On the way to LOUCINESS nirvana

audio levelling with EBU R 128

Florian Camerer *ORF*

This article describes one of the most fundamental changes in the history of audio in broadcasting: the change of the levelling paradigm from peak normalisation to loudness normalisation. This change is vital because of a problem that has become a major

Loudness normalisation is a true audio-levelling revolution!

source of irritation for television and radio audiences around the world – that of the jump in audio levels at the breaks within programmes, between programmes and between channels. Loudness normalisation is the solution to counteract this problem.

EBU Recommendation R 128 [1] establishes a predictable and well-defined method to measure the loudness¹ level for news, sports, advertisements, drama, music, promotions, films etc., throughout the broadcast chain, and thereby helps professionals to create robust specifications for ingest, production, playout and distribution to a multitude of platforms. Four supporting documents have

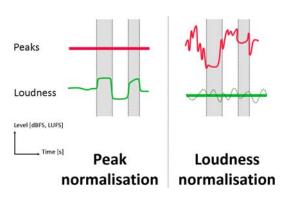


Figure 1
Peak-level normalisation vs. Loudness-level normalisation of a series of programmes

been created by the EBU to aid the audio industry to work with Rec. R 128. It is based entirely on open standards and aims to harmonise the way we produce and measure audio internationally. In addition to **Programme Loudness**, R 128 introduces two more audio descriptors, **Loudness Range** and **Maximum True Peak Level**. All three are designed to work together, forming a set of essential descriptors that characterise an audio signal.

Loudness metering and loudness normalisation signify a true audio levelling revolution (see Fig. 1). Furthermore, this loudness-levelling paradigm affects all stages of an audio broadcast signal, from acquisition to distribution and transmission. Thus, the ultimate goal is to harmonise audio loudness levels within broadcast

channels as well as between channels to achieve an **equal universal loudness level** for the benefit of the listeners. To be clear: the loudness level can (and should!) still vary according to artistic and tech-

^{1.} Loudness refers to the perceived strength of a piece of audio (music, speech, sound effects etc.). The loudness depends on the level, frequency, content and the duration of the audio, amongst other things.

nical needs within a programme. Rather, the loudness normalisation method uses the **average loudness** of a programme, to make the level consistent between programmes and between channels.

The experience of several EBU Members has shown that working with the loudness paradigm is both liberating and satisfactory. The fight for "Who is the loudest" disappears, mixes can be more dynamic, there are fewer dynamic compression artefacts, such as "pumping", and thus there is an **overall increase of audio quality!** Programme makers who favoured dynamic mixes in the past are now relieved from potential compromises because their programme no longer sounds softer than more-compressed ones. With loudness normalisation, this compromise is gone. Nirvana is closer than ever!

The origin of the problem – the "loudness war"

Audio metering in broadcasting today is based typically on quasi-peak programme meters (QPPMs – more usually known as just PPMs). It is "quasi" because of its finite reaction time of 10 ms

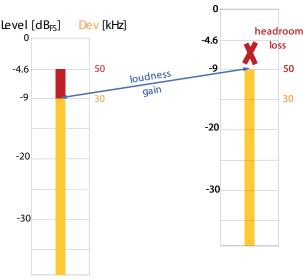


Figure 2
Relationship between Permitted Maximum Level (with QPPM) and FM Deviation and the abuse of it: gaining loudness, but losing headroom

(although 5 ms is also found). In practice, this means that signal peaks shorter than this reaction time won't be displayed correctly, if at all (for example, transients such as those created when keys are jangled). In order to provide headroom for these transients, which one wouldn't see on the meter, but which should nevertheless be there so as to contribute to the "openness" of the audio signal, the agreed *Permitted Maximum Level* (PML) was set at **–9 dBFS**.

This value was based on the familiar – and in many places, still extant – method of delivering sound to the home on an FM carrier. The carrier's maximum deviation for TV was standardized in many countries at 50 kHz and the PML at 30 kHz deviation (equating to –9 dBFS), which thus allowed 20 kHz, or 4.4 dB of headroom.

However, commercial pressures have taken a hold and the response to the pressure to stand out from the competition has been to be louder than it. Modern peak metering and powerful dynamic-range processing allowed organizations

to realign the PML to equate to the maximum deviation (50 kHz) of the FM carrier (see Fig. 2). All the transients have to be chopped at the PML to avoid distortion, but that has been thought an acceptable compromise by those who have implemented it.

When someone at home switches between one of these stations and a station which has not joined the "war" or is taken to a loud advertisement in a well-balanced programme with a wider dynamic range, the audio level jumps and there is a grab for the remote control to adjust the sound to a more acceptable level. In the case of loud advertisements, the sound then has to be re-adjusted when the main programme returns. It is no wonder that so many complaints are received by the broadcasters. Other people solve the problem by muting the sound during the advertisements and so the message is at least greatly diluted.

A Standard Emerges – and the EBU develops it

The International Telecommunication Union (ITU) recognized the problem and its work gave rise to ITU-R BS.1770 [2]. The purpose of that standard was to establish an agreed algorithm for the measurement of loudness and the true peak levels of programmes. It is a robust standard which has the benefit of a simple implementation. In brief, it defines a "K-weighting" curve (a modified sec-

ond-order high-pass filter) which forms the basis for matching an inherently subjective impression with an objective measurement.

This curve is applied to all the channels (except the Low-Frequency Effects (LFE) channel which is discarded from the measurement), the total mean square energy is then calculated (with different gain factors for the front and surround channels, see Fig. 3) and the result is displayed as "LKFS" (Loudness, K-Weighting, referenced to digital Full Scale).

For relative measurements, Loudness Units (LU) are used, where 1 LU is equivalent to 1 dB.

A more detailed study of the algorithm can be found in ITU-R BS.1770 [2] as well as in EBU Tech Doc. 3343 – "Practical Guidelines" [3]).

BS.1770 also defines and recommends the use of a true peak meter for measuring peaks. Such a meter runs at a

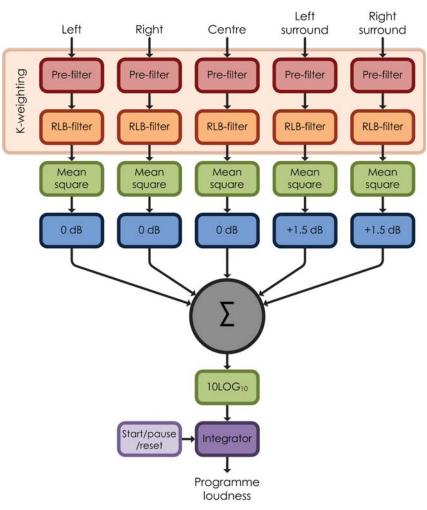


Figure 3
Channel processing and summation in ITU-R BS.1770

higher sampling rate than the audio signal (usually 4x oversampling) to catch inter-sample peaks which might otherwise exceed 0 dBFS and thus cause distortion later in the chain.

R 128 and ITU-R BS.1770 are the basis. Four more EBU technical documents provide details.

ITU-R BS.1770 provides the basis for EBU Recommendation R 128 which extends the ITU standard by actually defining a specific **Target Level** (see below) for loudness normalisation as well as a **gating** method to improve the

loudness matching of programmes which contain longer periods of silence or isolated utterances. The EBU's development was required to accommodate the needs of programme makers, with particular regard to having a means to measure complete mixes (rather than just one component, such as the dialogue) and the loudness range of the programme. To do this, the EBU has specified three new parameters:

- **O Programme Loudness**
- Loudness Range
- O True Peak Level

Programme Loudness

Programme Loudness describes the long-term integrated loudness over the duration of a programme. (In R 128, the definition of the word "programme" is also used to refer to advertisements,

bumpers and other interstitials.) The parameter consists of one number (in LUFS²), with one number after the decimal point) which indicates "how loud the programme is on average". This is measured with a meter compliant to ITU-R BS.1770 with the addition of a gating function. The gate serves to pause the loudness measurement when the signal drops below a certain threshold. Without this gating function, programmes with longer periods of silence or low-level background sounds or noise will get too low a value for the integrated loudness level. Such programmes would subsequently be too loud when broadcast.

Following a series of tests, a gate of –8 dB relative to the ungated LUFS measurement, with a block length of 400 ms, was agreed. The tests also confirmed, along with other findings, the choice of the **Target Level** to which every audio signal will be normalised. It is:

-23 LUFS (-8 rel gate)

-23 LUFS is the new centre of the audio-levelling universe.

A deviation of ± 1 LU is acceptable for programmes where an exact normalisation to the Target Level of -23 LUFS is not achievable practically (such as live programmes or ones which have an exceed-

ingly short turn-round). In cases where the levels of a programme's individual signals are to a large extent *unpredictable*, or where a programme consists of only background elements (for example, the music bed for a weather programme), this tolerance may be too tight. It is therefore anticipated for such cases that the integrated Loudness Level may lie outside the tolerance specified by R 128.

Loudness Range

Another major consideration was the Loudness Range which would be needed to accommodate **all** programmes (provided that they don't exceed the tolerable loudness range for domestic listening). The **Loudness Range** (LRA) descriptor quantifies (in LU) the variation of the loudness measure-

ment of a programme. It is based on the statistical distribution of loudness within a programme, thereby excluding the extremes. Thus, for example, a single gunshot is not able to bias the outcome of the LRA computation.

Loudness Range is a generic descriptor that helps to decide if dynamic compression is necessary.

EBU Recommendation R 128 does not specify a maximum permitted LRA, as it is dependent on factors such as the tolerance window of the average listener to the station, the distribution of genres of the station etc. R 128 does, however, strongly **encourage the use of LRA** to determine if dynamic treatment of an audio signal is needed and to match the signal with the requirements of a particular transmission channel or platform. More details about LRA may be found in EBU Tech Doc 3342 [4].

First experiences at broadcasting stations suggest a maximum LRA value of approximately 20 LU for highly-dynamic material, such as action movies or classical music. The majority of programming will never need to fully use such a high LRA value or, indeed, be able to reach it!

True Peak Level

The True Peak Level of an audio signal indicates the maximum (positive or negative) value of the signal waveform in the continuous time domain; this value is, in most cases, higher than that shown by a quasi-peak meter or even a sample-peak meter, both of which would miss the true peaks which potentially lie between samples. The use of an oversampling meter, which is compliant to BS.1770, allows those peaks to be detected.

LUFS indicates the value of K-weighted loudness with reference to digital full scale. The EBU recommends this unit to overcome an inconsistency between ITU-R BS.1770 and ITU-R BS.1771. This unit complies with ISO 80000-8.

An oversampling meter may still slightly under-read the actual peak value (depending on the oversampling frequency) and so the *Maximum Permitted True Peak Level* for Production is:

Oversampling peak meters provide a good estimate of the true peak of an audio signal.
Sample peak meters don't.

-1 dBTP

Note that some parts of the chain, such as analogue re-broadcasters and users of low-bitrate coders require a lower True Peak Level. The PLOUD Distribution Guidelines (EBU Tech Doc 3344 [5]) contains comprehensive coverage of the topic.

Strategies for loudness normalisation

Fig. 4 shows two approaches, mainly for **production**. The first is more relevant for the early stages of the transition and it is perhaps especially useful to those who work on live programmes. Existing

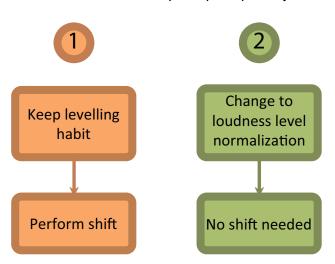


Figure 4
Two principal working methods to achieve uniform loudness in production and post-production

meters, limiters and mixing practices are retained and a level shift is done at the output of the console (after the main meters) to achieve the loudness Target Level of **-23 LUFS**. A loudness meter is placed after the level shift to enable the engineers to understand the exact amount of shift (which initially is still a bit of guesswork).

Using a loudness meter on past programmes of the same genre gives good guidance as to where the levels sit. Early experience at NDR, ORF and RTBF has shown that it is certainly possible for live mixes to fall into the \pm 1 LU window permitted by R 128.

Those who work with files have an easier task as the whole programme may be normalised quickly and easily to **-23 LUFS** by means of a level shift.

Those organizations or sections which can make the move to loudness metering straight away can be expected to reap the benefit straight away. The greater dynamic range possible will be a welcome bonus for crowd noise in sports programmes – for example, enhancing the impact of a game for the viewers and listeners.

EBU Tech Doc 3343 [3] gives practical guidelines for the new way of working with audio levels.

There are also basically *two ways* to achieve loudness normalisation for the **consumer**. One is the actual **normalisation of the source** itself, so that the programmes are equally loud by design. The other method is with the **use** of **loudness metadata** that describes how loud a programme is. For the latter, the actual average programme loudness levels don't need to be changed to a normalised

Abbreviations			
dBFS	dB relative to digital Full-Scale	LU	Loudness Unit
dBTP	a true peak meter	LUFS	K-weighted Loudness Unit with reference to digital Full Scale
FM	Frequency Modulation	PML	Permitted Maximum Level
LKFS	Loudness, K-weighting, with reference to Full Scale	PPM	Peak Programme Meter
LRA	Loudness Range	QPPM	Quasi-Peak Programme Meter

Loudness normalisation of the source is recommended in production because of simplicity and potential quality gain. value and can still vary a lot. For those with upto-date equipment, the normalisation can be performed at the consumer's end, using the individual loudness metadata values to gain-range the programmes to the same replay level.

Within the EBU R 128 loudness-levelling paradigm, the first solution is encouraged due to the following advantages:

- O simplicity;
- O potential quality gain at the source.

The second solution is not forbidden (see also the Distribution Guidelines document – EBU Tech Doc 3344 [5]), but having one single number (-23 LUFS) has great strength in spreading the loudness-levelling concept, as it is easy to understand and act upon. And the active normalisation of the source in a way "punishes" overcompressed signals and thus automatically encourages production people to think about other, more dynamic and creative ways to make an impact with their programme. In other words, the actual **technical** change of the source level through active normalisation to -23 LUFS has direct consequences on the **artistic** process – and in a positive way!

Nevertheless, it must be stated that both methods can complement each other: they are not to be seen as opponents or a black-and-white view of the same issue. Both approaches are a part of R 128 – but because of the advantages listed above, the **normalisation of the source** is recommended.

Working with loudness meters

Until now, mixing with QPPMs and normalising the peaks have been done with reference to a permitted maximum level (typically, –9 dBFS) and a peak limiter at that level provided a "safety ceiling", which could be hit as hard as desired – at the expense of less-engaging sound, of course.

In contrast, the loudness-levelling paradigm more resembles "floating in space" as the schematic representation of a bar-graph meter in *Fig.* 5 shows. *Fig.* 6 shows how a software meter based on a "needle" could look.

The EBU has deliberately not specified any graphical or user-interface details of a loudness meter, but it has specified enhancements of the algorithm described in ITU-R BS.1770 and two scales:

- O EBU +9 Scale which ought to be suitable for most programmes;
- O EBU +18 Scale which may be used for programmes with a wide LRA.

Both scales can either display the relative Loudness Level in **LU**, or the absolute one in **LUFS**. The meter manufacturers in the PLOUD Group have agreed to implement the "EBU mode" set of param-

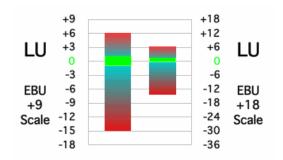


Figure 5
A schematic representation of the two loudness scales (here in LU) as described in EBU Tech 3341

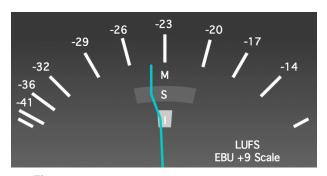


Figure 6
A schematic representation of an emulated loudness meter with a "bendy needle"

eters, to make sure their meters' readings will be aligned. Many more manufacturers have adopted "EBU mode" too, or are in the process of doing so.

An "EBU Mode" loudness meter as defined in EBU Tech Doc 3341 [6] offers three distinct time scales:

- O Momentary Loudness (abbreviated "M") time window: 400ms
- O Short-term Loudness (abbreviated "S") time window: 3s
- O Integrated Loudness (abbreviated "I") from "start" to "stop"

The **M** and **S** time windows³ are to be used for the immediate levelling and mixing of audio signals. Initial level setting may be performed best with the Momentary Loudness Meter, adjusting the level of key or anchor elements (such as voice, music or sound effects) to be around the Target Level of **–23 LUFS**. It is advisable to set levels with a bit of caution initially, as it is easier to gradually increase the integrated loudness level during a mix than to decrease it. Usually, a slight increase in the course of a programme is also more natural – and an initially "defensive" strategy leaves the engineer room to manoeuvre in case of unexpected or unpredictable signals and events.

Once levels are set, the audio engineer can switch to **mixing only by ear**. Watching the Momentary or Short-term Loudness, and an occasional glance at the Integrated Loudness value, should give confirmation that the mix is within the tolerance allowed around the Target Level. With a numerical readout of the I-value, with one decimal point precision, or a graphical display of similar resolution, trends can be anticipated and the appropriate countermeasures taken.

To summarize ... improving the audio metering by replacing the PPM with the loudness meter is a step closer to the best measurement tool – the human ear.

Loudness in the distribution chain

EBU Tech Doc 3344 [5] specifies the relevant settings and processing of the audio after it has left the broadcast centre and takes the loudness paradigm all the way to the consumers' equipment, including set-top boxes and AV receivers. In doing so, it also encourages non-compliant broadcasts to become compatible with EBU R 128. Digital transmissions, analogue transmissions, rebroadcast transmissions, locally-inserted advertisements, newly-added services and much more are covered in this comprehensive document.

Rather than measuring the loudness of individual programmes, the distribution company monitors a service over 24 hours, with special care being taken when considering switched or shared services. As well as the programmes, any loudness metadata provided with the digital services is monitored so that the actual loudness level of the service may be compared with the stated level, both of which should be **–23 LUFS** of course! Once a day, the data is analyzed and where a service is more than 0.9 LU from its Target Level (the exact value and method is still under discussion), a steering unit applies a correction factor so that the long-term average of the audio remains at the Target Level ±1 LU.

Business consequences

Because the loudness-levelling paradigm affects all stages of an audio broadcast signal, from acquisition to distribution and transmission and, because the ultimate goal is to harmonise audio loudness levels within broadcast channels as well as between channels to achieve an equal universal loud-

 [&]quot;M" and "S" are commonly used in stereophony for "Mid" and "Side". To distinguish the integration times
"Momentary" and "Short-term", the versions "ML_K" and "SL_K" (as well as "IL_K") may be used. "L_K" stands
for "Level, K-weighted", and complies with the international naming standard ISO 80000-8.

ness level for the benefit of the listeners, all audio professionals and audio-metering equipment in all parts of the chain will be affected by this change.

For many, a major question will be whether all the existing QPPM meters will have to be replaced and whether all the relevant personnel will need instruction to accommodate the new way of working. The answer is "In the long run, undoubtedly, yes", but this transition needn't happen all at once, although at least some loudness meters should be brought into service at the earliest opportunity, to sit alongside the QPPMs. The replacement of meters can happen in line with the technical refresh cycles of other equipment, refurbishing of production facilities or as separate step-by-step projects (whatever is feasible on a reasonably short time scale), and people can be trained at the appropriate time.

Those responsible for the purchase of equipment should also be aware that safety limiters to prevent over-modulation will have to be able to work in *true-peak mode* and they will need to be adjusted to the appropriate maximum permitted true peak level, in production as well as at the output of master control, at the distribution head-end and at the transmitter site.

Conclusions

EBU R 128 and the four supporting documents provide a means to end the "Loudness War" – at last! The use of audio dynamics becomes a creative art again. There are still things to learn and people will take a while to get used to the new ways of working, but it will be worth the effort.

Over 230 participants have joined the EBU PLOUD group (as at August 2010), the e-mail reflectors have shown a level of activity never seen before and meter manufacturers have prepared units for display at IBC 2010 before the specification was even published.

It is now time to put R 128 into action!



The EBU R 128 logo

Acknowledgements

The author especially wants to thank Ian Rudd, independent technology strategist and media consultant, for his substantial editing contribution, as well as Frans de Jong, EBU, and Andrew Mason, BBC R&D, for valuable input to this article.

References

[1] EBU Technical Recommendation R 128: Loudness normalisation and permitted maximum level of audio signals



Florian Camerer is a senior sound engineer at ORF, the Austrian Broadcasting Corporation, in Vienna. He started work in the area of production sound, expanding to sound editing and mixing. His main field was documentaries, where he also developed a special interest in surround sound. He mixed the first documentary of the ORF in 5.1 multichannel audio in 1995 and has been active in the area of surround sound ever since, helping ORF to become the first European broadcaster to transmit a live 5.1 surround sound signal (New Year's Concert, 2003)

In 2008, Mr Camerer asked the EBU to create a working group to study loudness issues which led to the group PLOUD which he chairs. PLOUD is the biggest and most active EBU group ever, which is reflected in the substantial output published.

Florian Camerer lectures on an international basis in surround sound and loudness matters.

- [2] ITU-R BS.1770: Algorithms to measure audio programme loudness and true-peak audio level
- [3] EBU Tech Doc 3343: Practical Guidelines for Production and Implementation in accordance with EBU Technical Recommendation R 128 publication awaited
- [4] EBU Tech Doc 3342: Loudness Range: A descriptor to supplement loudness normalisation in accordance with EBU R 128
- [5] EBU Tech Doc 3344: Practical Guidelines for Distribution of Programmes in accordance with EBU R 128 publication awaited
- [6] EBU Tech Doc 3341: Loudness Metering: 'EBU Mode' metering to supplement loudness normalisation in accordance with EBU R 128

This version: 6 September 2010

Published by the European Broadcasting Union, Geneva, Switzerland

ISSN: 1609-1469

Editeur Responsable: Lieven Vermaele

Editor: Mike Meyer

E-mail: tech@ebu.ch



The responsibility for views expressed in this article rests solely with the author