Audio Loudness and Dynamic Range in the Compressed Digital Environment

Craig C. Todd
Dolby Laboratories

Abstract

Digital audio compression systems are capable of delivering audio into the home or the headend with well over 100 dB of dynamic range. Will this potential dynamic range be exploited or squandered? This paper shows how features provided by the AC-3 audio compression standard allow a single encoded bit stream to supply an appropriate dynamic range to all listeners.

INTRODUCTION

Current NTSC broadcast practice is often to provide audio suitable for the worst case listening environment (audio received via demodulation of RF and reproduced by tv with a built in loudspeaker). The audio is highly compressed, and the level of the highly compressed speech is such that peaks often hit 100% modulation. There is little headroom for left for music and sound effects to make a dramatic impact. The situation is somewhat better for feature films delivered by premium cable movie services, where less compression is applied, and some headroom above dialogue level does exist for music and effects. However, in order to leave some headroom the average level of dialogue must be reduced compared with that of broadcast channels and this leads to some consumer dissatisfaction when level variations are encountered while channel hopping.

Future delivery systems will incorporate digital audio compression (bit rate reduction) systems which can provide very wide dynamic range audio (>100 dB). These new systems must be interoperable with current broadcast systems where common speech often reaches 100% modulation, and must be able to interface with a conventional TV set by means of a Ch3 RF signal. This leads to a quandary where either the newly available >100 dB of dynamic range will be squandered so that speech can be heavily modulated to match current broadcast practice (to achieve reasonable level matching on Ch3 while avoiding the potential of overmodulation), or the new services will leave an appropriate headroom for music and effects above the level of speech in (which case they will not match the level of speech of off-air broadcast stations).

Consideration of these practical problems has led to the development of solutions. The AC-3 audio compression standard[1,2] provides a number of syntactical elements in the encoded bit stream which provide practical solutions to many of the problems of audio delivery. Programs may be encoded with differing amounts of headroom and still reproduce at a consistent level. A Ch3 remod output may be provided which is comparable (in dynamic range and loudness) to current broadcast practice. A baseband audio output may be provided with a dynamic range which is somewhat compressed (suitable for most listeners). The listener may be provided control to optionally reduce or eliminate the amount of compression which was applied by the program provider. In this case the listener (perhaps a home theatre enthusiast) can choose to listen to the original soundtrack free of any dynamic range signal processing. In short, a common bit stream can be used to supply audio service to wide range of listeners with differing needs, without severely compromising the audio delivered to any particular group of listeners.

This paper will go into some detail on these practical problems, and how the features of the AC-3 audio coding standard may be used to solve them. Some product design guidelines will be provided.

AUDIO PRODUCTION

This section will review production practice followed in the creation of the highest quality soundtracks - those for major motion pictures.

Current film sound practice is to produce a multichannel soundtrack in the so called 5.1 channel format which is the subject of an ITU-R Recommendation [3]. This format provides five full bandwidth channels for left, center, right, left surround, and right surround. An additional limited bandwidth channel (the 0.1 channel) is available for high level low frequency effects (LFE). The LFE channel is limited to 120 Hz, and is adjusted to reproduce 10 dB louder than one of the full bandwidth channels. Reproduction of the LFE channel, in the context of the home environment, is considered optional.

During audio production, audio recorders are set up with a *reference level*. There is a match between this reference level and a sound pressure level (SPL) in the mixing studio of 85 dB SPL. That is, bandwidth limited pink noise, recorded at reference level, will measure an SPL of 85 dB on a sound level meter placed at the location where the mixing engineer sits. The sound recorders, depending on their type, typically have 20 dB of headroom available above reference level. Analog recorders using Dolby SR noise reduction have a soft overload characteristic and typically provide slightly more than 20 dB of headroom. Digital recorders, with their inherent hard overload characteristic, typically have exactly 20 dB of headroom.

The overload point of the audio recorders establishes a maximum SPL level for each individual channel. This level is (assuming 20 dB headroom above reference level) 105 dB SPL for each individual full bandwidth channel, and 115 dB SPL for the LFE channel. Assuming power addition in the room, all main channels driven to maximum level would result in an SPL in the room of 112 dB.

Film Release Formats

The 70mm magnetic 6-track film format, using Dolby A noise reduction, can essentially deliver a film sound production to a theatre without compromise. (Actually there is some compromise at the frequency extremes where headroom on the 70mm print is reduced.)

The 35mm stereo-optical film release format has severe constraints and cannot deliver the full soundtrack to a theatre. This format only delivers two discrete channels of sound. A 4-2-4 matrix system is employed to produce the illusion of four channels (left, center, right, and surround). Headroom above reference level is only about 6-8 dB when Dolby A noise reduction is employed, or about 9-11 dB when Dolby SR is used. A substantial amount of peak limiting must be applied to make the original 5.1 channel sound production fit into the 35mm optical sound track. First, each individual track is capable of overloading the 35mm format. Second, when the tracks are mixed down (and matrix encoded) to the final matrix encoded two channel LtRt signal, there is gain in the potential peak level due to the fact that channels are being combined. In the theoretical worst case, about 20 dB of limiting can be required to produce the 35mm version of the LtRt from the discrete 5.1 channel master. In typical operating practice the worst case signals do not occur and only about 10 dB of peak limiting is actually required.

The AC-3 digital audio compression system was designed to encode the full dynamic range of a 5.1 channel film soundtrack into a data rate of 320 kb/s (although occasional peak limiting is necessary when encoding a soundtrack which has signal peaks in excess of +20 dB re reference level.) The Dolby SR• D digital film system was designed to place this data optically onto a 35 mm film print, and allow the full sound field to be decoded and reproduced in the theatre.

Home Video Formats

Video media differ in their sound capability, although today they are all limited to two channels. The consumer may choose to employ a matrix surround decoder in order to decode the delivered LtRt two channel signal into four channels.

When film soundtracks are provided for video media, the amount of limiting applied when the 2-channel matrix encoded LtRt signal is produced varies. In some cases the same version that was used for the 35mm film release is reused. In other cases, the transfer may be redone with a different amount of high level limiting, and perhaps some low level compression which brings up the loudness of the quietest parts of the soundtrack. The techniques employed depend on the intended purpose of the soundtrack. Is it intended for broadcast television? For a premium cable movie service? For an airline? For VHS (Hi-Fi tracks or linear tracks), or LaserDisc? Each media and customer may have a different requirement or request.

When the 2-channel LtRt signal is actually used for a home video format, further signal processing may occur. One example to explore is network broadcast television. The LtRt signal provided to a broadcast network may be (based on network request) highly compressed. When this signal is broadcast, the network will pass the signal through a signal processor in order to protect against overmodulation and to "improve" the signal for their audience. When the signal arrives at the local affiliate station the signal will again be passed through a signal processor to prevent overmodulation and to "improve" the signal for their audience. The result is often that the signal received by the poor consumer has been so modified that it bears only a faint resemblance to the original soundtrack which was so painstakingly produced.

THE AUDIENCE AND THE SET-TOP BOX

The problem in audio delivery to a diverse audience is that different members of the audience have different needs. This is in contrast to the original goal of a film sound production which is to produce a soundtrack destined for reproduction in the cinema to a captive audience. The original soundtrack needs to be altered to be useable by most consumers in a home environment. Unfortunately, conventional delivery systems (as well as some new ones) require that a single version of the audio program serve the entire audience, even though different members of the audience have differing wants and needs.

This section will look at the problem from the point of view of an audience member receiving service from a set-top box decoding a compressed digital service employing Dolby AC-3 audio coding. The audio component of this service may contain all of the 5.1 channels of sound originally produced. The set-top box will contain an AC-3 audio decoder which decodes the audio into an appropriate 2-channel stereo representation.

The audience will be segmented based on the type of interface used between the set-top box and the audio/video reproduction equipment. Three types of interface will be considered: RF remod on Ch3; audio line level baseband; and raw audio bit stream.

RF Remod on Ch3

Program channels will be received by the set-top box either as conventional NTSC signals (the off-air channels) or as digitally compressed signals. The off-air NTSC channels can be frequency converted to Ch3. If an off-air channel has stereo audio the RF output of the set-top box will contain stereo audio. The digitally compressed channels will be modulated on Ch3. In this case the RF output will invariably be monophonic due to the high cost to generate a BTSC stereo signal.

The consumer using the RF output will thus only receive monophonic sound on the digitally compressed premium channels. This monophonic sound will generally be reproduced by the loudspeaker inside the TV set. It is appropriate that this audio be quite compressed. The speakers inside a TV set are generally not capable of reproducing audio very loud, and so the volume will rarely be turned up very high. It is desirable that low level sounds not be reproduced so quietly that they are lost in the background noise typical of a home. The modulated level of dialogue in the RF signal should be close to that of off-air channels. This allows channel hopping without wild level fluctuations. The peak modulation level allowed may be 6-7 dB over 100% since a TV set can accept this level of audio overmodulation, and the audio modulation level on the output of a set-top box is not subject to regulation.

The type of audio appropriate for this consumer is similar to that appropriate for broadcast, except that peak levels may be permitted to be 6-7 dB higher. Loud sounds may (fortunately) be subject to less peak limiting than in the case of off-air broadcasts. Quiet sounds have a similar need to undergo low level compression (where the quieter sounds are brought up in level to improve their audibility). The level of speech should match that of off-air broadcasts.

Baseband Audio Output

Many consumers will wish to enjoy higher quality audio and video than can be supplied over the RF output. These consumers should be encouraged to access the baseband audio and video outputs. The baseband audio output will be a 2-channel stereo signal which, in the case of a 5.1 channel broadcast, may be matrix decoded into a 4 channel signal. This output will generally be connected to a much higher quality sound reproduction system than that found in a simple TV set. This sound reproduction system may include a matrix surround sound decoder, and may be capable of playing as loud as a cinema. When receiving an off-air channel, the baseband output will invariably be monophonic due to the cost to decode a BTSC stereo signal (although some manufacturers might supply BTSC decoding capability as an option).

The audio delivered out of the baseband audio output is not subject to the modulation limits which apply to the RF output. It is still desirable for the level of speech to match between off-air broadcast channels and digitally compressed channels. The lack of a modulation limit means that audio on digitally compressed channels can potentially be reproduced without any peak limiting. The fact that the audio reproduction system connected to the baseband audio output may be able to play very loud means that it is useful, for some of the listeners some of the time, to be able to deliver the soundtrack free from low level compression. However, most of the listeners most of the time will not be reproducing the soundtrack at anywhere near its intended (by the original mixers mixing for the cinema) loudness. Thus some peak limiting and low level compression is still very useful on the baseband outputs. The limiting and compression do not need to be as severe as in the case of the RF output and, ideally, they should be selectively defeatable by the home theatre enthusiast who wishes to enjoy the sound in all its intended glory.

AC-3 Bit Stream Output

Some channels will supply digitally encoded audio bit streams containing the full 5.1 channel audio mix originally produced for the cinema. It is not anticipated that any set-top box decoder would ever provide all of these channels as discrete outputs as that isn't necessary. The consumer who wishes to enjoy multichannel surround sound will have a surround sound decoder. Today these units invariably contain a Dolby Pro Logic decoder which can decode the LtRt matrix encoded stereo signal. The set-top box can supply such a matrix encoded signal when receiving a multichannel audio transmission.

While the set-top box can supply a 2-ch signal on the analog baseband audio outputs, it can also supply a 2-ch PCM digital signal on an S/PDIF (IEC 958) digital audio output. This type of output (commonly found on CD players and DAT recorders) is preferred in order to deliver a 2-ch signal into a matrix surround decoder which employs digital signal processing, as it avoids an unnecessary stage of D-A and A-D conversion.

Today the home entertainment audio visual (A/V) receiver provides surround sound by means of matrix decoding of 2-ch signals. In the near future (late '95) models of these receivers will become available which also can decode raw AC-3 encoded bit streams into discrete 5.1 channel signals. The AC-3 bit stream interface on these new A/V receivers will be a modified form of the S/PDIF interface. This version of the interface is capable of conveying either the raw AC-3 bit stream or 2 channels of PCM audio. A particular input on the A/V receiver will thus accept analog 2-ch audio, digital 2-ch audio, or a raw AC-3 bit stream.

It is appropriate that the set-top box be capable of delivering either the decoded 2-ch PCM signal or the raw AC-3 bit stream out of the single S/PDIF output pin. The consumer using this output would (potentially) have access to all features provided by the AC-3 bit stream. The consumer using the AC-3 bit stream output will want access to the original soundtrack exactly as it was produced. However, like the consumer using the baseband output, much of the time this user will wish to reproduce the audio at a lower than intended level, and will thus wish to take advantage of peak level limiting and low level compression.

THE AC-3 SOLUTION

When the AC-3 coding system initially developed for the cinema application was adapted to a form suitable for consumer use, the practical and differing needs of consumers were carefully evaluated. Early on in the development of the consumer version of AC-3 the designers decided to create a signal representation syntax which would allow a single encoded bit stream to be decoded into a form useable by nearly every potential listener. Factors which were considered included:

- How each listener might receive the audio signal.
- How many loudspeakers each listener might have.
- The dynamic range available on an RF interface link.
- The dynamic range which is subjectively desired.
- Level matching between programs and channels.
- Absolute traceability back to original mixing room setup.

Decoder Downmix

When a 5.1 channel audio program is reproduced over a monophonic or stereophonic reproduction system, it is necessary to produce a 1 or 2 channel downmixed version of the original signal. The downmix can occur at the encoder if one is willing to employ matrixing techniques. This is the approach which was taken by MPEG multichannel audio, the so-called backward compatible (BC) system.

In the BC approach adopted by MPEG a multichannel audio transmission consists of two components. The first component is a 2-ch downmix. The second component is a set of three additional signals. The stereo decoder only has to deal with the first component. The multichannel decoder has to decode both components and, by means of a linear matrix operation, combine the two components to realize the original 5 channels. The approach is actually a 5-5-5 matrix system, where 2 of the intermediate 5 channels is the 2-ch downmix.

If the BC approach is used, all receivers end up with the same identical downmix. This can be a problem as different listeners will require different downmixes. Listeners who will be Pro Logic decoding a 2-ch downmix into 4 channels need a matrix surround encoded downmix. Listeners who listen in conventional 2-ch stereo or mono will benefit from a different type of downmix. In particular, if the mono listener receives a downmix intended for use by a matrix surround decoder, the information from the surround channels will not be present at all! Another problem with the BC approach is that the matrixing technique employed reduces the coding gain so that (as shown by the MPEG multichannel audio tests) a higher bit rate is required to achieve a given level of audio quality.

AC-3 is a non-backward compatible coder (NBC) which simply encodes all 5 channels. This allows the appropriate downmix to be performed in the decoder. Several types of downmix are available. One downmix is suitable for Dolby Pro Logic matrix surround decoding. Another available downmix is suitable for conventional stereo or (upon further summation) mono reproduction. This downmix features adjustable downmix coefficients so that the precise downmix can be adjusted to best suit the particular program which is being broadcast. The values of these coefficients are selected by the program provider and are delivered in the AC-3 bit stream to the decoder. The intent is to allow the program to be mixed for optimal 5.1 channel reproduction, and then to adjust the downmix coefficients for optimal stereo or mono reproduction. This is different from practice today, where a matrix surround mix is often severely compromised in order to adjust the result obtained in a mono or stereo reproduction of the mix. Today there is no way around the need to compromise; tomorrow, with AC-3, there will be.

Another problem with an encoder downmix has to do with the peak level buildup which occurs when signals are combined. Assuming one is trying to achieve some sort of level uniformity (where the speech is encoded at a common level) there will be a fixed amount of headroom available. When multiple channels of sound are downmixed into two channels, the potential peak level increases. If the available headroom is fixed, peak limiting will be required. This limiting will alter the signal and the original program dynamics will be lost.

In the case of the MPEG BC coder, if peak limiting is applied to the 2-ch compatible signal, in order for the full 5-ch decoder to unmatrix the 5-5-5 matrixed signals back into a discrete form, the same peak limiting must be applied to the 3 additional channels. This ruins the potential for the stereo or multichannel listener to ever exactly recover the original sound track dynamics. If a large amount of headroom is provided so that the peak levels after downmixing may be handled by the coder, there is still the problem of the RF remod with its limited headroom and the desire to match the modulated level of off-air speech. Either the level of speech will be much lower compared to off-air, or the peaks will severely overmodulate the remod, or a peak limited soundtrack must be supplied to the entire audience. The MPEG BC coder is constrained such that it cannot supply a proper version of the soundtrack to a wide audience.

While the AC-3 system does have a potential problem with the peak level of the 2-ch downmix occurring in the decoder vs maintaining a match to the off-air speech level, an elegant solution is provided in conjunction with the level normalization and compression control signals (compression now meaning level compression not bit rate compression).

Level Uniformity

We define level uniformity in terms of human speech spoken in a normal tone of voice (neither shouting nor whispering). When one hops between channels, or one program segment ends and another begins, there is typically someone talking. The subjective level of that speech should be a constant, or else we will be annoyed by the level fluctuations.

One way to establish level uniformity is to require every program to encode speech at a common level. This requires broad agreement and strict enforcement (i.e. this is impossible). Another method is to allow speech to be encoded at any level (which is easy) but to indicate in the bit stream the level at which speech had been encoded (which is possible). This allows different programmers to operate differently and simply asks them to tell us, by means of a value in the bit stream, how they operate. The AC-3 bit stream contains a bit field which allows the level of encoded speech to be indicated. This value is used by the decoder to adjust the reproduction level. The result is that all programs and channels may be reproduced with a common level of dialogue. In the case of the Ch3 remod, this allows the modulation level to be set so that the modulation level of speech matches that typical of off-air channels.

Dynamic Range Compression

The AC-3 bit stream contains an encoded representation of each individual audio channel free from any signal processing such as limiting or compression. No multichannel matrix processing is employed. The dynamic range of the encoded signals is much larger than desired by much of the audience much of the time.

The bit stream also contains compression control signals. These signals may be used by the decoder to reproduce the audio with an altered dynamic range. The control signals are generated by an intelligence up stream of the decoder. One type of control signal is generated to control the level of a mono downmix intended to be used by a 75 μsec emphasised FM modulator. The other type of control signal is intended to be used to provide an artistically reduced dynamic range signal for a 2-ch or multichannel reproduction using the baseband or bit stream interconnect.

Compression for Ch3 Remod

The AC-3 bit stream contains a control signal (called *compr*) which is used when the decoded audio is intended to be modulated onto an FM carrier. The compr control signal is generated to assure that peak modulation is controlled to an acceptable value when speech is modulated to a value consistent with off-air broadcast channels. The use of this control signal by the decoder in the set-top box is well defined.

The compr control signal generator algorithm takes into account the indicated dialogue level, and assumes 75 μsec emphasis in the modulator along with a decoder downmix to mono. The algorithm may be told the acceptable peak deviation (typically 6-7 dB over 100%). The decoder, when optimizing the RF remod output, will use compr and the static dialogue level indication to adjust the decoder output level into the FM modulator. The result will be that the modulation level of speech will match that of the off-air channels, and peak deviation will not exceed the acceptable level. This will be the case no matter how many channels of audio are included in the bit stream. Of necessity, the amount of peak limiting (indicated by compr) which will be applied in the decoder will increase as the number of audio channels is increased.

The consumer using the Ch3 RF output of a set-top box is the type of consumer who will appreciate receiving audio which has had relatively more low level compression applied. Besides indicating gain reduction to prevent high level signals from causing overmodulation, compr can indicate gain increases which may be used to bring very low level sounds up in level. Thus the signal delivered at the Ch3 RF output may be tailored for the portion of the audience using that output, without affecting the signal delivered to other members of the audience.

Artistic Compression

Members of the audience who use either the baseband output or the bit stream output will typically not wish to be subjected to the full dynamic range of the original soundtrack, but will not wish to have the dynamic range restricted as tightly as is done by compr. AC-3 has a second dynamic range control signal (called dynrng) which is intended to be used by this portion of the audience.

The dynrng control signal indicates to the decoder that the signal level should be modified by up to ± 24 dB. This control signal is generated to provide a pleasing amount of gain reduction for loud sounds (those above dialogue level). Thus the maximum reproduced loudness may be controlled. For quiet sounds (those below dialogue level) dynrng can indicate a gain increase. This will prevent quiet sounds from becoming so quiet that they will be lost in the typical ambient noise level of a domestic listening environment. When not generating an output suitable for the Ch3 RF output, the AC-3 decoder will, by default, use the dynrng control signal and reproduce a restricted dynamic range as intended by the program provider. The implementation of the decoder may provide the individual listener the option to disable, in whole or in part, the use of this control signal. If the control signal is disabled in the decoder, the full dynamic range of the original soundtrack will be reproduced. The dynrng control signal can also be attenuated, so that is only has a partial action. If it is set for 50% scaling, then an indicated gain reduction of 10 dB will result in an actual gain reduction of 5 dB. The scaling can be adjusted differently for values of dynrng which indicate gain increases and those which indicate gain reduction. For example, values which indicate gain reductions could be scaled by 80% while those which indicated gain increases could be scaled by 30%. This would result in loud sounds reproducing with slightly less gain reduction than indicated (a 10 dB indicated reduction would become an 8 dB gain reduction), and quiet sounds being reproduced with much less gain increase than indicated (a 10 dB indicated gain increase would become a 3 dB gain increase). The program provider provides a control signal suitable for the mass audience. Individual audience members may select to adjust the reproduced dynamic range to suit themselves. The home theatre enthusiast is able to reproduce the original soundtrack exactly as it was produced.

There is one limitation to the user control of the use of dynrng. When the decoder is receiving a multichannel audio program and performing a downmix a to 2-ch stereo signal, there is the potential of overload due to channels being combined. For multichannel audio programs, the value of dynrng will be generated such that adequate gain reduction will be indicated to prevent overload in the downmix. In this case, the optional scaling of values of dynrng indicating gain reduction is prohibited in order to assure that the 2-ch downmix does not overload.

Audio Production Information

The most accurate reproduction of the original soundtrack occurs in the mixing studio where it is produced. After all, that is where the artistic decisions are made which result in the final soundtrack. For the soundtrack to be accurately reproduced, the sound reproduction system must match that of the original mixing studio, and the soundtrack must be reproduced at the same loudness (volume setting) used in the mixing studio. To facilitate this, the AC-3 bit stream may carry information about the mixing room characteristics and the volume level of the original mix.

The sound reproduction system of a mixing studio will be calibrated one of two ways. Small studios mixing for reproduction in a small room are set up with a flat monitor characteristic. Large mixing rooms which produce soundtracks for the cinema are set up with a monitor characteristic known as the 'X' curve[4], which is has an approximate 3 dB per octave high frequency roll-off beginning at 2 kHz. Large rooms are equalized to the 'X' curve to make them subjectively match reproduction in small rooms which are equalized to have flat response. If the 'X' curve were ideal, then a movie soundtrack mixed in a large room could then be played back in a small room and the frequency response would appear to be correct. Unfortunately, the 'X' curve is not perfect, and when programs mixed to the 'X' curve in a large room are reproduced in a small room equalized flat, the sound is perceived as being slightly too bright. Thus some equalization may be required when soundtracks mixed in a large room are reproduced in a small room, or vice versa. The AC-3 bit stream can carry a bit field which indicates what room size was used to create the original mix. Any decoder has the option to use this information and implement appropriate equalization.

Sound is perceived differently as its reproduction volume is changed. Accurate (to the original artistic intent) reproduction is only possible if the final sound is reproduced with the same absolute loudness as that used in the original mixing session. The AC-3 bit stream can carry a value indicating the absolute loudness calibration of the original mixing room. By using this value, a sound reproduction system can automatically set the reproduction volume to match that of the original mixing session. If the listener chooses to reproduce the soundtrack at a lower level, appropriate equalization can be provided to compensate for the ear's amplitude dependent frequency characteristic.

CONCLUSION

There are a myriad of issues which must be confronted in order to deliver an audio signal simultaneously optimized for many different types of listeners. The AC-3 coder has taken the approach which provides the original soundtrack in a form which may be reproduced exactly as intended by the original mixing engineers. Extra control elements in the bit stream allow independent optimization of the sound reproduced by the TV set receiving an RF input on Ch3, or the hi fi stereo system receiving a stereo baseband signal. The full featured AC-3 decoder receiving the bit stream output can take full advantage of all features provided by the AC-3 bit stream. The dialogue normalization feature will end, once and for all, consumer complaints about uneven loudness between channels. The final result will be more enjoyment of the audio soundtracks by all listeners.

---

1 "Digital Audio Compression (AC-3)", ATSC Standard, Doc. A/52, 10 Nov. 1994

2 Todd, et. al., "AC-3: Flexible Perceptual Coding for Audio transmission and Storage", 96th AES Convention Preprint 3796, Feb. 26, 1994, Amsterdam.

3 "Multi-channel stereophonic sound system with and without accompanying picture", ITU-R Recommendation BS 775-1.

4 ANSI/SMPTE 202M 1991.