

DVD

AUDIO REALITY

The DVD-Audio disc specification is at last

finalized, and the DVD-Audio Format Book is due to be published in the spring.

RICHARD ELEN
looks at the specifications, challenges and benefits that the next audio distribution format will bring.

NEXT GENERATION AUDIO

It has been some time in the making, but the next generation of consumer audio distribution formats is finally with us — or at least it will be in just a short time. The DVD Forum, an ad hoc association of manufacturers and other interested parties, announced in early February that its Steering Committee had approved Version 1.0 of the DVD-Audio Disc specification, making it the fifth major format in the DVD hierarchy after DVD-Video, DVD-ROM, DVD-RAM and DVD-R. The DVD-Audio Format Book will be published in the near future.

Working Group

The DVD Forum's Working Group 4 was formed, in December 1995, to formulate specifications for the DVD-Audio format. It began deliberations in 1996 under the chairmanship of JVC, and the initial draft specification, Version 0.9, was released in June 1998. The final version has been expected for a while; it was originally slotted for release last fall at a

conference in San Francisco shortly after the AES Convention. The majority of the hold-ups have been due less to technical format matters than to copy protection, copyright management and anti-piracy issues.

There have been some changes along the way, however. The music industry's International Steering Committee requested that a number of specific requirements be included, one of which was the addition of lossless compression. It had become clear that the original goal of over an hour of multi-channel, high-definition music on a single-sided disc was going to be very difficult to achieve. The reason was lack of space on the disc (which has about seven times the capacity of a CD) due to multi-channel capability and the use of longer word lengths and higher sample rates (such as 24/96). Compression was going to be necessary to meet the requirement of a similar length of performance to that of CD, but at higher resolution and in multi-channel format.

We have been using compression on DVD since the beginning, of course, in the form of Dolby Digital (AC-3) and latterly DTS. Both, however, are 'lossy' compression systems in which psychoacoustic research is utilized to reduce the data rate — thus they are also known as 'perceptual coding' systems. Putting it bluntly, what you do is remove data in the areas where its removal is least audible (ie. hopefully not at all). For the next generation of audio disc, the industry — perhaps remembering other compression schemes in the past that were all too audible — rightly insisted that lossy compression was insufficient. If compression was going to be used, it had to be lossless.

It's likely that the only reason we ever got into lossy compression in the first place was that it was the best we could do at the time. Lossless compression simply needed too much processing power to be affordable. But, despite the undoubted quality of the results of the two leading systems, it seemed incongruous to go to all the trouble of using high-density linear PCM and then squashing it with a lossy technique, however effective. In addition, thanks to the rapid development of digital processing technology, by the time lossless compression was needed, it was possible.

MLP — The Heart Of DVD Audio

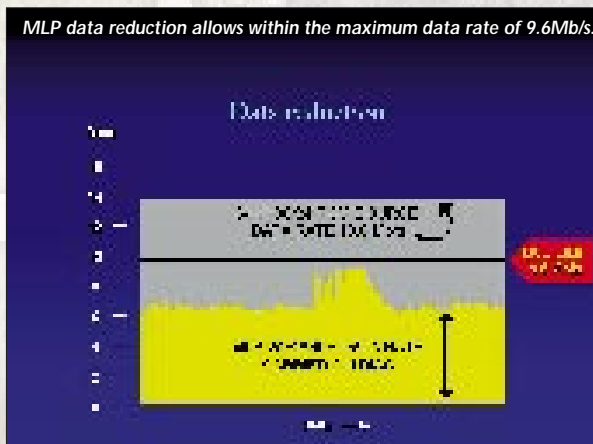
A major breakthrough offered by DVD-Audio is in the form of MLP — Meridian Lossless Packing. Like ZIP or Stuffit on your computer, MLP takes a PCM data stream and 'packs' it at one end of the chain, 'unpacking' it at the other to provide a completely accurate replica of the original (this includes accuracy of sample-timing too, by the way).

The technique was chosen in a shoot-out in Japan last year between four competing lossless systems. Initially, the focus fell on the encoding side of the equation. Here, MLP is fairly comprehensive, and it requires a respectable amount of processing power. Other systems required less power in this area. However, when the decode side was considered, MLP was the hands-down winner: it is much simpler to process. This is because most systems treat the unpacking and the extraction of the PCM data as two separate operations; first you do one, then the other. MLP, on the other hand, treats both as one operation. The result: a simple-to-decode system that requires very little expense in the player — where it really matters. Indeed, even the encode side is not prohibitively power-hungry. If you have an existing authoring suite with an AC-3 encoder based around Dolby's Onyx chip-set, you'll be able to run the MLP encoder on it.


MLP was developed by Robert Stuart of Meridian, Peter Craven, and revered audio genius, the late Michael Gerzon. Bob Stuart chaired the Acoustic Renaissance in Audio group, which had been proposing a very sensible set of ideas for DVD-Audio for some time, including the need for lossless compression. Gerzon and Craven — who often cooperated on digital audio projects — were behind the ARA on the technical side from the beginning. Gerzon will be remembered by many readers as one of the driving forces behind the development of the Ambisonic surround-sound system in Britain in the 1970s.

The MLP development team realized that a 24/96 PCM allocated a significant amount of space to recording very little information. According to the MLP preliminary data sheet, modern, high-rate formats are capable of "carrying more information than is strictly necessary for the human listener, or than is available from modern microphone and

converter techniques. This means that such audio streams contain true redundancy... MLP is an audio coding scheme that discovers the redundancy and packs the audio into a smaller space — but in such a way that a simple decoder can recover the original signal exactly, bit-for-bit." For more details on Stuart's research on this topic, see his paper, *Coding High Quality Digital Audio*, available from the ARA website (www.meridian-audio.com/ara/).



In an admirable political move, Stuart made a deal with Dolby Laboratories to act as the licensor for the system: in a single stroke, he insured the cooperation of a company with a major stake in multi-channel audio and compression, and placed both companies on the same side. At the announcement last June, Dolby's Director of Technology Strategy, Roger Dressler, stated: "We see the licensing of MLP as a perfect fit with our on-going licensing program. MLP not only has a role in the DVD format alongside Dolby Digital, but the system's unique capabilities will round-out our portfolio of audio coding solutions for a wide range of uses."



Data transmitted with MLP is only the remaining 'music' plus predictor parameters.

In a very interesting way, MLP was initially presented to the public this past June at the Hi-Fi '98 show in Los Angeles. Stuart replayed some very unusual and impressive CDs with the system, including some encoded with MLP in Ambisonic B-Format. As far as the multi-channel capabilities of MLP are concerned, the official DVD-Audio spec discusses up to eight audio channels. However, the MLP data sheet (available on Meridian's website) lists 'up to 64...' with 'flags for speaker feed identification, [and] flags for hierarchical



► feeds (eg. M&S, Ambisonic B-format and others). B-Format is a particularly elegant way of transferring surround audio information, as it uses just four channels to transfer everything about a real or imaginary acoustic environment — including height. With a decoder at the receiving end, you can drive just about as many speakers as you might wish.

The data sheet also suggests several applications besides DVD-Audio, including 'three or four channels on CD; two channels

20-/24-bit on CD; [and] 88.2kHz two-channel on CD'. It will be interesting to see whether anyone actually does this commercially. If they do, Bob Stuart's own remarkable range of Meridian home theater systems will play them back — his units already offer MLP decoding.

A Look Inside MLP

Meridian also gives some of the qualitative features of MLP, and not all have direct or obvious applications in the realm of DVD-Audio as we might understand it.

What they mean, however, is that DVD-Audio as a multi-channel audio distribution format can be made virtually open-ended; it's able to cope with yet-to-be-developed surround audio systems. Some of the features listed by Meridian include:

- Longer playing time than allowed by LPCM (essential to meet ISC requirements for DVD-Audio).
- Higher quality by delivering more channels, or bits, for the same playing time.

- Multi-channel audio at 192kHz.
- Guaranteed quality: the Lossless Decoder delivers bit-accurate data.
- High-quality, mix-down options: longer playing time with multi-channel material.
- Options to deliver artistically correct two-channel alongside multi-channel with fewer penalties on playing time or data rate.
- Fine control over delivered quality and playing time.
- Easy prediction of playing time as a result of constant packed data rate.
- Additional data channel in the stream to carry copyright and signature information.
- The large reduction in the audio data rate means that many more options for audio with pictures are possible.

Exactly how the flags for hierarchical encoding are implemented has not been revealed, but one hopes that they are more like MIDI IDs and less like DVD-Video flags. The latter are essentially individual bits set 'on' or 'off' to indicate a function. With MLP, you would run out of these fairly quickly. A MIDI ID-like flag system, however, would allocate an ID code bit-pattern to each registered system. If you came up with a new surround scheme and you wanted decoders to be aware of it in the future, you would ask for an ID to be allocated by the equivalent of the MIDI Manufacturers' Association. We might find this useful, for example, in incorporating height information into surround recordings.

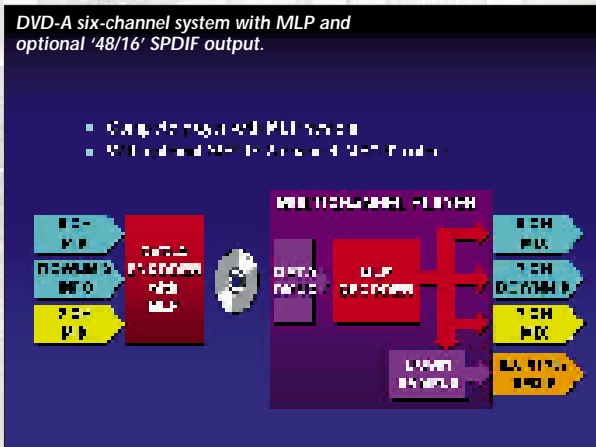
With MLP offering such an effective solution to the compression problem, it was only reasonable that, to meet ISC requirements, MLP be made the mandatory compression scheme for DVD-Audio. You can use uncompressed linear PCM instead if you wish, but if you need to use compression to get the program on the DVD-Audio part of the disc, MLP is what you will be using.

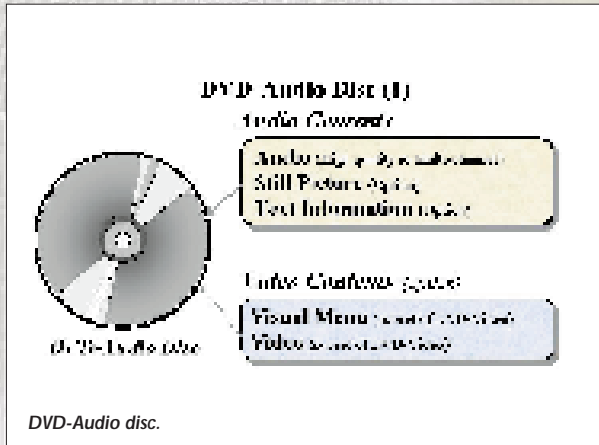
Digital Audio Implementation

Although DVD-Video is primarily used in 24/48 format to allow 5.1 transmission, DVD-Video already — theoretically — allows for four channels of PCM at 96kHz sampling and eight channels at 48kHz, the maximum bit-rate being 6.144Mbps. However, while you can do four channels of 96kHz sampling, you can only do it if you are happy with 16-bit word lengths and, in most peoples' view (including the present author), increasing the word length beyond 16 bits is more important than doubling the sample rate. (Only stereo at 24/96 is practical on DVD-Video, and note that Bob Stuart, in the paper referenced above, has good reason to believe that not only is 96kHz excessive but so is 24 bits.) The whole scheme goes like this:

PCM options for DVD-Video:

- 16/48 — up to eight channels
- 20/48 — up to six channels





remembered that, had we chosen sample rates that were not simple multiples of our current ones, even if they were more efficient, we would have possibly condemned ourselves to a life of unsatisfactory, non-integral, sample rate conversion. One big benefit of converting between 96 and 48kHz is that it does not require a perfect 200+ pole digital filter, so it may even sound decent.

In passing, one more thing is worth noting on the subject of very high sample rates. If 96kHz is excessive, what about 192? My view, until recently, was that it was 'completely absurd'. However, it has been brought to my attention that recent research noted by Sonic Solutions indicates a different conclusion. If you play a subject the same signal to each ear and vary only the time delay between the signals reaching the two ears, most people can say which signal arrived first, down to delays of a few tens of microseconds. This degree of accuracy is in the same order of magnitude as the delay between samples at 96kHz. Apparently, some people can hear

- 24/48 — up to five channels
- 16/96 — up to four channels
- 20/96 — up to three channels
- 24/96 — up to two channels

Dolby Digital on DVD-Video may deliver a maximum bit-rate of 448kbps, with six channels sampled at 48kHz. The other alternatives are as follows: MPEG-1, two channels of 48kHz sampling for a bit rate of 384kbps; and MPEG-2 multi-channel, eight channels of 48kHz for a bit rate of up to 912kbps.

When we get to DVD-Audio, we find that, although the same technology is in use, the absence of so much significant video content allows the maximum bit rate to reach 9.6Mbps. Fundamental to the specification is the ISC requirement that a disc be capable of delivering 74 minutes of material, using LPCM or lossless compression (the latter probably being necessary to deliver multi-channel program at this kind of length).

Players will support a hybrid CD/DVD-A disc, meaning that if you want to release a disc that will play on both a regular CD player and a DVD-Audio player (like a Super-Audio CD), you can. In addition, DVD-A includes the following capabilities:

- Scalable linear PCM multi-channel audio — 48/96kHz, 44.1/88.2kHz and 16/20/24 bits; six channels maximum.
- Super high-quality linear PCM audio — 192/176.4kHz and 16/20/24 bits; two channels maximum.
- System Managed Audio Resource Technique (SMART) down-mixing for stereo presentation of multi-channel contents in two-channel form. Each of the six channels may be mixed down to stereo by means of coefficients that can be set on a track-by-track basis. Coefficients determine level, panning and polarity.

You can also offer dual-side or multi-layer discs, so, for example, you could release a disc with 74 minutes of six-channel, MLP-compressed 24/96 audio on one side/layer, and 66 minutes of stereo 24/192 on the other. Or the 24/192 stereo mix could also be MLP-coded, giving a length up to about 86 minutes.

It may well be that Stuart is right about 24/96 being excessive. However, it should be



- ▶ 15 or even five-microsecond differences, which is way up in the 192kHz region as far as sample rates are concerned. This research indicates that, while high sample rates are not necessary for frequency response purposes, they may have a bearing on the accuracy of stereo, and thus also surround, localization.

You can produce a disc that will meet these specs and play back only on a special DVD-Audio player, or on one that plays both DVD-Audio and DVD-Video discs (called a 'Universal' player). Unfortunately, these do not actually exist yet. In addition, as Dolby has gone to great pains to point out, people in their right minds with enough space left on the disc after they've laid down the super-quality DVD-Audio section will take the opportunity to put down an AC-3 stream as well. This will mean that the disc can also be played on a DVD-Video player. Since there is a growing number of these already out there, this makes good sense. This type of disc, officially called a 'DVD-AudioV' disc, is likely to become the single, most popular type of DVD-Audio disc, because it will play back on any DVD player and also allows the inclusion of video clips.

The only problem, it appeared until recently, is that there might not be room to do it. To ensure that the most basic DVD-Video player would play back a

DVD-AudioV disc, it was originally thought that all such discs should have plain, ordinary, uncompressed linear 48kHz PCM on them, at least in stereo, as part of the DVD-Video content. At first this seems logical enough, since this is a requirement of DVD-Video discs so that they will play back on the most basic DVD-Video player. Unfortunately, this might take up as much as 700MB on the disc, possibly severely limiting the potential of the DVD-Audio part of the disc. A last-minute alteration in the specification, announced at the turn of the year, solved the dilemma. The mandatory requirement for PCM in the DVD-Video part of the disc was withdrawn (there will be a PCM version in the DVD-Audio portion anyway). Now the only mandatory audio format in the DVD-Video area is Dolby AC-3, although you can use PCM if you wish.

No doubt this pleases Dolby, who now has a handle on the two major compression technologies mandated for DVD-Audio. It possibly upsets DTS, however, even more than before. They were already complaining about being given equal billing with Dolby — more than they deserve in some people's estimation — instead of equal billing with MLP. Of course, they couldn't have equal billing with MLP because the ISC required lossless compression and, although

DTS has a little-known lossless mode, in the decision-making process they were not the ones chosen, because MLP offers so much more.

At this point, however, let us spare a thought for the poor people who want to give the listener the best of all possible worlds: a single-inventory disc, holding on it everything you could possibly want. This disc would probably include the following: a high-density DVD-Audio stream of up to six channels at 24/96 with MLP, maybe with a stereo LPCM stream; a simple sample-rate converted AC-3 5.1 version at 48kHz on the DVD-Video portion; and a Red Book CD layer at 16/44.1 so you don't even need to release a separate CD version. This sounds very nice, doesn't it? Unfortunately, if you are going to get the Red Book version from the same master, you must do a very nasty bit of sample-rate conversion. So maybe that won't happen. After all, even if you did a Red Book-compatible layer, the disc would be way more expensive than a regular CD. Thus, plenty of would-be listeners on CD wouldn't buy it anyway — but that's economics, not technology.

Getting back to our discussion of the DVD-Audio specifications, many will be gratified to learn that there is an enormous amount of flexibility built in for the content producer. Let's take a look at some of the options.

Producer's Choices

The DVD-Audio specification happily leaves it up to the content producer to determine a number of very significant factors. Here are the main ones:

Whether a center-front channel (hard center) or two front speakers only (virtual center) are used. Plenty of engineers hate a center-front channel and don't mix to it. Certainly, in many cases, it is hard to get a center channel to 'gel' in a surround-sound-for-music ('SSfM') mix. Brought up on stereo, as we were, using a center channel (which, after all, is only there so you can hear the dialog properly in a movie) is also fairly alien to most engineers, who prefer the smoothness of a front stage with a virtual ('soft') center. This is particularly true for recording techniques that don't rely solely on level-based localization, such as coincident pairs of mics. Thank goodness; if you don't want to use the CF channel, you don't have to. (I would suggest you make a note on the box, though, so that people don't send back your discs complaining that the CF channel is broken.)

Whether the sample rates and bit depths used for front and rear are the same or different. For example, you could use 24/96 at the front and 16/48 at the rear. The sample rates must be in the same 'family' (ie. related by a simple multiplier, such as 48/96/192kHz or 44.1/88.2/176.4kHz). I would recommend that people use the longest word lengths they can and avoid weird hybrids like 16/96. There probably isn't much point.

Whether a separate, two-channel mix or multi-channel fold-down are provided. If a separate stereo mix is provided, it is used automatically instead of folding down the multi-channel mix. This is important. Even though the fold-down parameters are far more flexible than they are for DVD-Video, there are still plenty of producers who would rather do the stereo mix themselves, rather than allow some set of down-mix coefficients in a machine do it for them, even if they define the coefficients. Well, if you have the room for a separate stereo mix, you can do it.

Whether lossless compression is used for longer playing time, or only uncompressed LPCM. Either MLP or LPCM are mandatory on the disc (but see the DVD-AudioV discussion above).

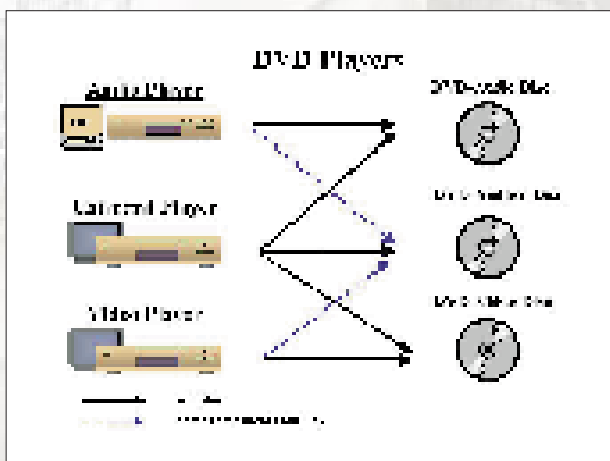
Whether LPCM is used or a lossy compression technique such as DTS or AC-3 (see above). DTS has equal billing with Dolby on the subject of optional perceptual coding. You have to have an LPCM stream somewhere, either uncompressed or with MLP but, apart from that, you can use either AC-3 or DTS as an optional method.

Value Added Options. Producers can optionally display liner notes, discographies

and other information, as well as video clips and slide shows. The facility exists to provide hot-links from the content to URLs on the web, for accessing bonus contents if the disc is used in a web-capable environment.

What Will DVD Audio Players Offer?

There are several areas in which the player has more strictly defined requirements than those placed on the producer — so that, whatever a producer decides to do within the guidelines above, they can rest assured that any player will play it. SMART



downmixing, for example, is optional for the producer, but mandatory in the player.

DVD two-channel LPCM is the minimum playback requirement, and DVD-Audio is to be compatible with DVD-Video (as discussed above) and DVD-ROM. DVD-Audio players will play back Red Book CDs. Players will also support a hybrid dual-layer Red Book/DVD-Audio disc, while this is optional for producers.

There is something else that players should offer, too. At the moment we are seeing major manufacturers gearing-up for another format war. This time it's DVD-Audio on one side and Sony/Philips' Super-Audio CD on the other, with its high sample-rate DSD bit-stream technology. The actual disc technology involved is identical. It's just that one format is based on PCM and the other on bit-stream technology. Theoretically, the DVD-Audio spec provides for something almost indistinguishable from an SACD, but it was only put there for political reasons (and Sony/Philips went off on their own anyway).

Such a format war involves expensive production equipment, and it also involves incompatible consumer gear. Once again, both ends of the chain risk getting tied up in a wasteful format war. Remember VHS and Betamax? Or DASH and Pro-Digi? Arguably, we would all have open-reel, digital machines today if there had been just one format, not two.

Part of the solution is a simple one that at least protects the consumer. The DVD-Audio spec, as noted above, supports an SACD-like disc as an optional format. The solution is to keep SACD support optional for the producer, but mandatory for the player. The technologies are pretty similar anyway, but there will need to be some low-cost changes on the D/A front. This will mean that the consumer can buy any DVD-Audio player and know that it will also play SACD discs.

Rumor has it that the reverse is already the case: the 11 manufacturers currently lined up behind SACD are committed to ensuring that their players also play DVD-Audio discs. If so, DVD-Audio player manufacturers will be obliged to follow suit, or they won't sell anything.

This still leaves an open question over the technology in the studio. My guess is that several producers who specialize in high-end classical recording and other material, which has a relatively simple signal path — mics, pre-amps, converters, storage, editing, master — will opt for DSD recording. Those who prefer to record multitrack and perform a great deal of mixing and DSP work may prefer to stick with PCM. There are still very few signal processors for DSD, and to perform DSP operations you must turn the signal into something very much like PCM and then back again. So, many will ask, why bother?

Something that initial DVD-Audio players will not do is to give you a high-density digital output. It looks as though, for copy-protection reasons, the first generation of DVD-Audio players will deliver 'CD-quality' digital outputs only, just as today's DVD-Video players, with 24/96 converters on board, only output 48kHz digital. However, the second generation of players is expected to provide a FireWire digital output including the entire MLP stream for external processing (eg. external digital decoders for B-Format discs) and possibly the video too. By that time, it is expected that all of the copy protection issues will have been sorted out.

In the meantime, DVD-Audio offers an exciting future for the production facility and for the consumer. For those of us who have been complaining about the poor quality of 16/44.1 CDs, here's the answer. □

Richard Elen is a frequent writer on professional audio and can be contacted at relen@igc.org. He is Vice President of Marketing at Apogee Electronics in California. Opinions expressed in this article are the author's own and not necessarily those of his employer.