HPS-4000®

HIGH PERFORMANCE STEREO

Sound IS the Experience !™

DOLBY SR-D DIGITAL

BY

JOHN F. ALLEN

HIGH PERFORMANCE STEREO™

HPS-4000®

HIGH PERFORMANCE STEREO

FIRST IN DIGITAL STEREO

HIGH PERFORMANCE STEREO™

NEWTON, MA 02459 USA • TEL: 1-617-244-1737

DOLBY SR-D DIGITAL

by
JOHN F ALLEN

Though primarily known for their analog audio products, Dolby Laboratories has been working with digital sound for over ten years. Even while talk about digital movie sound increased, Dolby continued to concentrate on analog products such as Dolby SR. Within just a few years following their introduction of the Dolby Spectral Recording (SR) noise reduction system and its use with analog 35 MM stereo optical soundtracks, Optical Radiation Corporation and Eastman Kodak marketed the all digital soundtrack format, Cinema Digital Sound (CDS). While capable of delivering excellent sound on both 70 MM and 35 MM release prints, the CDS prints had no analog soundtrack. This required a dual inventory of special digital prints along with the standard analog prints.

The CDS system and other proposed digital motion picture formats presented a clear challenge to Dolby's market position. With the advantage of years of experience in the production and exhibition ends of the business as well as observing the difficulties that CDS encountered, Dolby created a list of what they thought the "ideal digital soundtrack" should have.

Most important, they believed that there must always be a normal analog soundtrack, to eliminate the need for a dual inventory. Impact on the printing labs should be kept to a minimum. The digital soundtrack itself must be durable enough to last for the life of the print. And of course, the sound quality must be equal or superior to 70 MM magnetic prints.

Dolby certainly had the right idea. Simply announcing a single inventory digital print format was enough to put CDS out of business two years before the first production SR-D processors were to arrive.

To understand Dolby's digital soundtrack, it is best to break it down into two components: The first half is the Dolby AC-3 transform, a low bit digital coding scheme. The second half is the SR-D soundtrack format itself, which is a storage and retrieval implementation for the AC-3 coded digital data.

AC-3

A brief reference to the familiar digital compact disc reveals a relatively simple digital

audio storage system. First, there are only two channels. Each channel has a fixed data storage capacity. This means that if one channel is silent while the other contains a crash of thunder, the silent channel's data capacity will sit unused.

The CD's digital sampling rate is 44,100 per second and the resolution is 16 bits. So for each second, the signal is "chopped up" into 44,100 samples of audio. Each of these samples is then measured, or "quantized". The measured signal is stored as a code of 16 ones and zeros or bits.

Therefore, for every second, the CD stores about 1 1/2 million bits of digital audio coding. For each hour, the two channels require over five billion bits. A two hour 6-track motion picture soundtrack would require over 25 billion bits of storage capacity, if the CD format was used. However, this format is generally not used in professional systems.

Professional digital products usually employ a sampling rate of 48,000 and 18 bit resolution. Therefore a two hour "5.1" channel soundtrack (five channels plus a subwoofer) in the professional format would require a storage capacity of over 31 billion bits on the film. This works out to about 180,000 bits per frame. Such density requires that the tiny spots on the film that will represent the data must be extremely small. The smaller the spot size, the more difficult and expensive the printing and reading becomes. Less data would allow for a larger data spot size, making the printing and reading easier and less costly.

SPOT SIZE

There are thus two conflicting issues: A large and therefore easy to print and read spot size and achieving good quality multi-channel sound with an extremely limited data capacity.

After a careful study of print wear rates and distribution patterns, Dolby decided to place their digital data in between the sprocket holes on the soundtrack side of the film. The size which a spot in this area had to be in order to be reliably read (through scratches) after 1000 plays was felt to be about the size of an analog soundtrack "wiggle" corresponding to 6,000 Hertz, which is well within high speed printing capabilities.

The number of spots of this size which could fit within the .1 by .1 inch area between two perforations became 5776. Thus the data block would be a square of bits, 76 by 76, which provides a data capacity of 23,104 bits per frame. So how does one accurately represent and store 180,000 bits of digital audio with a capacity of only 23,104?

DATA REDUCTION

Several digital audio data reduction schemes have been developed. The problem is that such schemes can diminish the sound quality. To avoid this, any audio information which is to be deleted must represent inaudible sound. If sounds which are too soft to be heard in the presence of louder ones are to be left unrecorded, it must be well understood what will constitute and inaudible sound and under what conditions this occurs. So any digital sound data reduction system must select which data is important and which is not.

Dolby's approach is called AC-3, as it is the third generation of their low bit digital audio coders. The six channels of audio are each converted to 48,000 sample, 18 bit digital data. At this point, the data is in its simplest form, much like a six channel CD format might be. But rather than store all these bits, AC-3 performs several mathematical operations and extensive data analysis, all aimed at the goal of an inaudible data reduction of about 13 to 1. Most of the data reduction is accomplished by the mathematical transformation of the digital data into a code requiring fewer bits. The deletion of sounds which are masked by louder events is said to be used very little and with little impact.

THE BIT POOL

On average, a single audio channel can be represented by 128,000 bits per second with the AC-3 system. Rather than treating each channel as a fixed capacity of data storage, AC-3 treats all the channels as one channel for the purposes of data storage. This means that any channel's momentary unused capacity would be utilized by other channel's as needed. This is a much more efficient storage system and reduces the overall capacity required.

According to Dolby, it turns out that the average bit demand for multiple channels is roughly proportional to the square root of the total number of channels. In other words, the more channels you have, the less bit capacity per channel is required.

5.1 channels would, on average, need 128,000 times the square root of 5.1, or 289,000 bits per second. This is well below the 554,496 bits per second available on the film. The remaining bits are used for data redundancy to aid in error correction, and decoding information required by the decoder to reconstruct the transformed digital data information back to the standard 18 bit digital data, for conversion to analog audio and actual playback. The extra bits are also needed to code high density situations such as something very loud in all the channels at once.

THE SOUND NEGATIVE

Whether in digital or analog form, when the final discrete 6-track, and the matrixed left-total right-total two track SR soundtrack print masters are complete, both are used at the same time to create a single composite analog-digital sound negative. This will allow distributors to print and release all prints in one sound format. Up to now, this has not occurred. Only a few Dolby SR-D digital prints have been struck for each release. Dolby predicts that as the labs gain experience and the remaining printing difficulties are conquered, larger numbers of SR-D prints will be available and that at some point, all prints will indeed be in the same format.

DECODING SR-D

Reading and processing the SR-D soundtrack in the theatre is a rather amazing story by itself. Each second, 96 blocks of digital data fly past the reader. Of course, this also means that the data is read among 96 interruptions, as the perforations also pass by. All 5776 spots in each block mean something. But their meaning, such as channel assignments and other instructions, is constantly changing. Turning all this into continuous multi-channel sound, and doing so correctly, is an impressive feat.

The data block is read by a 512 element charged coupled device (CCD), which is, in essence, a television camera. The output of the CCD is an analog video "picture" of the spots on the film. This is converted into the digital video domain at a rate which tracks the speed of the film.

This video data is scanned for synchronization information and to determine where the four corners of the data block are located. If two of the corners are known, the position of the entire block is known and the proper location of the data bits is known. Determining the whether or not a real data spot exists and preventing false data from entering the process is best understood by imagining a piece of window screen being placed directly over the data block. The corner bits verify the proper Position for the screen. Looking at the "screen", one sees in the small square openings exactly where the spots should be.

Using a digital version of such a screen, the data spots are identified and then sent to a "thresholding" stage for further enhancement. Based on the average density of the spots, a threshold value is determined, above which the spot is recognized as a "1" and below which, a "0". This stage reduces the possibility that a scratch could be mistaken for actual data and establishes what, on the film, is a spot and what is not.

Such a critical and clever operation provides such cleaned up data that the next stage of error correction is considerably eased. During the error correction process, the data blocks

are also combined into a more continuous data stream.

The corrected data is still a long way from becoming audio. For one thing, the speed of different projectors varies and the digital decoding clock must be adjusted accordingly. To do this, the digital data is fed into a first-in first-out buffer, where it is smoothed into still a more continuous data stream. The decoding clock rate is determined by regulating the fullness of the buffer and adjusting the output rate to equal the input rate which changes with the speed of the projector.

FINAL DECODING

We now have digital data that we can work with in the sense that the vagaries of a projector and data retrieval have been dealt with. At this point, the data is delayed by an amount set by the installer to synchronize the sound with the picture. But it is only after all this that the actual audio is decoded.

The delayed digital data is passed onto an AC-3 transform decoder. Some of the data is for the sound and the rest instructs the decoder on what the data reduction encoder did, so the standard 18 bit digital data can be reconstructed.

The data representing all six channels has, up to now, been treated as one channel. Once it has been restored to 18 bit data words, the data is separated into the individual left, center, right, subwoofer, left surround and right surround channels. Each 18 bit channel is now converted back to analog audio and sent to the cinema processor for playback.

AUTOMATIC UPDATES

One additional feature of the SR-D system is important to note. As improvements in the AC-3 process are developed, field units are automatically updated by data stored in the digital soundtrack data. When such an update occurs, the film's first pass through the projector will play in analog for about 20 seconds while the new software is being loaded into the digital processor. The system then switches to digital. The remainder of the first run and all subsequent runs (from the beginning) will be played in the updated digital format.

SOUND QUALITY

Placing digital soundtrack information on the film is doing it the hard way, but it does allow digital and analog shows to be built up and treated in the identical manner. Considering all the processing that Dolby has employed to make SR-D as reliable as possible, it's amazing (to me at least) that the sound can be any good at all, let alone as superb as it is.

In my opinion, this industry desperately needs the quality of digital sound to attract audiences into theatres. Properly played back in a modern theatre, this is the kind of sound that "takes you there", transporting you to the scene, wherever and whatever it may be.

John F. Allen is the founder and president of High Performance Stereo in Newton, Mass. He is also the inventor of the HPS-4000® cinema sound system and in 1984 was the first to bring digital sound to the cinema. John Allen can be reached by E-mail at JohnFAllen@aol.com.