in a random manner. In the original PCM equipment, a parity bit is added to each sound-sample word to enable the presence of an error in the more significant of the bits to be detected. An erroneous sample is then replaced by the previous correct sample. This technique of "error concealment" is simpler to implement than error correction (in which the actual bit in error must be identified), but of course it is less than perfect. The previous sample is not necessarily a particularly good replacement for the erroneous sample, particularly at high audio frequencies, and if the error is in a relatively insignificant bit of the sample, it is often better to retain the erroneous sample rather than to replace it. An improved concealment system, in which a sample with an error in one of the more significant bits is replaced by a sample with a value derived by interpolation between the previous and the following correct samples, is employed in the BBC's latest NICAM[2] digital sound transmission equipment. This is illustrated in Fig. 10.

In the NICAM system [4], which is designed to code six high-quality sound-program signals for transmission over a digital bearer at the European primary multiplex bit rate of 2048 kbit/s, there is insufficient bit rate available to provide one parity bit per sample. Each parity bit therefore protects the five most significant bits of three samples, interleaved with other samples protected by other parity bits, so that a single error causes three samples to be replaced. This is clearly rather less satisfactory than using one parity bit per sample, because two perfectly good samples are replaced unnecessarily for every sample which has a most significant bit error, and also because the very act of

[2] NICAM stands for "near-instantaneous companding audio multiplex."

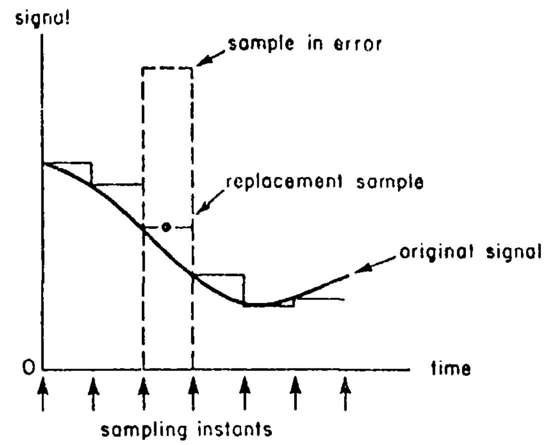

Fig. 10. Error concealment by interpolation.

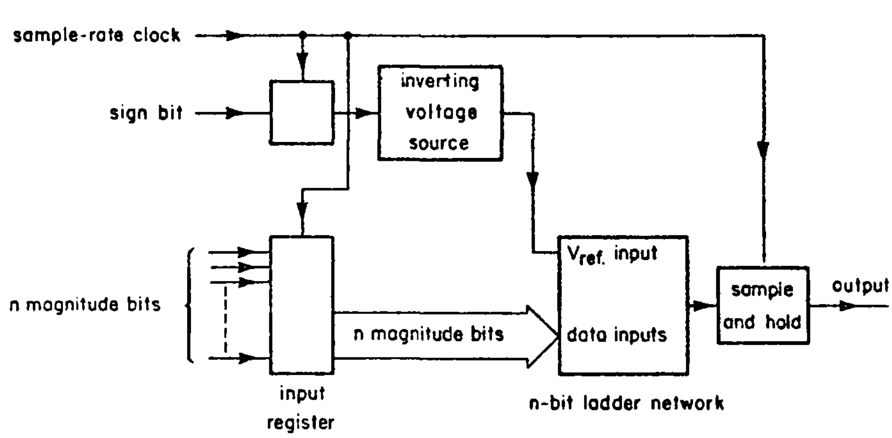

Fig. 9. Sign-magnitude D/A converter.

combining several samples in one parity group increases the size of the protected group (and therefore the risk of a double error occurring within the group with consequent failure of the error-detection process). Despite these shortcomings, the interpolation concealment works well in the NICAM-3 system, maintaining a performance substantially free of impairment at random bit-error ratios up to 1 in 10^5. Fig. 11 shows, in graphic form, the results of preliminary listening tests conducted to determine the impairment suffered by audio signals conveyed by NICAM-3 at different bit-error ratios. For some of the tests the error concealment was disabled. At a bit-error ratio of 1 in 10^4 the impairment is "perceptible, but not annoying" with this concealment system, and even at 1 in 10^3 it is little worse than "slightly annoying." Each plot point in Fig. 11 is the mean of 28 results, and the standard error is about 0.1 grade.

In the NICAM-3 transmission system, the audio signals are converted initially into 14-bit linearly coded digital samples. These digital samples are then compressed to 10 bits for transmission [4] in order to reduce the bit rate (and so accommodate six sound channels in a 2048-kbit/s bit stream). At the decoder, a digital expansion process re-expands the digital signal in a manner which complements the digital compression. There is a penalty to be paid for this compression–expansion, however. High-level audio signals are accompanied by a higher level of quantizing noise, because they are coded with lower resolution than low-level signals. This program-modulated noise must be kept at a sufficiently low level to be unobtrusive, and preferably masked by the sound of the audio program signal for most of the time.[3] Extensive listening tests were carried out by the BBC several years ago to compare the impairments introduced by different digital audio companding systems [5]. NICAM-3 gave the best performance of any system providing six high-quality sound channels in the 2048-kbit/s bit rate, and was therefore selected by the BBC for development into the equipment that is now in service.

[3] NICAM-3 employs CCITT pre- and deemphasis to reduce the audibility of program-modulated noise.

3 IMPAIRMENTS CAUSED BY ERRORS IN DIGITAL AUDIO TAPE RECORDING

Bit errors may be introduced into digital audio signals during the recording of these signals on tape and the subsequent replaying. Unlike a digital transmission system, in which individual errors may result from unrelated (and therefore random) events or disturbances, the most likely cause of errors in a recording on magnetic tape is a defective area on the tape (a "dropout"). Under these circumstances, hundreds of bits may be corrupted and a single-error correcting code or a simple error-concealment system is likely to be ineffective in overcoming the impairment, unless a relatively complex interleaving arrangement is adopted involving the storage of several hundred samples, all with their associated parity bits.

The tape defects which cause dropouts are normally up to 20 thousandths of an inch (0.5 mm) in diameter. The multichannel digital audio recorder built by the BBC [6] uses 40 tracks on 1-in (25-mm) magnetic tape to provide 10 recording channels, with the recorded bits of data protected by a Wyner–Ash (2, 1) convolutional error-correcting code. It is rare for the tape dropouts to exceed 20 thousandths of an inch (0.5 mm) in diameter, and so it is unlikely that a single dropout will affect more than one track when there are 40 tracks across a 1-in (25-mm) tape. Such a dropout can corrupt 300 consecutive bits in the track affected at a tape speed of 20 in/s (0.5 m/s), with 32-kHz sampling rate and 14-bits-per-sample coding.

By arranging for the data bits and associated parity bits to occupy alternate tracks in a four-track-per-channel arrangement, as illustrated in Fig. 12, one can ensure that a single dropout is most unlikely to affect more than one in any group or four consecutively coded bits. The Wyner–Ash (2, 1) code forms parity (indicated by parity bit P_n) from data bits D_n and D_{n+1}, as shown in Fig. 12. P_1 is the parity formed from D_1 and D_2; P_2 is the parity formed from D_2 and D_3, and so on. Figs. 13 and 14 show the parity failure pattern caused by errors in a data track or a parity track. The former always cause parity failures involving an even number

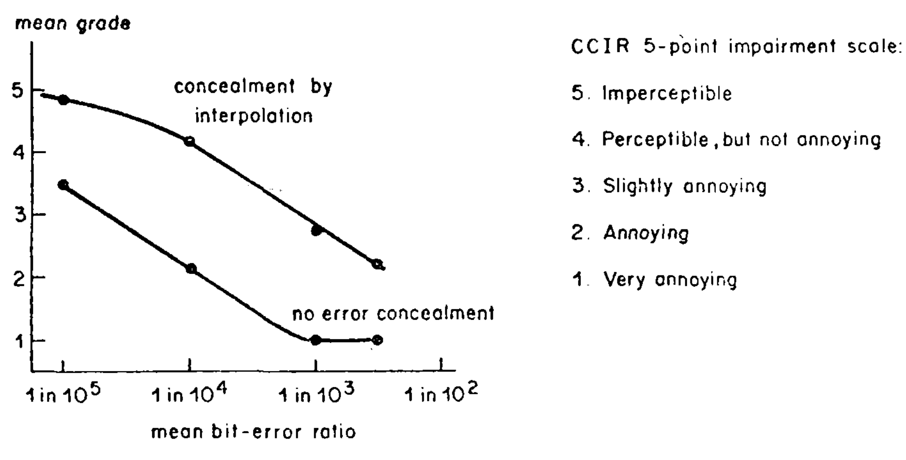

Fig. 11. Subjective impairment caused by errors (NICAM-3).

of parity bits, and by examining the parity bits in pairs from the start of the error, the data bits which are wrong can be detected and put right. Errors in a parity track, however, involve only single-parity bits, and no action need be taken. Bursts of errors of any length can be corrected, provided that no more than one track of the four for any channel suffers errors at a time [7].

Because four tracks are used to provide each audio channel in the BBC recorder, the maximum number of channels which can be provided economically on 1-in (25-mm) tape is limited—in this machine it is 10. The need for machines providing two or three times as many channels has led to the development of error-correction systems which can be used in a one-track-per-channel arrangement. Tape speed needs to be increased in order to record the bit rate for each channel within a single track, and the BBC has cooperated with the 3M Company in the development of a single-track-per-channel error-correction system for digital sound recording [7], [8].

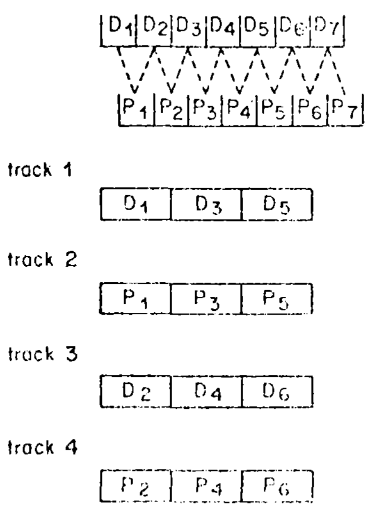

Fig. 12. Format for Wyner–Ash (2, 1) code in four-track-per-channel recorder.

The ten-channel recorder has proved an extremely useful laboratory tool for research work in the BBC. The limited number of channels has not proved a serious constraint in this application, and it is frequently used both as a source of high-quality, low-noise recorded program material and as a store for digital audio signals at various stages of digital processing. It has, for example, been particularly useful for recording sequences of digitally processed sound signals for use in listening tests and demonstrations.

The error-correction system is most effective, and it is very rare for there to be any impairment of the audio signal resulting from bit errors.

4 CODING RANGE REQUIREMENTS FOR DIGITAL AUDIO

Having considered the impairments which may be caused by instrumental imperfections or bit errors in digital audio apparatus, it is also important to consider the need to provide an adequate coding range (that is, an adequate number of bits per sample for coding the audio signal in digital form). This should be sufficient to permit the highest level signals to be reproduced without the impairment caused by "clipping," implying that a margin must be provided in order to avoid the inadvertent overloading of the system by signal peaks exceeding the expected maximum level, while preserving a satisfactory signal-to-noise ratio for the lowest level signals. To this coding range requirement (covering, effectively, the "dynamic range" of the signal) one may need to add an additional requirement to cover the wide range of signal levels encountered within the studio from different signal sources, assuming that all control of the signal level takes place in the digital studio desk. (In other words, no gain controls are provided in the analog sources.)

The BBC's experience has shown that once a signal has been controlled in level within the studio, and any unexpected occasional signal peaks in excess of the

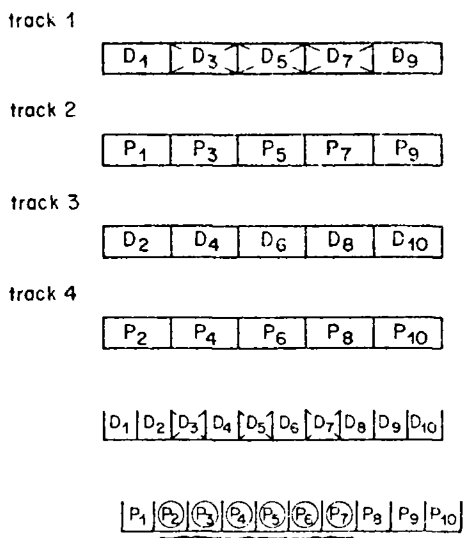

Fig. 13. Wyner–Ash (2, 1) code. Parity failure pattern for information track errors.

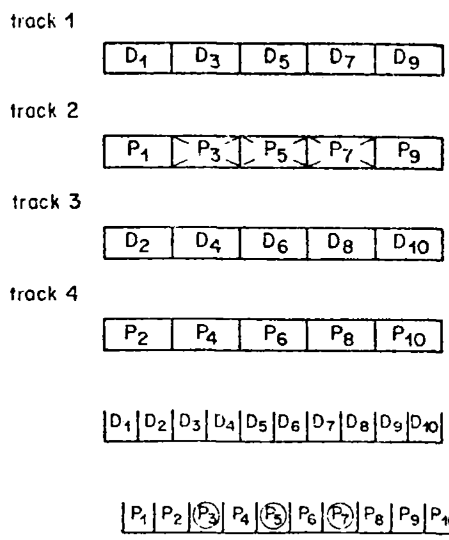

Fig. 14. Wyner–Ash (2, 1) code. Parity failure pattern for parity track errors.

nominal maximum sound-program level have been controlled by a limiter, a coding range of 14 bits per sample is sufficient. The BBC's NICAM equipment (mentioned in Section 2) codes audio signals initially with 14 bits per sample, and protection against overloading is provided by limiters which take account of the CCITT preemphasis applied to the audio signal. These variable-emphasis limiters are designed to control the signal peaks in a manner that minimizes audible impairment [9].

Without such a protective limiter an "overload margin" of at least 2 bits per sample should be provided (requiring a coding range of at least 16 bits per sample) for signals which have been subjected to control in level, in the studio desk. Prior to this level control, the signals from the various studio sources can have a dynamic range requiring up to 20-bits-per-sample coding, assuming that the maximum signal levels from different sources are adjusted to be the same, by means of some preset control for the analog source signal. If a preset control is not provided, or cannot be adjusted reliably prior to the commencement of a program, a coding range of up to 24 bits per sample may be necessary.

Fig. 15 shows a simple block diagram of a digital audio studio with the suggested coding range requirements indicated. The digital audio master recorder requires a coding range of at least 20 bits per sample, because it may not be possible to ensure that the level of every signal applied to the input will be controlled to reduce its dynamic range. It should, however, be possible to ensure that signals from different sources are all preset to the same nominal level. It is most important for the multichannel master recorder to be able to record the individual source signals without significant impairment, to avoid a need to rerecord performances or parts of performances solely for technical reasons.

Audio digital converters with a resolution of 24 bits per sample are not yet available. However, it is not necessary to provide such a high resolution in order to achieve a 24-bit coding range. Floating-point techniques can be used to obtain a wide coding range, using A/D and D/A converters with relatively low resolution [3]. The resultant program-modulated noise accompanying high-level audio signals will not be obtrusive, provided that the minimum resolution, determined by the converter at the "heart" of the floating-point system, is sufficient to give a satisfactory signal-to-quantizing-noise ratio for signals at the maximum amplitude.

There might also be some value in making the recorders operate in the floating-point mode, if the need to obtain a large number of channels conflicts with the requirements for coding range and low tape speed.

5 CONCLUSIONS

Studies to maintain a high level of technical quality form an important part of engineering research in the BBC. In the field of digital audio, listening tests have provided valuable information on the effect of imperfections in A/D converters, on the effectiveness of different error-correction and error-concealment systems, and on the impairments caused by program-modulated noise in companded transmission systems. In fact, digital audio equipment under development or study at BBC is always subjected to listening tests, in addition to objective measurements, to determine nonlinear distortion and quantizing noise levels. In work of this nature it is important to relate the results of measurements made on equipment in the laboratory or in service

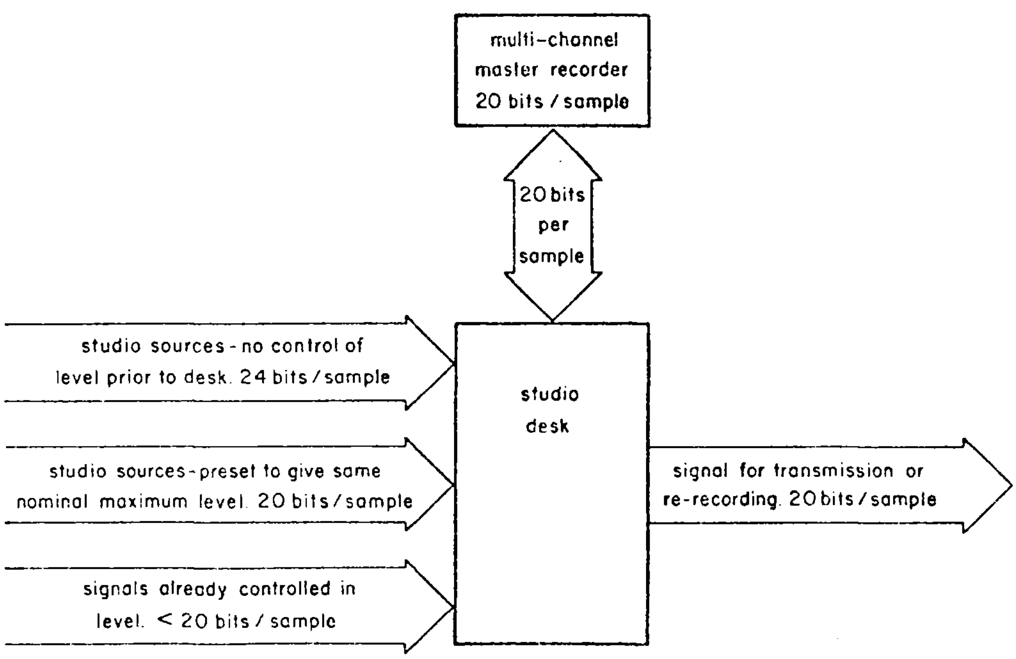

Fig. 15. Studio coding range diagram.

to the degree of impairment experienced by the listener, but this is not an easy task. Whenever possible, staff experienced in listening critically to reproduced program material are included among the listeners in subjective tests.

The study of errors in audio digital converters, mentioned in this paper, is still continuing, and particular attention will be given to the results of subjective tests and their relation to the measured results.

6 ACKNOWLEDGMENT

The author is indebted to the Director of Engineering of the BBC for permission to publish this paper. He also wishes to thank his many colleagues who have contributed to this work, and in particular F. A. Bellis who provided information on the digital audio recorder.

7 REFERENCES

[1] M. G. Croll, "Pulse Code Modulation for High-Quality Sound Distribution: Quantising Distortion at Very Low Signal Levels," BBC Research Dept., Rep. 1970/18, 1970.

[2] S. R. Ely, "Idle-Channel Noise in p.c.m. Sound-Signal Systems," BBC Research Dept., Rep. 1978/4, 1978.

[3] N. H. C. Gilchrist, "Analog-to-Digital and Digital-to-Analog Converters for High-Quality Sound," presented at the 65th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 28, p. 372 (1980 May), preprint 1583.

[4] C. R. Caine, A. R. English, and J. W. H. O'Clarey, "NICAM 3: Near-Instantaneously Companded Digital Transmission System for High-Quality Sound Programmes," *Radio and Electron. Eng.*, vol. 50, pp. 519–530 (1980 Oct.).

[5] G. W. McNally and N. H. C. Gilchrist, "The Use of a Programmed Computer to Compare the Performance of Digital Companding Systems for High-Quality Sound Signals," *EBU Rev.—Tech.*, no. 178 (1979 Dec.).

[6] F. A. Bellis, "A Multichannel Digital Sound Recorder," presented at the Int. Conf. on Video and Data Recording, Birmingham, England, *IERE Conf. Proc.*, no. 35.

[7] F. A. Bellis, "An Error-Correcting System for a Multi-Channel Digital Sound Recorder," BBC Research Dept., Rep. 1978/28, 1978.

[8] M. R. Brookhart and F. A. Bellis, "An Error-Correcting System for a Multi-Channel Digital Audio Recorder," presented at the 58th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 25, p. 1072 (1977 Dec.), preprint 1298.

[9] W. I. Manson, "Frequency-Dependent Limiters for f.m. Sound Transmitters," BBC Research Dept., Rep. 1973/25, 1973.

THE AUTHOR

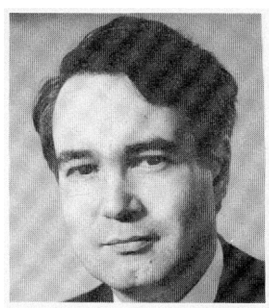

Neil Gilchrist joined the BBC after graduating from the University of Manchester with a degree in physics and electronic engineering. In 1967 he obtained an appointment as an engineer in the BBC's research department, working initially on aerials and the measurement of low-level harmonic signals produced by television transmitters. At the present time he is a senior engineer concerned principally with digital sound and television signals.

Mr. Gilchrist contributed papers to the 1978 and 1980 International Broadcasting Conventions, and to the 65th AES Convention. He has had papers published in *The Radio and Electronic Engineer*, the *EBU Review*, and *BBC Engineering*.