Mitsubishi X-86

Studio Sound

April 1990

Mitsubishi's X-86 2-track digital recorder

THE TECHNIQUES

A number of interesting questions are raised by the recent introduction of 'super digital converters', such as the Data Conversion System *DCS-900* used by engineer Tony Faulkner at recent sessions in Walthamstow Town Hall in conjunction with Mitsubishi's *X-86* 2-track digital recorder. The *DCS-900* is an analogue-to-digital converter for professional audio applications that aims to offer the user a potential improvement in sound quality over the converters supplied as standard in most audio equipment. Using oversampling techniques, it claims to offer increased low level resolution, quantisation to more than 16 bits, and very low noise and distortion.

In order to take advantage of the increased number of bits per sample, which the *DCS-900* can provide from its AES/EBU outputs, a tape machine which has the format space to accommodate 20 bit samples is required. Currently, there are very few commercial formats with this capability, the 2-track ProDigi format being the only readily available tape format that will handle 20 bits when fed from its digital inputs. The point of this article is to consider what advantages there may be in such an approach when the final master is going to end up in the compact disc format, quantised to 16 bit resolution, as it may well be supposed that any advantage in recording to 20 bits might be lost in the mastering process.

Background

To date 16 bits has been accepted as the norm for digital audio sampling in professional audio, as it is the resolution used in CD, and this is the format that nearly all consumers of digitally recorded music will use. Sixteen bits is also considered by many to be an adequate resolution for a digital tape recorder although there are those who believe that nothing less than 18 or 20 bits could be enough. A similar argument has raged about the sampling rates used in professional audio (44.1 or 48 kHz), with many claiming that they are nowhere near high enough. The question that has to be answered is to do with whether the push for more bits and higher sampling rates is simply a marketing exercise, together with some dubious subjective reasoning, or whether there is any ground for supposing that there is value in increasing the quality of conversion.

An important factor in all this is that many converters claiming to have 16 bit resolution do not in fact live up to this, because of non-linearities in the conversion procedure, critical tolerance of resistive components, accuracy of counters and so forth. For this reason, there is certainly at least an argument for having 16 bit converters that are truly accurate to 16 bits, as it is likely that the inadequacy of some earlier converters may have contributed to criticisms of digital audio sound quality. A second point to consider is that the traditional linear 16 bit converter suffers from increasing error as a proportion of total audio signal the lower the level of the signal is (since the quantising error is fixed at a maximum of ±0.5 LSB, and with high level signals this is relatively insignificant). This error manifests itself as quantising noise or distortion (depending on the nature of the signal, the distribution of the error in the frequency domain, and whether or not dither is employed), which gets greater as a proportion of the signal at lower signal levels, although in absolute terms it is still at a very low level.

Problems of low level resolution have sometimes been wrongly interpreted in the past and many false arguments have been perpetrated but, as with most false arguments there is some value in them because they are usually based on real evidence of a problem. Tony Griffiths of Decca, among others a great believer in the value of dither in digital audio, has demonstrated very convincingly on numerous occasions the effect of adding dither at low levels in digital audio equipment, showing that it turns a digital fade

◁ down to nothing from a dirty-sounding event near the bottom of the fader to a smooth fade into noise. Most people agree that although the 'dirtiness' only begins to appear at very low levels in relation to normal programme, it becomes important to get rid of it, especially in digital processing systems such as mixers, where signals may be changed in level, mixed with each other and otherwise processed. It is thus very important that the signal is quantised as accurately as possible at the input conversion stage before anything is done to it, in order that any error is not compounded. Certainly, if DSP and digital mixing equipment is to succeed, there is a need for more accurate conversion.

In many ways, whether misguidedly or not, the hi-fi world has tended to exist one step ahead of the professional audio world in the area of digital conversion. Compact disc players have sported D/A converters with specifications far higher than those on most professional equipment, using oversampling to 16 times and more, when the most that the latest professional systems use is either two or four times! Since it is the professional equipment that makes the CDs which the hi-fi enthusiast listens to, the whole issue becomes quite ironic, since the engineer balancing the original recording may be hearing lower quality sound than the domestic listener. It is only recently, though, that more attention has been paid to the quality of A/D conversion, this process being more complicated than D/A conversion. Attention to this area could be of potentially greater benefit to the consumer of high quality digital audio, as it is likely to affect the quality of what is recorded rather than what is replayed. Domestically, the consumer will be largely uninterested in A/D converters (unless he has a DAT machine) as domestic digital audio is generally 'replay-only', and so it is the professional who should consider improving his act in this area.

Oversampling in D/A conversion

Oversampling was first introduced to audio in the hi-fi arena to improve the perceived quality of CD replay, and has been used as a tool by Philips to gain effectively 16 bit resolution out of 14 bit converters. Used in a system such as CD, which operates at a nominal sampling rate of 44.1 kHz, the benefits of an oversampling D/A converter can only be gained by multiplying up the replay sampling rate using digital filters so that it is perhaps four or more times the nominal rate. One effect of this is to spread any noise or distortion over a wider frequency band, much of which is out of the audio frequency range, thus the noise within the audio baseband is reduced. Coupled with this is the advantage that any anti-aliasing filters need not have such steep cut-off slopes (since components of the sampled spectrum have now been moved to a much higher frequency and are unlikely to alias) and thus the audible artefacts associated with steep filters are avoided.

Some manufacturers of CD players have used oversampling to claim 18 bit resolution, even though 16 bit converters are used. In the case of an oversampling CD player, the multiplication of samples at the original sampling rate of 44.1 kHz with appropriate coefficients so as to create new samples at the higher sampling rate may result in samples of a longer word length than 16 bits. If a 16 bit D/A converter is to be used then any extra bits resulting from the multiplication process must be removed before the sample can be converted, and it is the intelligent truncation of the less significant bits that is one of the keys to improved sound quality. Simply to chop off any bits less significant than the 16th will negate much of the advantage that might have been gained, whereas intelligent rounding based on the values of less significant bits will preserve much of the information that had existed in the less significant bits in the on/off modulation pattern of the least significant bit of the 16 bit word, once the output of the converter has been averaged by lowpass filtering. It is part of the principle of information theory that as much information can be transferred in 4 bit words at a sampling rate of n, as can be transferred in 1 bit words at a sampling rate of $8n$. In other words, it is possible to use a high sampling rate and less bits per word, or a low sampling rate with more bits per word to transfer the same amount of information, and thus potentially the same audio quality.

Despite the improvement in sound quality available through the use of oversampling at the replay stage, it is important to realise that oversampling D/A conversion cannot magically extract information (audio quality) which was never there in the first place. It is now necessary to understand how oversampling applies to A/D conversion, as this will help to show how more information can be stored in recordings.

Increasing the recorded resolution

It should be stressed that there is no escaping the fact that you can't get something for nothing. "There is no such thing as a free lunch," to quote the cynical business man! Digital audio sound quality can only ever be as good as the worst device in the signal chain, and the quality is limited by the amount and accuracy of information transferred, whether in the number of bits per sample or the number of samples per second.

The *DCS-900* A/D converter that sparked off this discussion uses oversampling at 128 times the nominal rate of 44.1 kHz, which amounts to a little over 5 MHz. It also uses only 4 bit converters, which is the reason this feat is made possible. This bears out the theory that you don't need a 16 bit converter if you're going to sample at a high enough rate. In order that the output of such a converter can be recorded on a normal digital tape recorder operating at 44.1 or 48 kHz, the audio data sampled at the high rate (128 times nominal) must be converted to samples at the lower rate (effectively 1 times nominal) using a process known as decimation, and this results in samples at the nominal rate which have a word length much longer than 16 bits (the original information is now held in a greater number of bits per sample rather than in an increased number of samples per second). This is where the claim of resolution greater than 16 bits comes from, since the *DCS-900* produces samples from its digital output with up to 24 bits set (although, as DCS's Peter Gingold pointed out, he wouldn't like to guarantee the accuracy of bit 24!). The unit also has the capability to truncate the output intelligently to either 16, 18 or 20 bits.

In a trial session using the *DCS-900* linked via an AES/EBU interface to a Mitsubishi *X-86* 2-track machine, 20 bit samples were recorded on the tape. The 2-track ProDigi format has left space for this eventuality although the converters on the *X-86* itself do not operate to 20 bit resolution. Such an arrangement gives the engineer somewhat more freedom in level setting, as he now has a dynamic range of over 100 dB to play with, and he may be more free to allow a certain amount of 'headroom' for unexpected input peaks. The question is: 'What happens to the 20 bits when the recording is edited and mastered onto CD?'

It is possible that such a recording could be edited in a number of ways to preserve its 20 bit integrity. It could be splice-edited, or it could be copied to an editing system such as the DAR *SoundStation II*, which has 18 bit capability. Mitsubishi's electronic editor for the *X-86*, the *XE-2* could also be used. After editing the recording must somehow be truncated to 16 bits to be mastered onto compact disc and unless this is done with care, much of the advantage of improved conversion will be lost. If an engineer had taken advantage of the additional dynamic range of 20 bit recording, and intentionally allowed 10 dB or more of headroom (such that the recording peaked well below peak bits), truncation to 16 bits without gain correction would result in perhaps only a 14 bit recording. The correct approach would be to feed the recording through a digital mixer such as Neve's Digital Transfer Console *(DTC)*, raise the level such that the recording peaked near the maximum and then truncate to 16 bits.

There is therefore ground for suggesting that, in some cases, intelligent rounding to 16 bits at the A/D conversion stage, such as by using the switchable resolution provided in the *DCS-900*, might be more appropriate as a means of making a highly accurate 16 bit original recording to avoid the risk that somewhere further down the post-production chain the additional quality will simply be lost by equipment failing to recognise bits lower in significance than the 16th, or by people failing to realise the way in which the recording was made. In this way the engineer could be sure that the quality of his recording (even though limited to 16 bits) would remain intact throughout post-production without him having to keep track of its progress. Certainly, if a 20 bit tape is played back on the *X-86* via its analogue outputs, the four additional bits will not be recognised because it uses a 16 bit converter, and if the tape is copied to any other 16 bit format such as the Sony *1630* via AES/EBU it will also lose the four LSB's.

If 18 and 20 bit recordings are to become commonplace then there will be increased emphasis on the need for intelligent post-production and mastering, using equipment and people which understand about the technical implications of truncation, oversampling and so forth. One thing is clear, though, and that is that no matter how many bits per sample the original recording has, the consumer who buys the finished product will still only reap the benefit of 16 of them, at the ordinary CD sampling rate of 44.1 kHz. Because of this we would be back to the situation in which professional recordings can be of a higher quality than domestic equipment is capable of reproducing, it being part of the skill of the mastering engineer to relay as much of the master tape's quality to the domestic listener as possible.

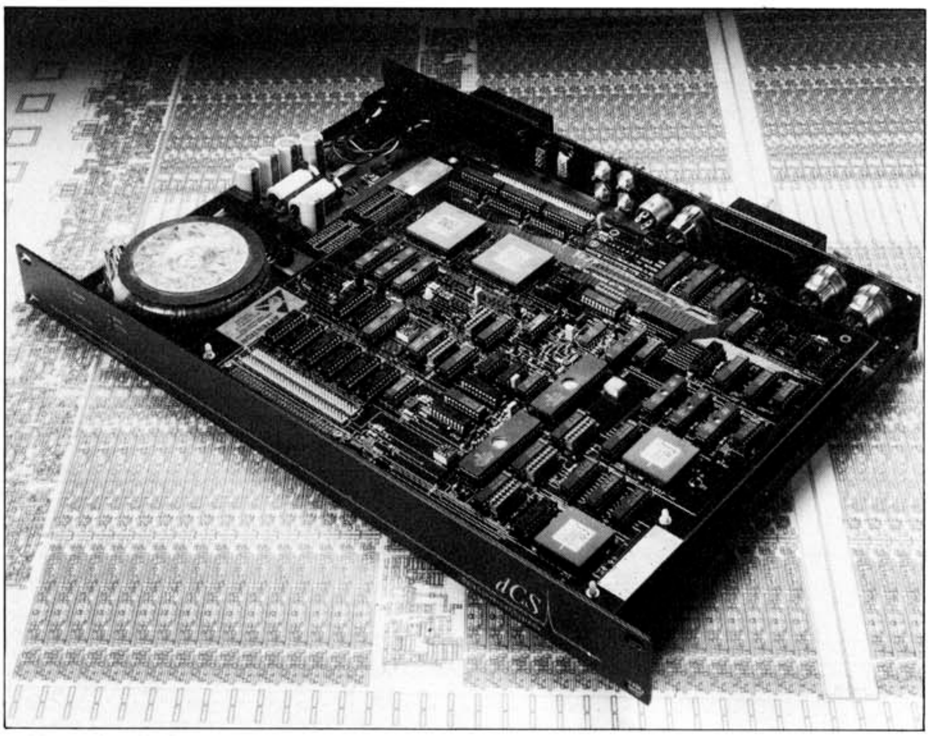

DCS-900 analogue to digital converter

Further benefits

What other benefits might there be in using a 'super converter' to make one's digital master recordings? Again, Peter Gingold of DCS, points to the fact that the additional bits produced at the digital output are not the whole key to the story. There are factors such as the accuracy of the converter over the whole dynamic range and the 'noise shaping' employed, which avoids the traditional effect of increasing distortion with falling level. In fact he indicates that although typical distortion components are at around −105 dB with respect to maximum modulation, at a single level of around −50 dB the distortion is in fact as low as −130 dB. This can be measured digitally using an FFT analyser. (It should be noted that DCS has indicated its intention to develop a D/A of similar quality to the *DCS-900*.)

Other advantages of high quality A/D conversion will be the improved phase linearity and freedom from ringing at the high frequency end of the audio spectrum, due to the oversampling technique described above.

David Ward of Mitsubishi emphasises the fact that high quality conversion will become increasingly important as audio signals are processed entirely in the digital domain. The more gain changes and equalisation that a digital audio recording experiences during post-production, the more chance there will be for low level errors to be compounded, and thus there is value in original material with as high a quality as possible.

For the future

The industry is without a doubt committed to 16 bit mastering, as this is the resolution of CD. It is hard to imagine anybody trying to change the CD format for the foreseeable future, unless they have a very hard hat, as its success is proven and millions of machines have been sold. It is just possible that someone may find a way of squeezing 20 bits onto a conventional CD while retaining its compatibility with existing 16 bit players but this is unlikely to find particular favour with the average consumer who thinks that the dynamic range of CDs is already too great for the average living room or car and possibly would rather see it compressed than expanded.

Converters such as the *DCS-900* will score their real points in returning to the audio engineer the ability to record audio of a higher quality than that which eventually needs to be reproduced in the home, leaving room for flexibility in processing and perhaps a margin for error. The engineer has been used to this situation with analogue recording and there are many who would like to return to the security of knowing that the material which they submit on digital tape is of a similarly superior standard. □