The CALM Act and Cross-platform Broadcast

Thomas Lund
TC Electronic A/S
Risskov, Denmark

Abstract – ***The ITU-R BS.1770-2 recommendation specifies a method of measuring Program Loudness and aims at aligning loudness levels across programs of various genres as well as commercials; a distinct improvement over an earlier idea of controlling loudness based merely on the level of speech. Because of the CALM Act, speech-centrism was also given up in ATSC A/85. This paper details the Loop spanning from production to multi-platform delivery that paves the way for high-quality audio across genres, across platforms, around the globe. Optimized normalization for Apple iPod and iPhone devices is reported, and the ongoing NoTube project is described. Finally, the paper explains True-peak measurement, a complementary part of new loudness standards, and novel ways of using True-peak assessment to prevent down-mix overload in the AC3 codec.***

PROGRAMS, PLATFORMS AND PROFIT

Recent year's proliferation of channels and platforms, where the number of listeners per stream goes down, combined with a more dynamic and thus less predictable consumer environment, makes it mandatory for a broadcaster to consider these five factors before committing to any change of station procedure, or to any new technology investment:

1) Are we addressing listener concerns?
2) How well does a technique cater to the station's majority of programs?
3) Does it bring content creation time down?
4) How does it facilitate cross-platform distribution?
5) Are we going down a one-way alley, or will we retain freedom to maneuver in the future?

The digital TV transition has created a bigger market for consumer devices and gadgets. Technology providers and patent-holders have also grown fat, but that's in sharp contrast to year reports from broadcast networks telling a unanimous story of eroding profits. It's therefore necessary to move focus to the broadcaster rather than industries such as film, music, IP or consumer electronics.

Consequently, the five questions above will be called upon when various ideas and technologies are scrutinized later. They should also be kept in mind as more countries consider initiatives/legislation on commercials inspired by the BCAP Codes and by the CALM Act.

ITU-R BS.1770-2

International broadcast is moving away from the two schemes that created systematic level jumps between regular programs and commercials, namely peak level normalization and speech based normalization. As illustrated by pro-active and legislative initiatives in several countries (UK, Italy, China, US and counting), such jumps are the home viewer's main audio concern.

Peak level measurement has typically taken two forms: Sample peak level or quasi-peak level based [9, 10, 13]. While quasi-peak metering with a reasonable headroom leads to less distortion than maximization based merely on sample peak restriction (a concept from the music industry), neither method is relevant across genres because peak/average ratio shows systematic variation depending on genre, see Fig 4. Peak level meters are of little use for platform adaptation, and a specialist is required for the reading. Consequently, these meters fail on 1), 2), 3) and 4).

Speech based normalization was proposed in the early stages of digital TV in the US, but that concept also didn't help preventing loud commercials, or work well across genres. It's a technique originating from film where the production process is less time-critical than in broadcast. Furthermore, regular speech can't be defined, unless a patent protected algorithm is accepted as the reference. Finally, speech measurement has been sold as a prerequisite for setting dialnorm metadata in the AC3 codec, but used that way, normalization has effect for one platform only. A speech based leveling system therefore fails on 1), 2), 3), 4) and 5).

In 2006, ITU-R specified a broadcast alternative to the two deficient normalization schemes from the music world and from the film world, but the BS.1770 recommendation wasn't really made application-friendly and efficient across genres until March 2011. At that time, it was updated to BS.1770-2 through the addition of a relative measurement gate. ITU received a Tech & Engineering Emmy Award for BS.1770-2 at the Consumer Electronics Show, January 2012.

After the publication of *BS.1770-2* [1], programs of any genre may be normalized transparently with other programs, and with commercials, using one and the same measurement based entirely on open standards. The new definition of Program Loudness is backwards compatible, and even improves the setting of dialnorm metadata in the AC3 codec.

The unit for expressing loudness is [LKFS] or [LUFS]. They're both the same, signifying *absolute* loudness. The reason for the different spelling is merely a question of either keeping the original unit (LKFS), or to comply with ISO naming conventions (LUFS). The loudness of a program measuring -23.5 LKFS is therefore exactly the same as the loudness of a program measuring -23.5 LUFS.

Loudness values may also be shown *relative* to Target level using the unit [LU]. If a station's Target level is -24 LKFS / LUFS, and a program measures -22.5 LKFS / LUFS, that value may be expressed as +1.5 LU on a meter. In either case, the value indicates that the program should be lowered by 1.5 dB. Selecting LKFS / LUFS or LU is merely a question of user preferences [5, 20].

While ITU-R BS.1770-2 is a significant improvement over its predecessors, European Broadcasting Union has added further complementary tools to cover all types of programming, including extra defenses against strident commercials, known as *EBU R128* [4]. Through open standards, additions such as Loudness Range, Momentary Loudness and Short-term loudness are harmonized across meter vendors. Furthermore, because compliance is secured through an abstract description in combination with a comprehensive selection of test signals, each meter manufacturer may design a platform-efficient implementation. Together, a guarantee against loudness meter prices getting excessive. At the time of writing, a host of compliant plug-ins, apps and hardware meters had become available, and some of the restrictions on older IP protected measurements had been somewhat relaxed. R128 includes four output documents: Meter harmonization (Tech 3341), definition of the Loudness Range descriptor (Tech 3342), practical guidelines (Tech 3343) and distribution guidelines (Tech 3344) [5-8]. In total, EBU R128 gives broadcaster-friendly answers to the five questions of the first section.

The Transparent Loop

Because the same measurement may be applied in production, ingest, transmission and logging, a transparent loop can be established all the way from production to delivery to any platform. The loop may even be closed, with feedback from logging used to improve production step by step, and to gradually back-off station processing.

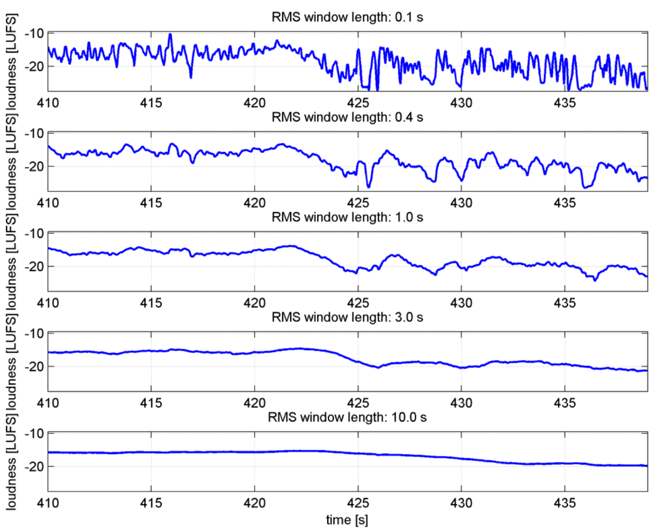

FIG 1. ABSOLUTE LOUDNESS ON DIFFERENT TIMESCALES OF THE MOVIE PULP FICTION FROM 00:06:50 TO 00:07:20.

To help this transparent loop, Loudness Range (LRA) is a new statistical tool for making objective mixing and processing decisions at the point where most options are available: During production. LRA is compliant with BS.1770 and designed to measure loudness variations *inside* a program or a music track using the unit [LU].

Broadcast normally should not be mixed like a cinema movie, nor like a pumped-up up commercial. LRA provides an understandable value to aim at, usable to audio engineers, video editors and journalists alike. Fig 1 shows loudness changes in a clip from the movie Pulp Fiction: Relatively loud music plays until halfway through the clip when the scene changes into dialog. Both scenes sound even in loudness, but the first scene is noticeably louder than the second. The 3 s time-scale seems ideal for measuring the magnitude of that macrodynamic change; the 1 s time-scale shows the same tendency but more noisily, and the 10 s time-scale blurs the change unnecessarily. LRA catches this difference because it's tuned to time-scales relevant to film, broadcast and to music; and not just to one genre [21].

LRA has proven useful in BS.1770-2 production, but also at stations where programming is mostly anchored to speech. Speech based anchoring has the side-effect of not catering particularly well to broadcast needs. LRA tells a clearer story of what the consumer hears, and it may also be used to tighten differences in speech level during production. Such differences can remain unnoticed in a movie, yet be annoying in broadcast.

This is the advice of how to use Loudness Range in Practical guidelines for Production and Implementation of R128 [7]: "With Loudness Range (LRA) it is now possible to quantify the dynamics of a programme. In the past, it had to be "educated guesswork" of experienced audio personnel to decide if a programme would fit into the loudness tolerance window of the intended audience. Using Loudness Range, the guesswork is over - at the end of the measurement period (usually the whole programme), a single number enables the mixer/operator to decide if further dynamic treatment is necessary."

Used during ingest or on a broadcast server, LRA is an objective measure for deciding when programs for delivery to certain platforms require range restriction. HD platforms may be set to tolerate any LRA value, though a limit such as 12, 15 or 20 LU, depending on genre, may be recommended to production and in delivery guidelines [20]. Programs for mobile platforms can be automatically pre-conditioned to keep LRA below, for instance, 8 or 10 LU. Other platforms may be based on different LRA limits, while range restriction is always performed in an audio-friendly manner around the Target level.

Downstream of production, LRA doesn't change as long as gain offsets only are applied (normalization), but the number reveals when any significant range processing has taken place between two points in the broadcast chain. Loudness Range may therefore also serve as an inspection and logging tool, verifying that no range processing has sneaked in during distribution, or unexpectedly in a codec.

For short programs, under 30 sec of duration, LRA is of less value. Short-term loudness or Momentary loudness are the metrics to use for preventing this kind of programs from becoming too loud. More LRA info can be found in [19-21].

Though an improvement over earlier broadcast production practice, it should be noted how a couple of issues in A/85 [3] currently contribute to a less transparent loop: 1) While the recent Annex J and K, driven by the CALM Act, are steps forward with regard to leveling of commercials based on all sources, and not speech only, the lack of clarity of which gate principle to apply is counterproductive. In case Annex J and K prescribe an un-gated measurement, the limitation is easy to trick by keeping parts of the program soft. 2) The measurement of regular programs in A/85 is not transparent. By using a vague anchor principle, nobody can tell precisely how normalization should be performed.

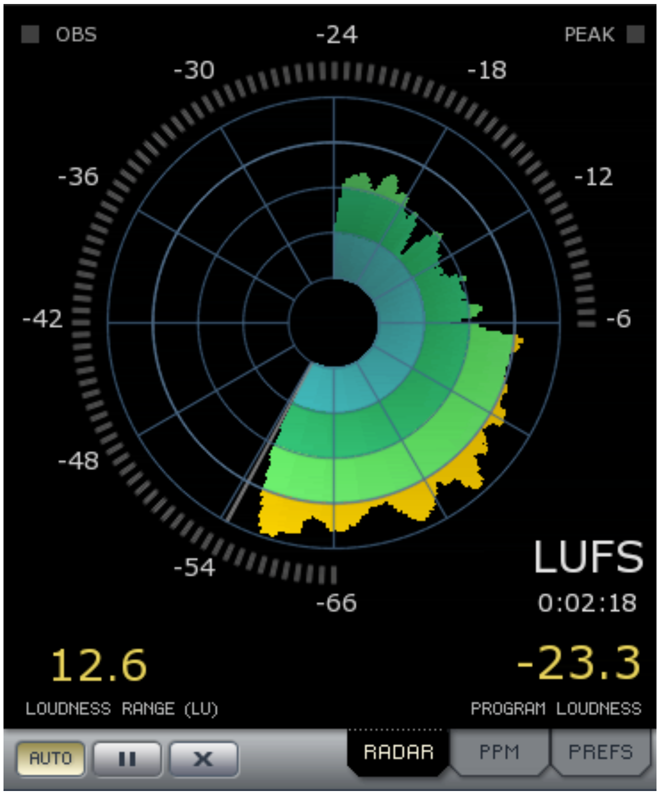

FIG 2. "ALMA" BY TOM LEHRER
BOLD RADAR LINE REPRESENTS -24 LUFS, 6 LU PER DIV
GREEN: SOFTER THAN -24, YELLOW: LOUDER THAN -24

Consider the loudness meter readout in Fig 2, Tom Lehrer performing live in San Francisco, 1965. It's not obvious if this stellar performance should be anchored to the talking part (1st quarter of the radar) or the singing part (2nd quarter of the radar). If the program is anchored to speech, the playing would get louder than a battle scene in Pirates of the Caribbean; and a commercial right after would drown. Such dilemmas are transparently taken care of in BS.1770-2. In general, it's a challenge to find examples where its normalization isn't reasonable [4, 12, 17].

IPOD, MOBILE TV AND NOTUBE

Mobile and computer devices generally have a different gain structure, and make use of different codecs, than domestic AV devices such as television, home theatre etc. [11] Systematic tests have therefore been carried out to determine the standard operating level on Apple devices. Broadcast aimed at these devices must sit at a target level suitable for the their gain structure. If not, the user may not be able to turn up the level high enough to hear a program. iPod, iPhone and MacBook share a normalization function known as "Sound Check". We decided to measure which loudness level, using an BS.1770 scale, Sound Check aims at. Given Apple's history for attention to audio detail, broadcast for iPods etc. should ideally aim at that level.

Apple audio is blessed with digital optical output so there is no doubt about calibration issues. Based on 1250 music tracks (rock, pop, jazz and classical) and 210 broadcast programs (news, scientific, drama, sports, gameshow), the Apple normalization number comes out as -16.5 LUFS / LKFS on a BS.1770-2 scale. It is therefore suggested to aim podcast at a Target level no lower than -16 LUFS. The easiest and best sounding way to accomplish this is to

1) normalize to target level (-24 LUFS or -23 LUFS),
2) limit peaks to -9 dBTP,
3) apply a gain change of +7 or +8 dB.

Applying this principle, differences between foreground sound and background sound isn't washed out, and the whole procedure stays codec-agnostic. Keeping clear of proprietary schemes is an advantage regarding future options, and is also cheaper than getting locked in.

Another project about alternative platform broadcast and the future of TV has been inaugurated. Within the scope of "NoTube", enlightening experiments have been conducted by BBC (UK), IRT (Germany) and 11 other organizations [17]. On the audio side, loudness normalization for multiplatform environments (including loudness harmonization on the web), Loudness Range adaptation, and listening conditions is being investigated systematically.

Evaluations are based on video clips of different genres like "Movie", "Commentary", "Concert", "Sport", "Show", "Commercial" and "News". The audio part of the clips selected for evaluation is varied with respect to Program Loudness and Loudness Range. Evaluating loudness normalization methods, "excellent performance" was found when using the ITU-R BS.1770-2 / EBU R128 Program Loudness metric.

At the time of writing, final results on Loudness Range were not available. However, preliminary results indicate there is "a tendency identifiable", where subjects preferred medium (or even strong) Loudness Range restriction rather than uncompressed audio when listening to web content.

TRUE-PEAK LEVEL

Program to program loudness is just one side of the coin. Overload also has to be taken into account during production, and every time positive gain may subsequently be applied. The upper illustration of Fig 3 shows a situation where reading sample peak level gives the same result as reading the max peak level after D to A conversion [15]. In the illustration below, this is not the case.

Note how sample level may systematically be lower than the level between samples [13-15]. Production of commercials and pop music has taken advantage of the deficient sample peak measure. Unfortunately, exploiting the sample peak meter leads to distortion in sample rate converters, data reduction codecs and in consumer equipment.

Loudness is where user focus should be now, but ITU, ATSC and EBU standards are complemented by an improved measure of peak level, known as *true-peak* metering. The concept is to consider only loudness as long as you're not overloading. To signify a true-peak measurement, the unit [dBTP] is used. In Fig 3, the upper trace would read 0 dBFS on a sample peak meter, and 0 dBTP on a true-peak meter. The lower trace would read -1.25 dBFS on a sample peak meter, but 0 dBTP using true-peak.

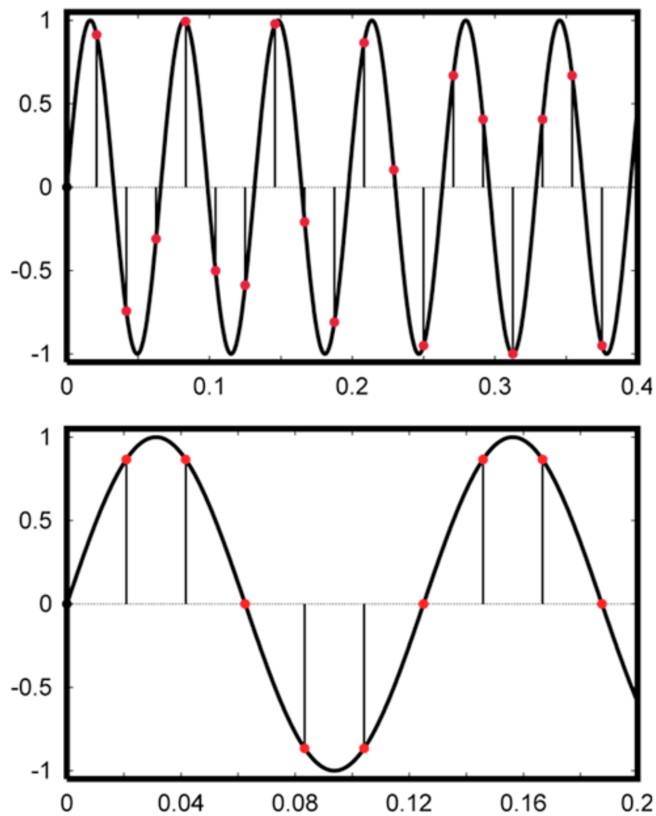

FIG 3. SAMPLES (RED DOTS) VS.
ACTUAL WAVEFORM (BLACK CURVE)

A true-peak meter estimates level between samples by up-sampling the signal at least four times. A true-peak meter may easily reach +3 dBTP or more when fed with a pop/rock CD. Hot signals like this are invisible to a sample peak meter, but need to be attenuated before conversion or data reduction takes place [13-15]. Contrary to a sample meter, a true-peak meter therefore can prevent distortion from building in digital to analog converters, sample rate converters and in lossy codecs.

Based on true-peak metering, what threshold do we need to keep below in order to stay clear of distortion? It's not a goal to be too conservative as that takes away precious headroom for no good reason. We don't wish to be too liberal either, as that can lead to distortion. In linear PCM, also known as "baseband audio", sample rate converters and DA converters perform as expected all the way up to 0 dBTP. With a four times up-sampled meter, it may under-read by approximately 0.5 dB, so that's as high one can safely go.

The picture is more muddy if we turn the attention to lossy codecs. Codecs exhibit extra peaking due to change in phase and bandwidth in the encoder. At low bit-rates, they generate more peaking than when less data is thrown away. Furthermore, lossy codecs are less predictable with regard to peaking than other elements of the signal-path [14].

The sensible and audio-friendly solution is to use linear PCM as the reference when deciding on a general max true-peak level. Bandwidth and storage is growing so rapidly these years that it makes sense to specify based on linear PCM, especially in new installations. This is the stance EBU has taken in R128 by setting the general production limit at -1 dBTP. That is 5-6 dB more conservative than today's pop music, yet doesn't waste headroom. ATSC and BS.1771 wear a bit more belt and braces by specifying max peaks at -2 dBTP [2].

Instead of implementing a lowest common denominator scheme, true-peak level may be restricted at the point of transmission depending on how well a codec for a certain platform behaves. Make an informed decision for your own specific conditions by reading a true-peak meter before and after a codec. Try different kinds of programs, and don't limit the pre-encoder peak level more than necessary.

In order to make true-peak meters better harmonized, ITU is considering conformance test signals, a specification method also used to specify various measurements in EBU R128. A comprehensive suite of test signals is generally a good solution because meter vendors can use an implementation suitable for a given hardware platform without readings becoming unpredictable. A fixed algorithm, on the other hand, can be impossible on some hardware, consequently at the risk of excluding efficient and low-cost solutions from performing a given measurement.

June 2011, a set of true-peak test signals was suggested by engineers from Dolby Corporation [16]. While the draft covers aspects of true-peak measurement, an important type of signal was not considered; namely the type of legal sine waves potentially creating havoc in DA converters, sample rate converters and in lossy codecs [14, 15]. If one cannot check for sufficient headroom in the instrument itself, the whole point of true-peak metering is jeopardized.

Headroom in Broadcast

For a signal-path, the ratio between max peak level and average operating level is called headroom. Using BS.1770, headroom can be regarded as the ratio between true-peak level and Program Loudness. The amount of headroom in the signal-path is quite genre dependent, see Fig 4.

In commercials and pop/rock music, the headroom requirement can be 6 dB or even lower, while a cinema movie may need over 20 dB. Furthermore, movies and classical music only need their headroom for a fraction of the time while "beat music" in general requires headroom from start to end.

When a signal-path offers less headroom than required for conveying a program, limiting or clipping will result. A person ignorant of audio can hear jumps in loudness, but headroom is more difficult to describe to him or her: You only hear when there's not enough, and sometimes only if the reproduction system is of a certain quality. With insufficient headroom, transients are distorted and good loudspeakers are wasted, regardless if one is listening to Johann Sebastian Bach, Lord of the Rings or Donald Fagen.

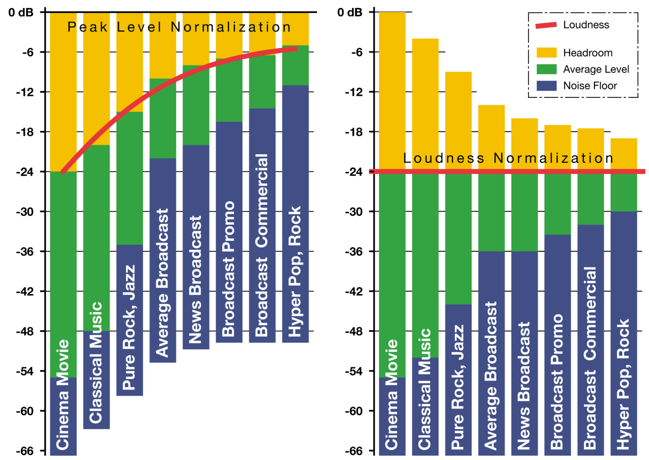

FIG 4. PEAK LEVEL VS LOUDNESS NORMALIZATION.
HEADROOM SHOWN IN YELLOW

Unfortunately, every part of the signal-path may constitute a headroom bottleneck, so broadcast doesn't sound better than its weakest link. This is why comprehensive Distribution guidelines, Tech 3344, are part of the output documents accompanying EBU R128 [8].

In analog broadcast, headroom is frequency dependent with less at high frequency because of transmission emphasis. Analog TV has a headroom of 10-12 dB, while FM radio is often operated with 8 dB or less.

In digital broadcast, noise is generally lower and emphasis is no longer part of the equation. Consequently, a lower average level in combination with a higher peak level is now a possibility. With target and peak level specified by ATSC A/85 (-24 LUFS / LKFS and -2 dBTP), a generous 22 dB of headroom is available, much more than ever before.

The headroom appears to be the same in EBU R128 (-23 LUFS / LKFS and -1 dBTP), but that only holds if the two RPs are based on the same measurement of Program Loudness. At the time of writing, EBU R128 is firmly rooted in ITU-R BS.1770-2, while A/85 (July 25, 2011) ambiguously quotes BS.1770-1, which is no longer in effect.

If A/85 actually stays with the application critical and genre critical BS.1770-1, headroom in A/85 is significantly lower than in R128 [20]. This is because the measurement gate of BS.1770-2 returns a Program Loudness number representing moderately loud parts of the program, disregarding the quiet parts as well as any pre-roll or post-roll. For Wide Loudness Range programs (drama, movie, classical music etc.), assuming a certain target level, headroom is 3 dB or more better with BS.1770-2 than with its predecessors. Under EBU R128 specs, The Matrix or Ravel's Bolero may be transmitted without the need for dynamics processing at all, where this would not be the case had BS.1770-1 been used for normalization.

Speech based normalization is less clear than measuring all sources, with regard to the headroom needed for distribution. Speech based anchoring is a concept practiced indirectly in film where there's more time for production. A calibrated listening environment, the facilitator in film, could be applied also in prestigious production for broadcast, but even under pristine production conditions, the Loudness Range of speech in a feature movie can exceed the loudness jump tolerance in broadcast [19]. Taking The Matrix again as an example, regular speech varies between -24 LUFS and -46 LUFS. The movie doesn't fit under the A/85 regime without processing, or without reverting to costly measures such as fitting the entire HDTV signal-path with floating metadata capability without any benefit to the end listener. Floating metadata is one more thing to go wrong, and only half a solution for one broadcast platform.

Headroom and AC3

Used for stereo only, the AC3 codec isn't more sensitive than other codecs at a similar bit-rate, with regard to the max true-peak level it handles without clipping. If a typical pop/rock track is encoded without attenuation, AC3 clips frequently like other lossy codecs. If the same track is attenuated so peaks don't exceed -1 dBTP, the problem is gone.

The real challenge with AC3 and headroom is the way it handles 5.1. A majority of consumers are listening in stereo, regardless if programs are in 5.1 or in stereo. In countries where AC3 is transmitted without an independent stereo stream, the decoder has to down-mix every time a 5.1 program comes along, and this is where problems start. The decoder doesn't include a transparent down-mix limiter, so option number one is to use *conservative* mix coefficients in order not to generate stereo overloads when all 5.1 channels are busy. Conservative settings means something like L, R: -6 dB; Center: -9 dB; SL, SR: -12 dB. Now there will be no mix overloads, but instead we will have systematic level-jumps when programming switches from native 5.1 to native stereo.

Consequently, the real peak level problem in AC3 doesn't come from the data reduction system itself, but from the down-mix section in the decoder. If only broadcasters could keep peak level low, decoder mix coefficients wouldn't have to be set conservatively. On the other hand, it would be a shame if a general restriction of headroom in broadcast was inflicted because of first generation codecs with technical design issues.

The most audio-friendly compromise with AC3 is therefore to peak-limit before the encoder in combination with coefficients that don't create level jumps, let's call them *benign* L, R: 0 dB; Center: -3 dB; SL, SR: -6 dB.

Recent experiments have pointed to a solution more tolerable from an audio point of view than using a general limit threshold at -6 dBTP. In 5.1 action movies, one channel generally uses more of its headroom than the others, namely the center. Dialog is challenged by *two* front channels, so it uses up more of its headroom. The AC3 down-mix solution is therefore simple: Use -6 dBTP limiting for all the lateral channels, but -3 dBTP for center. Keeping the offset, thresholds may be moved closer to 0 if the main concern is preservation of headroom rather than a 100 year storm.

Limiting should be applied pre encoder, and benign mix coefficients are specified as part of a static metadata structure where DRC safely may be set to off.

Conclusion

Loudness and true-peak based broadcast recommendations have been described. ITU-R BS.1770-2 provides more headroom for the broadcast chain than ever before, and it enables a listener to enjoy music and film unaltered, regardless of which distribution format is used. The last weak element of broadcast is the lossy codec. Once codecs are disposed of, another significant step up in quality may be immediately taken. In the meantime, broadcast should not be forced to carry all burdens of these improvements. Consequently, it is suggested to ask five questions to probe new RP and new legislation before implementation. Otherwise, stations are put at disadvantage against "web-blasters" serving one or two platforms only. Obviously, broadcast standards are written neither for the sake of music or commercials only, nor solely for the film industry.

Because of the CALM Act, annexes J and K have been added to ATSC A/85. While it's a step in the right direction to explicitly state the need for measuring all sources with commercials and promos, application and transparency aspects of these annexes could be further improved by the RP referencing BS.1770-2 rather than BS.1770-1 which is no longer in effect. Should more defenses against strident commercials be needed, procedures specified by BCAP may be considered, as well as additional tools from EBU R128.

The paper has also shown how gating differences in the loudness measurement directly influences the amount of headroom available in the signal-path; and that BS.1770-2 provides more headroom for the same target value than BS.1770-1 or speech gating. The normalization target of Apple devices has been determined using a BS.1770-2 scale. Results reveal a gain structure not suitable for a target level as low as normal broadcast. Therefore, an easy target transcode procedure, useful with any platform and any codec, has been described.

No-compromise audio for broadcast is a fact, including a capability to scale easily for cross-platform delivery. The NoTube project is systematically investigating listener preferences in this respect with preliminary results reported here. Optimized normalization and dynamics processing for mobile and IPTV platforms may be based entirely on objective criteria. With a novel BS.1770 compliant tool, Loudness Range, these criteria are transparent already in production.

References

[1] ITU-R, "Rec. ITU-R BS.1770-2, Algorithms to measure audio programme loudness and true-peak audio level", *International Telecommunications Union*, 2011.

[2] ITU-R, "Rec. ITU-R BS.1771, Requirements for loudness and true-peak indicating meters", *International Telecommunications Union*, 2006.

[3] ATSC, "Techniques for Establishing and Maintaining Audio Loudness for Digital Television", *Advanced Television Systems Committee Doc A/85*, 2011.

[4] EBU, "EBU Technical Recommendation R128 -- Loudness normalisation and permitted maximum level of audio signals", *European Broadcasting Union*, 2010.

[5] EBU, "Tech Doc 3341: Loudness Metering: 'EBU Mode' metering to supplement loudness normalisation in accordance with EBU R128", *European Broadcasting Union,* 2010.

[6] EBU, "Tech Doc 3342: Loudness Range: A measure to supplement loudness normalization in accordance with EBU R128", *European Broadcasting Union,* 2010.

[7] EBU, "Tech Doc 3343: Practical guidelines for Production and Implementation in accordance with EBU R128; version 2.0", *European Broadcasting Union,* 2011.

[8] EBU, "Tech Doc 3344: Practical guidelines for distribution systems in accordance with EBU R128; version 1.1", *European Broadcasting Union,* 2011.

[9] Klar, S. & Spikofski, G., "On levelling and loudness problems at television and radio broadcast studios", Paper #5538 of the 112th AES Conv., Munich, 2002.

[10] Rumsey, F., "If It's Loud, Does That Mean It's Bad", *JAES*, vol 56, iss 6 pp. 493-498, June 2008.

[11] Stoll, G., "Streaming-Audio@Internet: Perspectives for the Broadcasters", *Paper #17-003 of the 17th International AES Conference*, Florence, September 1999.

[12] Camerer, F., "On the way to Loudness nirvana - audio leveling with EBU R128", *EBU Technical Review*, vol. 2007, April.

[13] Nielsen, S.H. & Lund, T., "Level Control in Digital Mastering", *Paper #5019 of the 107th AES Convention*, New York, 1999.

[14] Nielsen, S.H. & Lund, T., "Overload in Signal Conversion", *Paper #11 of the 23rd International AES Conference*, Copenhagen, May 2003.

[15] Lund, T., "Stop Counting Samples", *Paper #6972 of the 121st AES Convention*, San Francisco, CA, 2006.

[16] Wilson, R. et al., "Proposal for a Set of True Peak Conformance Signals", *Tech document from Dolby Corporation*, June 2011.

[17] EU, "NoTube Project", Research project exploring Semantic Web and TV, *coordinated by Vrije Universiteit, Amsterdam*. http://notube.tv

[18] Skovenborg, E. & Lund, T., "Loudness Descriptors to Characterize Programs and Music Tracks", *Paper #7514 of the 125th AES Convention*, San Francisco, 2008.

[19] Skovenborg, E. & Lund, T., "Loudness Descriptors to Characterize Wide Loudness Range Material", *Paper #7948 of the 127th AES Convention*, New York, 2009.

[20] Lund, T., "BS.1770 Revisited", BE Paper of the NAB Convention, Las Vegas, 2011.

[21] Skovenborg, E., "Loudness Range (LRA) - Design and Evaluation", *Paper of the 132nd AES Convention*, Budapest, 2012.