

Audio Signal Visualisation and Measurement

Robin Gareus

Université Paris 8, linuxaudio.org
robin@gareus.org

Chris Goddard

Freelance audio engineer
chris@oofus.co.uk

ABSTRACT

The authors offer an introductory walk-through of professional audio signal measurement and visualisation using free software.

Many users of audio software at some time face problems that requires reliable measurement. The presentation focuses on the SiSco.lv2 (Simple Audio Signal Oscilloscope) and the Meters.lv2 (Audio Level Meters) LV2 plugins, which have been developed open-source since August 2013. The plugin bundle is a super-set, built upon existing tools adding novel GUIs (e.g. ebur128, jmeters,...), and features new meter-types and visualisations unprecedented on GNU/Linux (e.g. true-peak, phase-wheel,...). Various meter-types are demonstrated and the motivation for using them explained.

The accompanying documentation provides an overview of instrumentation tools and measurement standards in general, emphasising the requirement to provide a reliable and standardised way to measure signals.

The talk is aimed at developers who validate DSP during development, as well as sound-engineers who mix and master according to commercial constraints.

1. INTRODUCTION

An engineer reading and using audio level meters is comparable to a musician reading or writing sheet-music. Also, just like there are virtuoso musicians who can't read a single note, there are great sound-engineers who just go by their ears and produce great mixes and masters without ever looking at a single meter (although this doesn't imply that the use of meters is the option for those who can't hear).

Audio level meters are very powerful tools that are useful in every part of the production chain:

- When tracking, meters are used to ensure that input signals do not overload and maintain reasonable headroom.
- Meters offer a quick visual indication of activity when working with a large number of tracks.
- During mixing, meters provide a rough estimate of the loudness of each track.

- At the mastering stage, meters are used to check compliance with upstream level and loudness standards, and to optimise the dynamic range for a given medium.

Similarly for technical engineers, reliable measurement tools are indispensable for the quality assurance of audio-effects or any professional audio-equipment.

2. METER TYPES AND STANDARDS

For historical and commercial reasons various measurement standards exist. They fall into three basic categories:

- Focus on **medium**: highlight digital number, or analogue level constraints.
- Focus on **message**: provide a general indication of loudness as perceived by humans.
- Focus on **interoperability**: strict specification for broadcast.

For in-depth information about metering standards, their history and practical use, please see [1] and [2].

2.1 Digital peak-meters

A Digital Peak Meter (DPM) displays the absolute maximum signal of the raw samples in the PCM signal (for a given time). It is commonly used when tracking to make sure the recorded audio never clips. To that end, DPMs are calibrated to 0dBFS (Decibels relative to Full Scale), or the maximum level that can be represented digitally in a given system. This value has no musical connection whatsoever and depends only on the properties of the signal chain or target medium. There are conventions for fall-off-time and peak-hold, but no exact specifications. Furthermore, DPMs operate on raw digital sample data which does not take inter-sample peaks into account, see section 2.7.

2.2 RMS meters

An RMS (Root Mean Square) type meter is an averaging meter that looks at the energy in the signal. It provides a general indication of loudness as perceived by humans.

Bar-graph RMS meters often include an additional DPM indicator for practical reasons. The latter shows medium specifics and gives an indication of the crest-factor (peak-to-average power ratio) when compared to the RMS meter.

Similar to DPM's, there is no fixed standard regarding ballistics and alignment level for a general RMS meter, but




Figure 1. Various meter alignment levels as specified by the IEC. Common reference level is 0dBu calibrated to -18dBFS for all types except for DIN, which aligns +9dBu to -9dBFS. dBu refers to voltage in an analogue system while dBFS to digital signal full-scale.

various conventions do exist, most notably the K-system introduced by Bob Katz [3].

2.3 IEC PPMs

IEC (International Electrotechnical Commission) type Peak Programme Meters (PPM) are a mix between DPMs and RMS meters, created mainly for the purpose of interoperability. Many national and institutional varieties exist: European Broadcasting Union (EBU), British Broadcasting Corporation (BBC), Deutsche Industrie-Norm (DIN),... [4].

These loudness and metering standards provide a common point of reference which is used by broadcasters in particular so that the interchange of material is uniform across their sphere of influence, regardless of the equipment used to play it back. See Fig. 1 for an overview of reference levels.

For home recording, there is no real need for this level of interoperability, and these meters are only strictly required when working in or with the broadcast industry. However, IEC-type meters have certain characteristics (rise-time, ballistics) that make them useful outside the context of broadcast.

Their specification is very exact [5], and consequently, there are no customisable parameters.

2.4 EBU R-128

The European Broadcast Union recommendation 128 is a rather new standard, that goes beyond the audio-levelling paradigm of PPMs.



Figure 2. Various meter-types from the meter.lv2 plugin bundle fed with a -18 dBFS 1 kHz sine wave. Note, bottom right depicts the stereo phase correlation meter of a mono signal.

It is based on the ITU-R BS.1770 loudness algorithm [6] which defines a weighting filter amongst other details to deal with multi-channel loudness measurements. To differentiate it from level measurement the ITU and EBU introduced a new term ‘LU’ (Loudness Unit) equivalent to one Decibel¹. The term ‘LUFS’ is then used to indicate Loudness Unit relative to full scale.

In addition to the average loudness of a programme the EBU recommends that the ‘Loudness Range’ and ‘Maximum True Peak Level’ be measured and used for the normalisation of audio signals [7].

The target level for audio is defined as -23 LUFS and the maximum permitted true-peak level of a programme during production shall be -1 dBTP.

The integrated loudness measurement is intended to quantify the average program loudness over an extended period of time, usually a complete song or an entire spoken-word feature. [8], [9].

Many implementations go beyond displaying range and include a history and histogram of the Loudness Range in the visual readout. This addition comes at no extra cost because the algorithm to calculate the range mandates keeping track of a signal’s history to some extent.

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

- **Momentary** response. The mean squared level over a window of 400ms.
- **Short term** response. The average over 3 seconds.
- **Integrated** response. An average over an extended period.

¹ the ITU specs uses ‘LKFS’, Loudness using the K-Filter, with respect to Full Scale, which is exactly identical to ‘LUFS’.

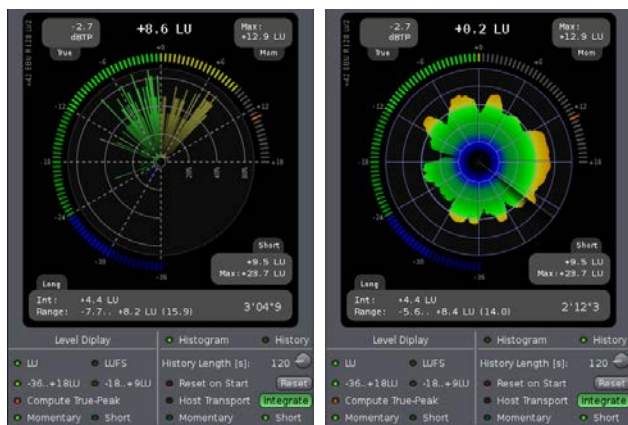


Figure 3. EBU R-128 meter GUI with histogram (left) and history (right) view.

2.5 VU meters

Volume Unit (VU) meters are the dinosaurs (1939) amongst meters.

The VU-meter (intentionally) "slows" measurement, averaging out peaks and troughs of short duration, and reflects more the perceived loudness of the material [10], and as such was intended to help program producers create consistent loudness amongst broadcast program elements.

In contrast to all the previously mentioned types, VU meters use a linear scale (in 1939 logarithmic amplifiers were physically large). The meter's designers assumed that a recording medium with at least 10 dB headroom over 0 VU would be used and the ballistics were designed to "look good" with the spoken word.

Their specification is very strict (300ms rise-time, 1 - 1.5% overshoot, flat frequency response), but various national conventions exist for the 0VU alignment reference level. The most commonly used was standardised in 1942 in ASA C16-5-1942: "The reading shall be 0 VU for an AC voltage equal to 1.228 Volts RMS across a 600 Ohm resistance"²

2.6 Phase Meters

A phase-meter shows the amount of phase difference in a pair of correlated signals. It allows the sound technician to adjust for optimal stereo and to diagnose mistakes such as an inverted signal. Furthermore it provides an indication of mono-compatibility, and possible phase-cancellation that takes place when a stereo-signal is mixed down to mono.

2.6.1 Stereo Phase Correlation Meters

Stereo Phase Correlation Meters are usually needle style meters, showing the phase from 0 to 180 degrees. There is no distinction between 90 and 270 degree phase-shifts since they produce the same amount of phase cancellation. The 0 point is sometimes labelled "+1", and the 180 degree out-of-phase point "-1".

² This corresponds to +4dBu

2.6.2 Goniometer

A Goniometer plots the signal on a two-dimensional area so that the correlation between the two audio channels becomes visually apparent (example in Fig. 8). The principle is also known as Lissajous curves or X-Y mode in oscilloscopes. The goniometer proves useful because it provides very dense information in an analogue and surprisingly intuitive form: From the display, one can get a good feel for the audio levels for each channel, the amount of stereo and its compatibility as a mono signal, even to some degree what frequencies are contained in the signal. Experts may even be able to determine the probable arrangement of microphones when the signal was recorded.

2.6.3 Phase/Frequency Wheel

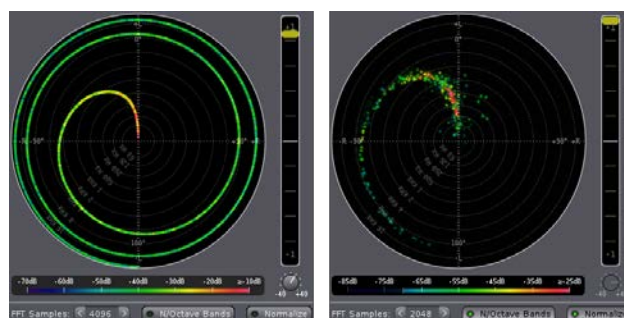


Figure 4. Phase/Frequency Wheel. Left: pink noise, 48KSPS with right-channel delayed by 5 samples relative to left channel. Right: Digitalisation of a mono 1/2" tape reel with slight head misalignment.

The Phase Wheel is an extrapolation of the Phase Meter. It displays the full 360 degree signal phase and separates the signal phase by frequency. It is a rather technical tool useful, for example, for aligning tape heads, see Fig. 4

2.7 Digital True-Peak Meters

A True-Peak Meter is a digital peak meter with additional data pre-processing. The audio-signal is up-sampled (usually by a factor of four [6]) to take inter-sample peaks into account. Even though the DPM uses an identical scale, true-peak meters use the unit dBTP (decibels relative to full scale, measured as a true-peak value – instead of dBFS). dBTP is identical to dBFS except that it may be larger than zero (full-scale) to indicate peaks.

Inter-sample peaks are not a problem while remaining in the digital domain, they can however introduce clipping artefacts or distortion once the signal is converted back to an analogue signal.

Fig. 5 illustrates the issue. Inter-sample peaks are one of the important factors that necessitate the existence and usage of headroom in the various standards, Table 1 provides a few examples of where traditional meters will fail to detect clipping of the analogue signal.

floating point audio data	mathematical true peak value
.. 0 0 +1 +1 0 0 ..	+2.0982 dBTP
.. 0 0 +1 -1 0 0 ..	+0.7655 dBTP

Table 1. True Peak calculations @ 44.1 KSPS, both examples correspond to 0dBFS.




Figure 5. Inter-sample peaks in a sine-wave. The red line (top and bottom) indicates the digital peak, the actual analogue sine-wave (black) corresponding to the sampled data (blue dot) exceeds this level.

2.8 Spectrum Analysers

Spectrum analysers measure the magnitude of an input signal versus frequency. By analysing the spectra of electrical signals, dominant frequency, power, distortion, harmonics, bandwidth, and other spectral components can be observed. These are not easily detectable in time domain waveforms.

Traditionally they are a combination of band-pass filters and an RMS signal level meter per band which measures the signal-power for a discrete frequency band of the spectrum. This is a simple form of a perceptual meter. A well known specification is a 1/3 octave 30-band spectrum analyser standardised in IEC 61260 [11]. Frequency bands are spaced by octave which provides a flat readout for a pink-noise power spectrum, which is not unlike the human ear.

As with all IEC standards the specifications are very precise, yet within IEC61260 a number of variants are available to trade off implementation details. Three classes of quality are defined which differ in the filter-band attenuation (band overlap). Class 0 being the best, class 2 the worst acceptable. Furthermore two variants are offered regarding filter-frequency bands, base ten: $10^{\frac{x}{10}}$ and base two: $2^{\frac{x}{3}}$. The centre frequency in either case is 1KHz, with (at least) 13 bands above and 16 bands below.

In the digital domain various alternative implementations are possible, most notably FFT and signal convolution approaches³.

FFT (Fast Fourier Transform, an implementation of the discrete Fourier transform) transforms an audio signal from

³ There are analogue designs to perform DFT techniques, but for all practical purposes they are inadequate and not comparable to digital signal processing.




Figure 6. 30 Band 1/3 octave spectrum analyser

the time into the frequency domain. In the basic common form frequency bands are equally spaced and operation mode produces a flat response for white noise.

For musical applications a variant called ‘perceptual analysers’ is widespread. The signal level or power is weighted depending on various factors. Perceptual analysers often feature averaging functions or make use of screen-persistence to improve readability. They also come with additional features such as numeric readout for average noise level and peak detection to mitigate effects introduced by variation in the actual display.

2.9 Oscilloscopes

The oscilloscope is the “jack of all trades” of electronic instrumentation tools. It produces a two-dimensional plot of one or more signals as a function of time.

It differs from a casual wave-form display, which is often found in audio-applications, in various subtle but important details: An oscilloscope allows reliable signal measurement and numeric readout. Digital wave-form displays on the other hand are operating on audio-samples - as opposed to a continuous audio-signal. Figure 7 illustrates this.

For an oscilloscope to be useful for engineering work it must be calibrated - for both time and level, be able to produce an accurate readout of at least two channels and facilitate signal acquisition of particular events (triggering, signal history) [12].

3. STANDARDISATION

The key point of measuring things is to be able to meaningfully compare readings from one meter to another or to a mathematically calculated value. A useful analogy here is inches and centimetres, there is a rigorous specification of what distance means. There are various standards and conventions, but there is no margin for error: One can rely on the centimetre.

Unfortunately the same rigour is not always applied to audio metering. On many products the included level meter mainly serves to enhance aesthetics, “make it look cool”,

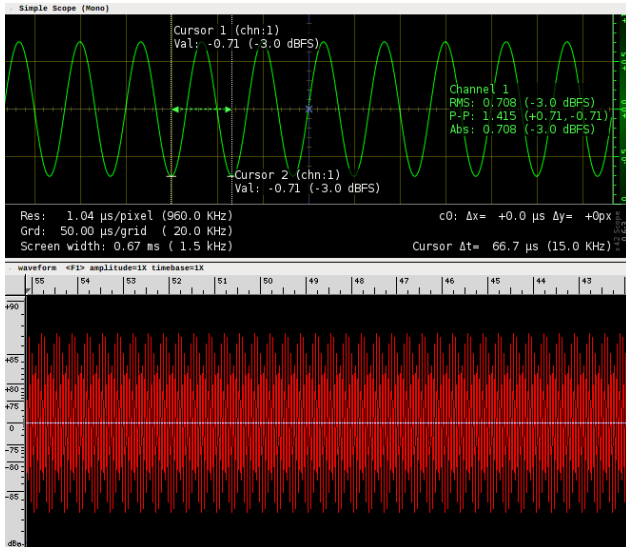


Figure 7. 15KHz, -3dBFS sine wave sampled at 48KSPS. The Oscilloscope (top) up-samples the data to reproduce the signal. The wave-form display (bottom) displays raw sample data.

rather than provide a reliable measurement. This trend increased with the proliferation of digital audio plugins. Those meters are not completely without merit, they can be useful to distinguish the presence, or otherwise, of a signal, and most will place the signal-level in the right *ballpark*. There is nothing wrong with saying “the building is tall” but to say “the building is 324.1m high” is more meaningful. The problem in the audio-world is that many vendors add false numeric labels to the scale to convey the look of professionalism, which can be quite misleading.

In the audio sphere the most prominent standards are the IEC and ITU specifications: These specs are designed such that all meters which are compliant, even when using completely different implementations, will produce identical results.

The fundamental attributes that are specified for all meter types are:

- Alignment or Reference Level and Range
- Ballistics (rise/fall times, peak-hold, burst response)
- Frequency Response (filtering)

Standards (such as IEC, ITU, EBU,...) govern many details beyond that, from visual colour indication to operating temperatures, analogue characteristics, electrical safety guidelines, test-methods, down to electrostatic and magnetic interference robustness requirements.

4. SOFTWARE IMPLEMENTATION

4.1 Meters.lv2

Meters.lv2 [13] is a set of audio plugins, licensed in terms of the GPLv2 [14], to provide professional audio-signal measurements according to various standards. It currently

features needle style meters (mono and stereo variants) of the following

- IEC 60268-10 Type I / DIN
- IEC 60268-10 Type I / Nordic
- IEC 60268-10 Type IIa / BBC
- IEC 60268-10 Type IIb / EBU
- IEC 60268-17 / VU

An overview is given in Fig. 2. Furthermore it includes meter-types with various appropriate visualisations for:

- 30 Band 1/3 octave spectrum analyser according to IEC 61260 (see Fig. 6)
- Digital True-Peak Meter (4x Oversampling), Type II rise-time, 13.3dB/s falloff.
- EBU R128 Meter with Histogram and History (Fig. 3)
- K/RMS meter, K-20, K-14 and K-12 variants
- Stereo Phase Correlation Meter (Needle Display, bottom right in Fig. 2)
- Goniometer (Stereo Phase Scope) (Fig. 8)
- Phase/Frequency Wheel (Fig. 4)

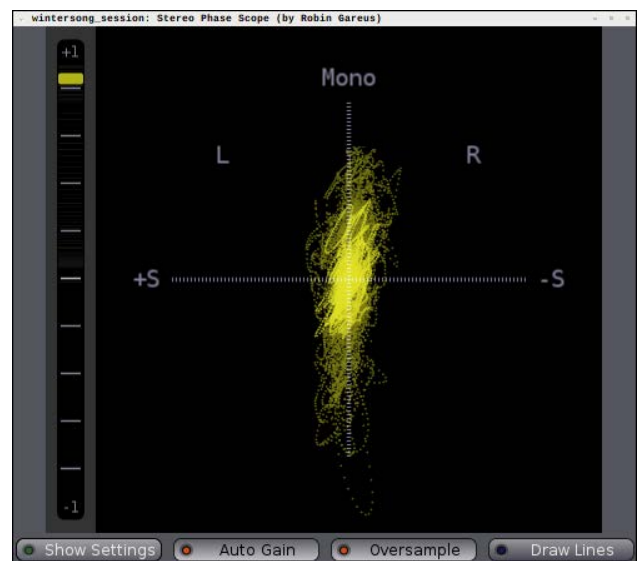


Figure 8. Goniometer (Phase Scope)

There is no official standard for the Goniometer and Phase-Wheel, the display has been eye-matched by experienced sound engineers to follow similar corresponding hardware equivalents.

Particular care has been taken to make the given software implementation safe for professional use. Specifically real-time safety and robustness (e.g. protection against denormals or subnormal input). The graphical display makes use of hardware acceleration (OpenGL) to minimise CPU usage.

4.2 Sisco.lv2

Sisco.LV2 [15] implements a classic audio oscilloscope with variable time scale, triggering, cursors and numeric readout in LV2 plugin format. While it is feature complete for an audio-scope, it is rather *simplistic* compared to contemporary hardware oscilloscopes or similar endeavours by other authors [12].

The minimum grid resolution is 50 micro-seconds - or a 32 times oversampled signal. The maximum buffer-time is 15 seconds. Currently variants up to four channels are available.

The time-scale setting is the only parameter that directly affects data acquisition. All other parameters act on the display of the data only. The vertical axis displays floating-point audio-sample values with the unit [-1..+1]. The amplitude can be scaled by a factor of [-10..+10] (20dB), negative values will invert the polarity of the signal. The numeric readout is not affected by amplitude scaling. Channels can be offset horizontally and vertically. The offset applies to the display only and does not span multiple buffers (the data does not extend beyond the original display). This allows the display to be adjusted in 'paused' mode after sampling a signal.

#	Title	Description
1	Signal Edge	Signal passes 'Level 1'
2	Enter Window	Signal enters a given range (Level 1, 2).
3	Leave Window	Signal leaves a given range (Level 1, 2).
4	Hysteresis	Signal crosses both min and max (Level 1,2) in the same direction without interruption.
5	Constrained	Signal remains within a give range for at least 'Time 1'.
6	Drop-out	Signal does not pass through a given range for at least 'Time 1'.
7	Pulse Width	Last edge-trigger occurred between min and max (Time 1,2) ago.
8	Pulse Train	No edge-trigger for a give time (max, Time 2), or more than one trigger since a give time (min, Time 1).
9	Runt	Fire if signal crosses 1st but not 2nd threshold.
10	LTC	Trigger on Linear Time Code sync word.
11	RMS	Calculate RMS, Integrate over 'Time 1' samples.
12	LPF	Low Pass Filter, 1.0/ 'Time 1' Hz

Table 2. Description of trigger modes in Fig. 9.




Figure 9. Overview of trigger preprocessor modes available in "mixtri.lv2". The arrow indicates trigger position.

The oscilloscope allows for visually hiding channels as

well as freezing the current display buffer of each channel individually. Regardless of display, data-acquisition for every channel continues and the channel can be used for triggering.

The scope has three modes of operation:

- **No Triggering** The Scope runs free, with the display update-frequency depending on audio-buffer-size and selected time-scale. For update-frequencies less than 10Hz a vertical bar of the current acquisition position is displayed. This bar separates recent data (to the left) and previously acquired data (to the right).
- **Single Sweep** Manually trigger acquisition using the push-button, honouring trigger settings. Acquires exactly one complete display buffer.
- **Continuous Triggering** Continuously triggered data acquisition with a fixed hold time between runs.

Advanced trigger modes are not directly included with the scope, but implemented as a standalone “trigger pre-processor” [16] plugin, see Fig. 9 and Table 2. Trigger-modes 1-5 concern analogue operation modes, modes 6-9 are concerned with measuring digital signals⁴. Modes 10-12 are pre-processor modes rather than trigger modes. Apart from trigger and edge-mode selectors “mixtri.lv2” provides two level and two time control inputs for configuration.

5. CONCLUSION

An overview of various instrumentation tools and measurement standards was presented. The various tools are available as free software and have already found their way into GNU/Linux distributions, making Linux even more suitable as a platform for Pro-Audio work.

Acknowledgments

Kudos to Fons Adriaensen, DSP expert, author of jmeters, jkmeter, ebur128 and related zita-* software that was used as the basis to facilitate the software implementation of meters.lv2.

Thanks go to Jaromír Mikeš who created the initial software packages for Debian/Ubuntu, Alexandre Prokoudine who took great care of public relations, and Axel Müller who spent countless hours on testing and quality assurance.

Further thanks go to harrison.com which caters for binary packages for OSX and Windows platforms, to be made available with their Mixbus3 Product.

6. REFERENCES

- [1] E. Brixen, *Audio Metering: Measurements, Standards and Practice*, 2010, ISBN 0240814673.
- [2] J. Watkinson, *Art of Digital Audio*, 2000, ISBN 0240515870.
- [3] B. Katz, “How to make better recordings in the 21st century - an integrated approach to metering, monitoring, and leveling practices.” <http://www.digido.com/how-to-make-better-recordings-part-2.html>, 2000, updated from the article published in the September 2000 issue of the AES Journal by Bob Katz.
- [4] Wikipedia, “Peak programme meter — wikipedia, the free encyclopedia,” 2013, [Online; accessed 1-February-2014]. [Online]. Available: http://en.wikipedia.org/w/index.php?title=Peak_programme_meter&oldid=547296518
- [5] IEC, “IEC60268-10 sound system equipment – Part 10 Peak programme level meters,” 1991, International Electrotechnical Commission.
- [6] ITU, “ITU BS.1770, Algorithms to measure audio programme loudness and true-peak audio level,” <http://www.itu.int/rec/R-REC-BS.1770/en>, 2006, International Telecommunication Union.
- [7] EBU, “EBU R-128 – audio loudness normalisation & permitted maximum level,” <https://tech.ebu.ch/docs/r/r128.pdf>, 2010, European Broadcast Union – Technical Committee.
- [8] F. Adriaensen, “Loudness measurement according to EBU R-128,” in *Proceedings of the Linux Audio Conference 2011*, 2011.
- [9] EBU, “EBU TECH 3342 – Loudness Range: A measure to supplement loudness normalisation in accordance with EBU R-128,” <https://tech.ebu.ch/docs/tech/tech3342.pdf>, 2011, European Broadcast Union – Technical Committee.
- [10] Wikipedia, “VU meter — wikipedia, the free encyclopedia,” 2014, [Online; accessed 2-February-2014]. [Online]. Available: http://en.wikipedia.org/w/index.php?title=VU_meter&oldid=588986200
- [11] IEC, “IEC61260 electroacoustics – Octave-band and fractional-octave-band filters,” 1995, International Electrotechnical Commission.
- [12] F. Adriaensen, “Design of an audio oscilloscope application,” in *Proceedings of the Linux Audio Conference 2013*, ISBN 9783-902949-00-4, 2013.
- [13] R. Gareus, “Meters.lv2 - audio plugin collection,” <https://github.com/x42/meters.lv2>, 2013.
- [14] “Gnu general public license, version 2,” <https://www.gnu.org/licenses/old-licenses/gpl-2.0.html>, Free Software Foundation, June 1991.
- [15] R. Gareus, “Sisco.lv2 - Simple Scope,” <https://github.com/x42/sisco.lv2>, 2013.
- [16] —, “Mixtri(x).lv2 - matrix mixer and trigger processor,” <https://github.com/x42/mixtri.lv2>, 2013.

⁴ Digital signal trigger modes are of limited use with a generic audio interface, but can be useful in combination with an adapter (e.g. Midi to Audio) or inside Jack (e.g. AMS or ingen control signals).