Audio Loudness in the Age of Digital Media:

Simplified Solutions for Providing the Best Listening Experience

Introduction

Loudness management is a vital step for content creators. Audio loudness control has evolved to include not only broadcast but various forms of digital media, such as streaming services—from OTT to music streaming platforms. Maintaining consistent loudness across all of the programs and services becoming available is a vital step for a better experience for content consumers. In this white paper, Cornelius "Corny" Gould, Audio Processing Architect of Angry Audio, explores optimal real-time adjustments of various types of audio signals, allowing broadcasters to get consistent audio levels across content, format, and platform.

ANGRYAUDIO

www.AngryAudio.com

Background

Broadcasters are very familiar with "normalization" of loudness levels, which gives consumers consistent loudness levels for a better and more enjoyable audio experience that ultimately makes them more inclined to engage with your content. In a pure broadcasting model, it is simple to control audio levels because programming, which is all under the broadcaster's control, is simply routed through an automatic level control device—aka, an audio processor. All the audio sources are fed into the processor in real time, creating consistent loudness even with dissimilar content.

However, in today's digital media era, the task of audio level normalization is a complicated challenge! Every video program and every song delivered on most streaming services originates from an audio or video file that is delivered directly to the audience to consume, bypassing the broadcast engineer who would normally tweak audio levels to be more palatable to consumers.

This means that each file delivered may have completely different audio levels, which becomes an irritation for the consumer who is constantly reaching for the volume control as the content changes. In fact, inconsistent audio levels were such a problem in television broadcasting that congress passed the CALM (Commercial Advertisement Loudness Mitigation) Act, essentially creating a loudness standard for TV.

Loudness standards are a hot topic for everyone now, especially live streamers. In the absence of an official streaming standard for loudness, major platforms like Apple, Amazon, and YouTube are demanding adherence to their own self-imposed loudness criteria. If content creators don't comply, they don't get carried! Many audio-only content creators, such as podcasters, jumped onto this standard as well. Now, maintaining consistent audio levels has become the number-one best practice for podcast audio production.

Proposed new legislation extends the CALM Act to streaming TV content. The CALM Modernization Act of 2022 would update that original act to include video streaming services. "Consumers don't like loud commercials any more than they did in 2010, when the original CALM Act...was signed into law. And they don't distinguish between high volume commercials aired on traditional television platforms versus the many streaming video services accessed by consumers in 2022," said Jonathan Schwantes, Senior Policy Counsel of Consumer Reports.

Until the new CALM act is passed, all of the processing solutions currently in place for streaming revolve around a standard system of measuring audio content's subjective or perceived loudness, then convincing content creators to voluntarily adhere to a unified loudness level standard for the sake of a better experience for listeners.

Understanding Loudness Units

dB, LUFS, LKFS...oh my! As we move into the digital era, it's important for broadcasters to understand the language of loudness control so they can adhere to different services' loudness guidelines. The standard method of measuring sensational loudness is the ITU BS.1770 loudness meter, which is designed around the BS.1770-3 loudness standard. Each entertainment delivery service today typically has a "target" BS.1770 loudness level to which all content must adhere.

The loudness level is referred to as LUFS, or Loudness Units (relative to) Full Scale. The term "LUFS" is also sometimes referred to as "LKFS," or "Loudness Units, K-Weighted, relative to Full Scale." Both acronyms mean the same thing. Just a slight difference in naming convention. Each "loudness unit" is a number that corresponds to a perceived loudness level. Similar in concept to the Volume Unit of level as expressed on standard VU style meters.

Keep in mind, LUFS is not equivalent to "dB." LUFS and dB are entirely different units of measurement. Some have the misconception that loudness normalization can be determined using the standard VU (dB) meter, but loudness measurements can only properly be determined using a LUFS meter.

What makes a LUFS meter so special? Generally speaking—and this is not a "lab coat" description—the LUFS meter tracks two variables to provide accurate representation of perceived loudness: filtered level and time.

Filtered Level: Humans do not experience loudness the same across the audio spectrum. We can tolerate loud bass far easier than high frequencies, such as cymbals and electric guitars. A LUFS meter uses a K-Weighted filter to shape the signal according to perceived loudness sensitivity. The filter is also somewhat deaf to bass energy. The lower the bass frequency, the less it matters for measurements.

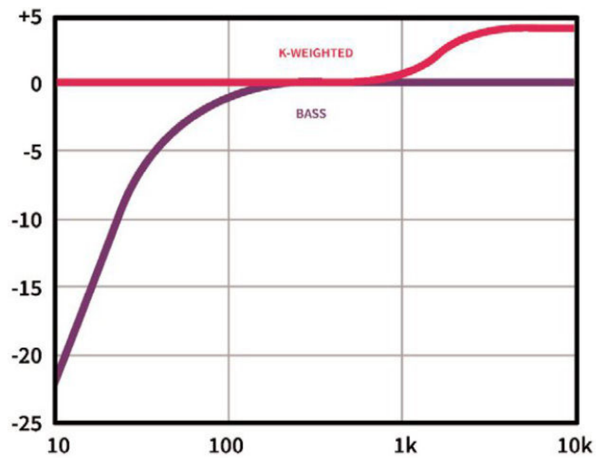

Time: The longer the duration of loud audio, the more we notice it. A single drum hit, for example, might be many times louder than a shrieking heavy metal guitar, but its extremely short duration means we don't perceive it as very loud. With this in mind, the LUFS meter has several methods for measuring the temporal effect of audio content levels.

The LUFS meter averages filtered audio levels in 400ms blocks. Each block is then applied to two additional averaging systems. The first one averages the blocks over a three-second period, and then uses the resulting data to display short-term loudness. The second is a long-term measurement which can be a fixed time window or will run until reset, depending on the meter and application.

For file-based loudness measurements, the entire file is read at faster than real time, and the long-term value is the loudness average across the entire audio file. For real-time, or live, measurements, some meters feature a rolling long-term average, where the end user can define a window from a few seconds to many minutes. The rolling average is very useful when used in live streaming applications.

the measuring differences between mono and stereo audio content.

Many loudness meters, such as the YouLean loudness meter, accept pre-produced files and give you the data you need. Other meters such as the

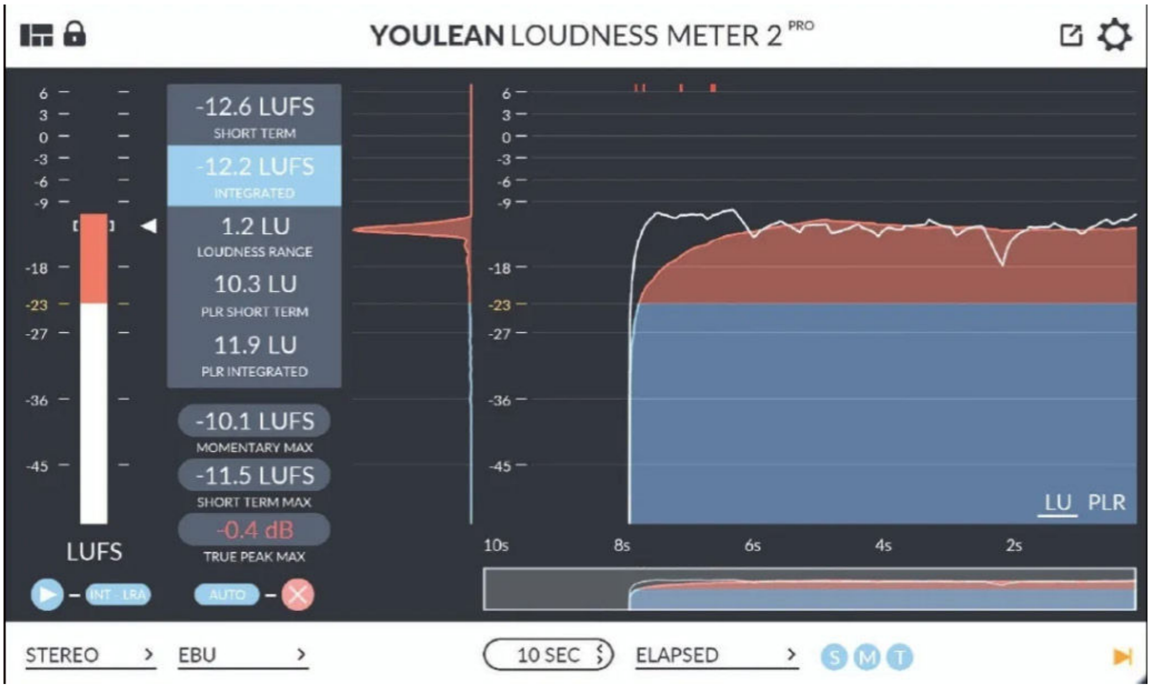

The Youlean meter is very popular for evaluating audio loudness.

The closer the meter gets to “0 LUFS,” the higher the sensational loudness. To achieve levels above -13 LUFS typically involves the heavy use of compression and limiting, and is generally frowned upon by audio purists. You will rarely see any recommendations above -14 LUFS.

For audio services featuring stereo tracks (like music), the LUFS levels mentioned in the above paragraph would represent maximum LUFS. When mono content (like speech) is present, the reference level is typically dropped by roughly 3 LUFS to compensate for

Orban Loudness Meter and the TC Electronic Loudness Radar can be used for audio files too but are very useful for real-time monitoring applications as they both feature a user-defined real-time rolling average.

True Peak

Loudness measurements are a bit worthless without defining an actual audio level. This is where the standard (dB) level measurement comes into the scene in the form of the True Peak meter.

True Peak meters are specialized level meters that show the actual peak level at the consumer's end of the chain. This might differ from what the producer's peak level meter shows because of certain errors that may occur with the reproduction of peak levels at the receive end. These errors can happen typically due to intersample peaks—peaks that occur between audio samples.

Many peak limiter designs can miss these peaks as they are not "seen" at the producer's side. On the consumer side, the reconstruction of digital audio data to analog can re-create those missing peaks due to the filtering processes used during conversion.

True peak meters are designed to show the effect of waveform reconstruction at the consumer end. This gives the content producer the ability to provide enough "headroom" in audio levels to prevent possible distortion on consumer devices.

With the above in mind, you will see that all loudness standards will also define a true peak specification in "dBFS," which defines the audio decibel level relative to full scale on the familiar dB peak-reading style VU meter.

How LUFS is used depends on your intended target. Amazon, for example, defines a -14 LUFS loudness target with a -2dB dBFS peak value for music and podcasts sent to its smart speakers and devices.

As of the writing of this article, some services are using this -14 LUFS loudness standard as well, while others are using a -16 LUFS. Other applications, such as inter-studio content delivery are using lower loudness levels entirely. For example:

- Broadcast TV: -23 LUFS
- Podcasts: about -16 LUFS
- Apple Music: -14 LUFS
- Amazon Music - 14 LUFS
- YouTube: - 14 LUFS
- Streaming Audio services (proposed by AES): -16 LUFS

For the distribution of content between studios, a different scale is used, which ranges from -23 to -18 LUFS, depending on the service used. File distribution for content consumers (the listening or viewing public) has been converging to -16 LUFS as a standard.

Dynamic Solutions for Streaming Loudness Leveling

Radio engineers are stretched very thin these days, and don't have time to tune and tweak their audio processing to comply with various content distributors' guidelines. I developed Angry Audio's Chameleon audio processing software, found in the Angry Audio Audio Chameleon C4 Livestream Audio Processor, to address this constraint. The C4 has unique architecture that combines four intelligent compression bands with a precision loudness controller, all orchestrated by an advanced form of artificial intelligence (AI) unsigned highly trained models for optimal real-time adjustments of any type of signals. The system continually monitors incoming audio and adjusts its parameters dynamically to fit. The technology eliminates the cost and complexity that prohibit many stations from processing their streams and HD channels. Users can set a target anywhere between -24 LUFS and -14 LUFS according to the required destination specifications. (It also supports the proposed AES loudness standard.)

Conclusion

Radio broadcasters understand that having consistent audio levels across content, format, and delivery method is extremely important to attract, delight, and retain listeners. The cost and complexity of processing streams and HD channels can be intimidating. Solutions that let users set targets within predetermined loudness ranges and continually monitor incoming audio, adjusts its parameters dynamically, help save broadcasters time, money, and make complex audio problems simple. Visit AngryAudio.com for more information on streaming solutions.

About the Author

Cornelius “Corny” Gould has been developing audio processors for decades. He is renowned for his groundbreaking work on the Telos Alliance Omnia.11 audio processor and is the mind behind Chameleon audio processing technology, Angry Audio’s innovative approach to processing for studios and streams.

Corny is also a member of the AES Broadcast and Online delivery committee, influencing and educating content producers on loudness standards and best practices for streaming audio.