Loudness measurement according to EBU R-128

Fons ADRIAENSEN,
Casa della Musica,
Pzle. San Francesco 1,
43000 Parma (PR),
Italy,
fons@linuxaudio.org

Abstract

This paper introduces a new Linux application implementing the loudmess and level measuring algorithms proposed in recommendation R-128 of the European Broadcasting Union. The aim of this proposed standard is to ease the production of audio content having a defined subjective loudness and loudness range. The algorithms specified by R-128 and related standard documents and the rationale for them are explained. In the final sections of the paper some implementation issues are discussed.

Keywords

Loudness, metering, mastering, EBU

1 Introduction

Most radio listeners and TV viewers will probably agree that having to reach for the remote control to adjust the audio volume every so many seconds is a nuisance. Yet it happens all the time, and there are many reasons for this.

One of them is the nature of contemporary broadcast content, a large part of which consists of sequences of 'bumpers', commercials, previews and teasers of upcoming features, etc. All of these originate from different production sources and none of those is particularly interested in the final listener experience, let alone responsable for it.

In the 'old days' there would be a trained audio technician taking care of continuity and levels. Such a person would preview upcoming content, be familiar with the avaliable metering and monitoring equipment, and above all, use his/her ears. In the worst case anything out of order would be adjusted promptly.

Today the situation is very different. You won't find an audio technician in a typical TV continuity studio - more often than not audio is slaved to the video switcher and there are no level adjustments at all. For radio in many cases the presenter takes care of everything (or tries to), and much play-out is just automated without any human monitoring it.

Even if a recording is made by an audio technician knowing his business there will be problems. Imagine you are recording a talk with studio guest that will used later 'as live' in some program with the same presenter. You know the presenter will just click the start button and the interview will be played out without any level adjustments. At what level should you record it ? The same problem occurs nearly all the time, for the simple reason that so much content is first produced 'out of context' and later used blindly and without any consideration of the context.

Broadcasters are aware of the problem but don't have the means to do much about it. Most large organisations have technical guidelines which may or may not be followed for in-house production, and with varying degrees of succes. Smaller ones usually just don't care. And all of them are to some degree involved in the 'loudness war', and forced to increase levels rather than control them.

What is missing is some standard way to determine the 'loudness' of audio content, and one which can be easily automated. Current metering systems are of little use for this, as will be seen in later sections.

Given such a standard, it would be possible to define target loudness levels for any particular type of content or program. Audio technicians would know what to do when recording, and automated systems would be able to 'measure' the loudness of audio files and store the result in a database (or add it to the file as metadata) for use during play-out. Consumer playback systems could do the same. This could even lead to a much needed improvement in music recording practices: if music producers know that broadcasters and playback systems will adjust the level of their records anyway, there is no more reason to push it up using the absurd amounts of agressive compression we see today.

2 An overview of current level metering practice

A number of quite different audio level measuring methods are being used today. Most of them do not provide any reliable measure of subjective loudness.

2.1 VU meters

The VU meter was designed in the days when audio equipment used tubes[1] and therefore could use only simple electronics, at most an amplifier stage to drive the passive meter. But it was quite strictly specified.

A *real* VU meter, as opposed to something just looking like one,[2] indicates the average of the absolute value of the signal (which is *not* the RMS level). For a fixed level signal, it should rise to 99% of the final value in 300 ms, and overshoot it by 1 to 1.5% before falling back. The small overshoot may seem a detail but it isn't — it has quite a marked effect on the actual ballistics. These are determined not by any electronics but only by the properties of the moving-coil meter which is a classic mass + spring + damping system equivalent to a second order lowpass filter.

A real VU meter does provide *some* indication of loudness, but not a very accurate one in practice. Apart from that its dynamic range is quite limited.

2.2 Digital peak meters

These indicate the peak sample value, with a short holding time and/or a slow fallback. They are found in most consumer and semi-pro equipment and in almost all audio software. They can be useful to indicate signal presence and check digital recording levels for clipping, but they provide no useful loudness indication at all. And in fact even as peak meters most of them fail, as the real peaks are almost always between the samples.

2.3 Pseudo peak meters

A PPM, also known as 'peak program meter' is driven by the absolute value of the signal (again not RMS), but with a controlled rise time (usually 10 ms, sometimes 5) and a slow fallback. This is the most popular type of meter in broadcasting (at least in Europe), and in many professional environments. Specifications are provided by various organisational or international standards.

A pseudo peak meter *can* provide some idea of loudness, but only to an experienced user. The reason is that the relation between indicated level and effective loudness depends very much on the type of content, and some interpretation is required. This makes this type of metering algorithm unsuitable for automated loudness measurement.

2.4 Bob Katz' K-meter

This is a relatively new way to measure and display audio levels, proposed by mastering expert Bob Katz. It displays both the digital peak and RMS values on the same scale. Since for normal program material the RMS level will always be lower than the peak value, and the intended use is based on controlling the RMS level (with the peak indication only as a check for clipping) the reference '0 dB' level is moved to either 20 or 14 dB below digital full scale. The ballistics are not specified in detail by Katz, but they should correspond to a simple linear lowpass filter, and not use different rise and fall times as for a PPM. Typical implementations use a response speed similar to a VU meter.

The K-meter provides quite a good indication of loudness, mainly because it uses the true RMS value, and because its response is not too fast. One way to improve this would be to add some filtering, and this is indeed what is done in the system discussed in the next sections.

2.5 Discussion

It should be clear that with the possible exception of the K-meter (which is not as widely used as it should be), current audio level metering systems provide a rather poor indication of actual subjective loudness.

Another issue is that all these level measurement systems were designed for interactive use, and only provide a 'momentary' level indication. What is really needed is a way to automatically determine the *average loudness* of a recording, e.g. a complete song, in a reliable way and without requiring human interpretation.

Apart from such an average loudness value, another one of interest is the subjective *loudness range* of some program material — how much difference there is between the softer and louder parts. This value could for example guide the decision to apply (or not) some compression, depending on the listening conditions of the target audience.

Surprisingly few application or plugins for loudness measurement are available and widely

---

[1]or *valves* for some of us
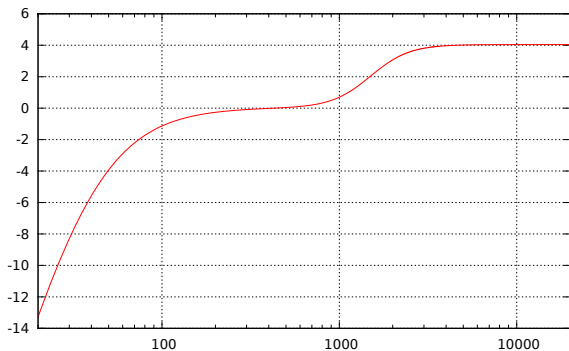[2]as do most of them

Figure 1: The K-filter response

used. The application presented at the LAC in 2007 [[Cabrera, 2007]] seems not to be actively developed anymore.

In the commercial world a notable exception is Dolby's *LM100*. Early versions only supported L_{eq} based measurements, while recent releases also offer a mode based on the ITU algorithm discussed in the next section.

3 The ITU-R BS.1770 loudness algorithm

The EBU R-128 recommendation (discussed in the following section) is based on the loudness measurement algorithm defined in [[ITU, 2006a]]. This specification is the result of research conducted in several places over the past 10 years. Listening tests using hundreds of carefully chosen program fragments have shown a very good correlation between subjective loudness and the output of this algorithm. Details and more references can be found in the ITU document.

The ITU recommendation specifies the use a weighting filter followed by a mean square averaging detector. The filter response is shown in fig.1 and is the combination of a second order highpass filter (the same as in the L_{eq}(RLB) standard), and a second order shelf filter. The latter is added to model head diffraction effects. The combination of the two filters is called the *K-filter*[3]

For multichannel use the mean squared values for each channel are multiplied by a weighting factor and added. This means that the powers are added and that inter-channel phase relationships have no effect on the result. For ITU 5.1

surround the weights for L,R and C are unity, +1.5 dB for the surround channels, and the LFE channel is not used. For stereo only L and R are used. In all cases there is just a single display for all channels combined — the idea is that a loudness meter would be used along with conventional per-channel level meters and not replace those.

The summed value is converted to dB, and a correction of -0.69 dB is added to allow for the non-unity gain of the K-filter at 1 kHz. This facilitates calibration and testing using standard 1 kHz signals. For a 0 dBFS, 1 kHz sine wave in either of L,R or C the result will be -3 dB. For the same signal in both L and R it will be 0 dB.

The ITU document does not specify if +3dB should be added when measuring a mono signal. Considering that a such a signal will in many cases be reproduced by two speakers (i.e. it is really the equivalent to a stereo signal with identical L and R) such a correction would seem to be necessary.

According to [[ITU, 2006a]] the output of a loudness measurement performed according to this algorithm should be designated 'LKFS' — Loudness using the K-filter, w.r.t. to Full Scale.

A second ITU document, [[ITU, 2006b]], provides some recommendations related to how loudness measurements should be displayed on a real device. In practice the LKFS scale is replaced by one that defines a 'zero' reference level at some point below full scale. To indicate that this is not a real level measurement the scale is marked in 'LU' (Loudness Units) instead of dB, with 1 LU being the same ratio as 1 dB. A linear range of at least -21 to +9 LU is recommended, but the reference level itself is not specified. This probably reflects the view that a different reference could be used in each application domain (e.g. film sound, music recording,...).

This document also recommends the use of an 'integrated' mode to measure the average loudness of a program fragment, but again does not specify the algorithm in any detail.

4 The EBU R-128 recommendation

Recommendation R-128 [[EBU, 2010a]] builds on ITU-R BS.1770 and defines some further parameters required for a practical standard. More detail is provided by two other EBU documents, [[EBU, 2010b]] and [[EBU, 2010c]]. Two more are in preparation but not yet available at

---

[3]This could result in some confusion with the 'K' from Bob Katz' name as used in 'K-meter'. There is no relation between the two.

the time of writing.

4.1 Reference level and display ranges

R-128 defines the reference level as -23 dB relative to full scale, i.e. a continuous 1 kHz sine wave in both L and R and 23 dB below clipping corresponds to standard loudness.

It also requires meters conforming to this standard to be able to display levels either relative to this reference and designated **LU**, or to full scale and designated **LUFS**. Two display ranges should be provided: one from -18 to +9 dB relative to the reference level, and the second from -36 to +18 dB. The user should at any time be able to switch between these four scales. This choice then applies to all displayed values.

4.2 Dynamic response: M,S,I

Three types of response should be provided by a loudness meter conforming to R-128:

The **M** (momentary) response is the mean squared level averaged over a rectangular window of 400 ms. An R-128 compliant meter should also be able to store and show the maximum of this measurement until reset by the user.

The **S** (short term) response is the mean squared level averaged over a rectangular window of 3 seconds. R-128 requires this to be updated at least ten times per second. No such value is specified for the **M** response, but it seems reasonable to assume that at least the same update rate would be required.

The **I** (integrated) response is an average over an extended period defined by the user using *Start, Stop* and *Reset* commands. It is detailed in the following section.

4.3 Integrated loudness

The integrated loudness measurement is intended to provide an indication of the average loudness over an extended period, e.g. a complete song.

It is based on the full history, within an interval specified by the user, of the levels used for the **M** response. The input to the integration algorithm should consist of measurements in 400 ms windows that overlap by at least 200 ms.

Given this input, the integrated loudness is computed in two steps. First the average power of all windows having a level of at least -70 dB is computed. This absolute threshold is used to remove periods of silence wich may occur e.g. at the start and end of a program segment. In a second step all points more than 8 dB below the first computed value are removed and the average power is recomputed. This second, relative threshold ensures that the integrated measurement is not dominated by long periods of relative silence as could occur in some types of program.

The result is the integrated loudness value, displayed as either **LU** or **LUFS** according to the scale selected by the user. This algorithm can be applied either in real time or on recorded audio. When a loudness meter is operating on real-time signals the indicated value should be updated at least once per second.

4.4 Loudness range, LRA

The purpose of the loudness range measurements is to determine the perceived dynamic range of a program fragment. This value can be used for example to determine if some compression would be necessary. The algorithm is designed to exclude the contribution of very short loud sounds (e.g. a gunshot in a movie), of short periods of relative silence (e.g. movie fragments with only low level ambient noises), and of a fade-in or fade-out.

The input to the LRA algorithm consists of the full history, within the same interval as for the integrated loudness, of the levels used for the **S** measurement. The windows used should overlap by at least 2 seconds.

First an absolute threshold of -70 dB is applied and the average value of the remaining windows is computed — this is similar to the first step for the integrated loudness (but using different input). A second threshold is then applied at 20 dB below the average value found in the first step. The lower limit of the loudness range is then found as the level that is exceeded by 90 percent of the remaining measurement windows, and the upper limit is the level exceeded by the highest 5 percent. In other words, the loudness range is the difference between the 10% and 95% percentiles of the distribution remaining after the second threshold.

4.5 True peak level

Both the ITU and EBU documents cited in previous sections recommend the use of true peak level indication in addition to loudness measurement.

Most peak meters just display the absolute value of the largest sample. There are two potential sources of error with this simple approach. First, almost all peaks occur between
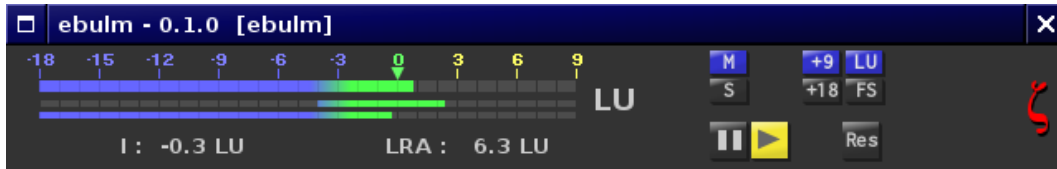
Figure 2:   The EBU mode meter

the samples. To reduce the error from failing to see these peaks, [[ITU, 2006a]] recommends to upsample the signal by a factor of at least four. Second, the peak level may be different if later stages in the processing include DC-blocking — in fact it could be either higher or lower. For this reason it is recommended to measure peak levels both with and without DC blocking, and display the highest value.

The EBU documents do not require a continuous display of peak levels. Instead they recommend the use of a true peak indicator with a threshold of 1 dB below the digital peak level.

5   Implementation

The *ebulm* application (fig.2) is written as a Jack client. The upper bargraph shows either the **M** or **S** response. The two thinner ones below display the loudness range and the integrated loudness which are also shown in numerical form. To the right are some buttons that control the display range and scale, and below these the controls to stop, start and reset the integrated measurements.

Two features are missing in this version (but will be added): the display of the maximum value of the **M** response, and the true peak indicator.

The ITU document specifies the K-filter as two biquad sections and provides coefficients only for a sample rate of 48 kHz, adding that implementations supporting other rates should use 'coefficients that provide the same frequency response'. It is in general not possible to create exactly the same FR using a biquad at different rates, but the code used in *ebulm* comes close: errors are less than 0.01 dB at 44.1 kHz and much less at higher rates. Another peculiarity is that the highpass filter has a double zero at 0 Hz as expected, but the nominator coefficients given for 48 kHz are just $+1, -2, +1$ instead of values that would provide unity passband gain. This has to be taken into account when using a different sample rate.

There is no specified limit on the lenght of

the integration period for the **I** and **LRA** measurements. A simple implementation of the algorithms would require unlimited storage size, and for the loudness range calculation the stored data would need to be sorted as well. The solution is to use histogram data structures instead — these require a fixed storage size, keep the data sorted implicitly, and make it easy to find the percentiles for the loudness range calculation. The current implementation uses two histograms, each having 751 bins and covering the range from -70 to +5 dB with a step of 0.1 dB. Points below -70 dB can be discarded, as the absolute threshold in both algorithms will remove them. Levels higher that +5 dB RMS over a 400 ms period mean the measurements will probably be invalid anyway. If such levels occur they are clipped to +5 dB and an error flag is set.

6   Acknowledgements

The work done by the EBU on loudness measurement and standardization is very much the result of an initiaitive taken by Florian Camerer, senior sound engineer at ORF and chairman of the PLOUD working group.

References

Andres Cabrera. 2007. Audio metering and linux. In *Proceedings of the Linux Audio Conference, 2007*. Available at `http://http://lac.linuxaudio.org/2007/papers`.

EBU. 2010a. *EBU Recommendation R-128, Loudness normalisation and permitted maximum level of audio signals*. EBU, available at `http://tech.ebu.ch/publications`.

EBU. 2010b. *EBU-TECH 3341, Loudness Metering: EBU Mode metering to supplement loudness normalisation in accordance with EBU R 128, Supplementary information*. EBU, available at `http://tech.ebu.ch/publications`.

EBU. 2010c. *EBU-TECH 3342, Loudness Range: A descriptor to supplement loud-*

ness normalisation in accordance with EBU R 128. EBU, available at `http://tech.ebu.ch/publications`.

ITU. 2006a. *ITU-R BS1770-1, Algorithms to measure audio programme loudness and true-peak audio level*. ITU, available at `http://www.itu.int/rec/R-REC-BS/e`.

ITU. 2006b. *ITU-R BS1771, Requirements for loudness and true-peak indicating meters*. ITU, available at `http://www.itu.int/rec/R-REC-BS/e`.