WHITE PAPER

Loudness Monitoring in TVM/VTM

The move to digital audio workflows over the past several years has produced many benefits, from clean multi-generation copies to the ability to easily implement digital effects. Many of the old analog problems, like certain noise and distortion issues, have been entirely eliminated.

Of course, not all audio processing problems have disappeared. There are still many ways to characterize audio levels — from VU (volume unit) to peak to true peak. None of these serves as a good measurement to represent the overall program volume to our human ears, generically referred to as "loudness."

Loudness measurement has been represented by different techniques, and work continues on this issue. One particular aspect of this problem is represented by different perceived audio levels across different portions of audio content — scene to scene, program to commercial, or channel to channel. This fundamental problem continues to vex audio content producers and broadcasters: how to make the average audio level seem consistent across a program, regardless of program type and content, and beyond that, consistent across channels.

This paper examines some of these loudness measurement techniques in detail, including some recent work by the ITU (published as ITU-R BS.1770). The problems with some past techniques are explained, and a new method of displaying audio loudness levels, consistent with ITU-R BS.1770, is described.

Different Methods to Describe Audio Levels
Audio loudness measurement fundamentally requires another method of defining audio levels — a "meter ballistic." There are, of course, other considerations that will be described, but since loudness only defines audio levels, some other common meter ballistics will be described first.

Most commonly used meter ballistics do a poor job of representing the average audio level perceived by a listener, since they were designed for some other purpose. Often, the main application of existing ballistics is to prevent saturation and distortion on a recording medium or transmission path. While certainly important, this is a very different application than determining loudness.

One of the most common audio measurement ballistics is the VU meter. This meter traces its roots back to the 1930s, so it is quite familiar to most audio producers. It is actually an early attempt to display an "average" audio level; this meter ballistic has an integration time of 300 ms, so it does not respond to fast audio peaks and troughs. This meter is defined by IEC 60268-17.

Another set of common audio ballistics attempts to show the peak excursions. The BBC developed its peak program meter (PPM) in the 1930s for this purpose. Technically, the PPM is a quasi-peak meter, since it uses a 10 ms integration time. The PPM is defined by IEC 60268-10.

True peak meters are commonly used with digital audio content. Unlike the PPM, this meter ballistic has zero integration time, responding to every digital audio sample.

These and other common meter ballistics can lead to widely varying loudness levels when used to monitor different types of audio content. The dynamic range of any audio content has a major impact on the perceived loudness, but is not reflected well by any commonly used meter ballistic. Peak reading meter types can give especially erroneous results, since there is (by definition) no attempt to average the audio level. The figure below shows an example.

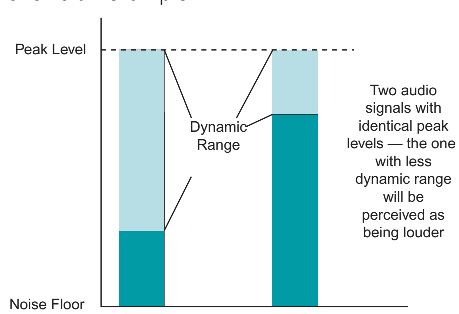

However, none of these ballistics accurately reflects the average level perceived by human ears. This loudness quantity is as important as other audio level measurements in terms of listening experience for the consumer.

Existing Methods to Measure Loudness

The industry has made previous attempts to quantify this perceived quantity of loudness. One of the earliest was the "CBS Loudness" measure, which was described in 1982. This was an attempt at making a perceptual, as opposed to a purely electrical, measurement. For many years, this was as close to a reference loudness specification as existed.

Dialog normalization is a metadata parameter in compressed audio content that is designed specifically to control dialog levels across channels at the consumer premises. This is a valid concept, but there have been implementation issues that so far have prevented it from being as effective as intended.

There are multiple steps needed to properly employ dialog normalization. The first step still requires a reading of the perceived loudness of the audio, regardless of content (with or without spoken dialog, for instance). This currently requires proprietary hardware and algorithms to compute. There have been questions about the proprietary means used to define the average dialog level with this technique.

This proprietary process outputs a dialog normalization level in dBFS, which then needs to be entered into the encoder so that it may be carried in the compressed audio stream as metadata. This is, however, one common place where the process fails — many encoders in use still have their default dialog normalization level being encoded regardless of what the actual value may be. Needless to say, this does not yield correct results.

A simple mean-squared power measure has also been developed, commonly labeled Leq. This measurement can be applied unweighted (usually just referred to as Leq) or can use a weighting filter before the power measurement. The type of filter is typically indicated along with the measurement means, such as Leq(RLB). Studies by the ITU have shown that Leq(RLB) actually matches perceived loudness quite well for monophonic material. This measurement means forms the basis of the ITU loudness method described in ITU-R BS.1770. This specification was developed to address some of these existing problems with other loudness methods (it also defines a new true peak measure as well, although that will not be discussed here).

This document was based on the results of extensive perceptual trials that tried to match an electronic loudness response to many different sets of listener results. The ITU has extended the Leq(RLB) ballistic, which proved effective with monophonic content, to work with multichannel content. To do this required adding an additional hi-pass filter to the already existing RLB response.

A key feature of the BS.1770 measurement technique is its multichannel capability. To achieve this, it uses the sum of five channels of a surround sound mix (the LFE channel is not used). All five channels pass through the same filtering and power summing process, with their individual results being scaled and then summed together for a final result. A block diagram of the measurement technique is shown below.

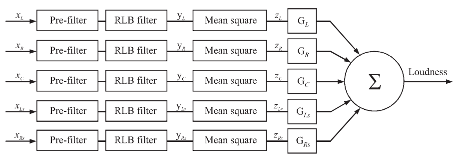

The pre-filter block is the hi-pass filter mentioned previously, which is then followed by an RLB response. The Leq (mean square) level measurement is next, followed by gain blocks (Gx) that account for the relative angle of each surround channel. Finally, the scaled results of each channel are summed.

Extensive tests by the ITU have proven that this technique has a strong correlation to perceived loudness across different program materials and across different listener groups.

BS.1770 Loudness Implementation in Modern Test Equipment

BS.1770 describes an effective measurement technique for determining audio program loudness, but the issue of how to effectively display and use that information remains unresolved.

Since any loudness measurement technique is fundamentally a meter ballistic, displaying the individual channel loudness results as standard audio meters conveys this information in a readily understood manner. Audio levels are still labeled in dBFS, and the meters themselves should be labeled in "LU" or loudness units. This differentiates these meters from "VU" or peak reading level meters.

Looking at loudness channel-by-channel over a short time span does not really tell much about the perceived loudness of the entire audio program, though. A longer-term display method is required for this, one that will also show the BS.1770 mandated sum of the surround channels.

Harris' Videotek® test and measurement equipment has recently added these features to its TVM and VTM Series™ audio hardware as well as its QuiC™ media analysis server. A screen capture of a BS.1770 loudness analysis screen is shown below.

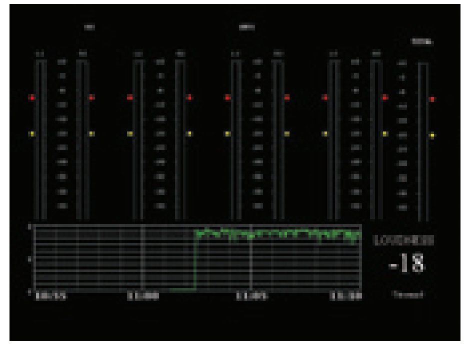

The first six audio meters each display one of the surround channels after proper filtering and mean squared level calculation, as defined by BS.1770. The right-most meter displays the sum of those channels (except for LFE, as defined by the spec). The chart below the meters offers a greater understanding of loudness levels over particular program segments. This is a trend plot or strip chart recorder, displaying a history of the total program loudness level over a time period from 15 minutes to 24 hours — allowing monitoring of loudness on a program segment up to a transmission output level. To simplify loudness monitoring even further, the average value of the historical loudness is displayed as a single numeric value.

Conclusion

Many useful audio ballistics exist, but none of them have accurately reflected how the human ear perceives the loudness of audio program content. Determined to rectify that situation, the ITU performed extensive research and from that created ITU-R BS.1770, a specification that defines a technique to monitor multichannel audio content for loudness.

Putting this specification into practice, the Harris Videotek® TVM, VTM Series™, and QuiC™ test and measurement equipment now offer audio loudness based on BS.1770. Harris extends the utility of this specification by allowing short-term, channel-by-channel examination of loudness as well as historical loudness trending, allowing loudness results to be viewed by program segment or the entire day of programming.

For more information please visit www.broadcast.harris.com.

HARRIS®
assuredcommunications®

Broadcast Communications Division
4393 Digital Way | Mason, OH USA 45040 | Tel: 1 (513) 459 3400
www.broadcast.harris.com

WP_TVMVTM_LOUD_0407