LOUDNESS NORMALISATION IN DISTRIBUTION THE SAFETY NET

Richard van Everdingen

Broadcast and Audio Consultant at Delta Sigma Consultancy
richard@delta-sigma-consultancy.nl

ABSTRACT

EBU R128 is a milestone in the history of broadcasting by specifying normalised loudness levels in production, in play-out systems and in transmission. However, not all broadcasters may follow the book, an act which threatens the position of the ones who do. This paper explains how distributors could support the good broadcasters while improving consumer satisfaction at the same time. But even if all broadcast stations were to transmit with the right loudness levels, there is no guarantee that the equipment at home would reproduce the signals correctly. Issues regarding loudness occur in consumer equipment. The resulting level jumps cause annoyance and spoil the result of R128 in production and play-out. This paper includes the solution how to deal with these difficulties.

1. LOUDNESS NORMALISATION

After three years of work, in 2010 the European Broadcasting Union (EBU) published Recommendation 128 [1]. It describes the revolutionary change from peak to loudness normalization, based on the algorithm published in ITU-R BS.1770 [2]. A number expresses the level in LUFS[1] (Loudness Unit referenced to digital Full Scale) on an absolute scale or in LU (Loudness Unit) on a relative scale.

The target loudness level according to EBU R128, which all broadcast stations should comply with, is specified as –23 LUFS. The measured loudness tells us how loud the programme is. That explains the term Programme Loudness. A programme is an individual event from start to finish and can be anything from a short commercial to a three hour film. By aligning (normalizing) Programme Loudness, the average loudness of all broadcasting events can be equal, which is the basis for a consistent audio level delivery to viewers and listeners.

1.1 Distribution guidelines

Besides the stages production and play-out, the EBU group PLOUD included the distribution and reproduction stage in the scope of its work. As a result, a document has been created which engineers and designers could never have found before: the EBU Tech 3344 guidelines for broadcast distribution and consumer equipment [3]. That report includes all processing stages and works around all known problems in the link between the broadcast center and the consumer. This paper offers an introduction to what can be found in detail in that almost 90-page document.

1.2 The second line of defense

At first sight, one would say that broadcasters 'just' need to switch to EBU R128 and that is it. Distribution parties can subsequently transfer the signals transparently and loudness consistency is guaranteed. However, it is not that simple. In many countries there is no legislation about loudness levels, which means that broadcast stations that do not follow R128 can keep transmitting at whatever level they want. However, even regulation is not ideal, as somebody needs to keep an eye on what is going on out there, acting like a loudness police – strictly and firmly if parties do not obey the rules. Forcing equal loudness by law can also lead to heavy audio processing, making the sound as dull as ditch water and the dynamic range as flat as a pancake as the objective focuses mainly on avoiding charges by the authorities.

Unfortunately, the loudness war has become part of the competition amongst stations. Viewing figures and ratings, along with the ever-repeated message to 'stay tuned' dominate in broadcast land. Normalizing the loudness to the EBU recommended target level typically means that stations go down in average

[1] The ATSC A/85 standard uses the term LKFS instead of LUFS. Other organizations, such as the EBU, use the term LUFS to comply with ISO 80000-8. Both units refer to the same measurement.

loudness in comparison to their previous level. Although they can gain a lot in audio quality, their decision to lower the loudness level actually puts them in a weak position against competitors that remain at their old – sometimes extremely loud – level. Instead of seeking refuge in legislation, PLOUD came up with a much more efficient approach in offering protection and comfort: Loudness Normalization in Distribution.

2. ALIGNMENT IN THE DISTRIBUTION STAGE

By making the distribution stage part of the R128 paradigm, several goals can be achieved at the same time. The challenge is to pass through signals untouched from those broadcasters that transmit by the R128 book, while correcting those that deviate, thereby intimidating the good broadcasters. Just as loudness normalization of all content cannot be achieved without the negative side-effects of using an end-stage processor in the studio of the broadcast station, similar fast-operating processing devices in distribution head-ends would cause the same drawbacks and would even spoil the quality of stations that transmit in a proper way. That is why a completely different approach has been designed.

2.1 Long-term alignment

The only acceptable solution – not least for the EBU Members themselves – is to align loudness levels integrated over a long term. This starts by performing measurements. Rather than assessing the loudness of individual programmes that would require individual start, stop and pause data being available, the trick is to monitor a service over a full day by looking at the loudest measurements during 24 blocks of one hour each. The outcome of that analysis is called 'Service Loudness' and represents the averaged maximum loudness of the broadcast station. The process used, in which the loudest blocks are selected, is meant to discard the blocks that are not representative, for example when playing just background music during the night while the video is displaying text pages.

The target level is –23 LUFS, with a maximum deviation of +1 LU. As R128 itself specifies a strict and narrow maximum range of ±1 LU for the loudness of individual live programmes, it is quite representative to take the loudest one-hour blocks from a full day as a reference. To improve stability and to avoid negative side effects of this estimation, the outcome of a full day's measurement is integrated yet again, over several days. After all, it is not the performance of a particular day that is of interest, but the long-term trend that the station applies. The outcome of the measurement feeds an adaptation device that applies the loudness correction. This is done by means of an offset value that remains static until it is updated, if necessary, at three o'clock in the morning. The reason for using this particular time is to have minimal influence on daily programming.

The gain change is done very efficiently in the digital domain, avoiding any quality loss, and it is described in Tech 3344 for MPEG-1 Layer II, Dolby Digital (Plus) and HE-AAC codecs. Dynamically switching metadata remains fully supported. The task can be implemented easily in a DVB multiplexer or a similar processing device. By separating the measurement and adaptation, the latter is reduced to a simple process without introducing risks for reliability. The measurement system could, for example, be switched off for maintenance without affecting the continuity of the distribution.

Broadcasters remain responsible for the consistency of the loudness of their own output over the day, as it should be. Inconsistency on an hourly basis – in other words, loudness levels in some blocks that are significantly louder than those in other blocks on a daily basis – will result in a less loud average transmission level overall, not the other way around. This encourages broadcasters to get their loudness level more consistent, without giving them the chance to deliberately peak loud during the competitive prime time. However, broadcasters which conform to the loudness requirements but which support a relatively wide loudness range will not be penalized and will be aligned to other stations that have reasons to compress heavily.

Dynamic Range Control (DRC) implemented in Dolby Digital (Plus) and HE-AAC can influence the perceived loudness. It is expected that the next update of Tech 3344 will include information on how to handle situations where the loudness differs substantially if several DRC settings are used. This includes movie channels that use loudness normalization based on voice level.

Loudness normalization in distribution is not only an answer to differences in television services, but can be used for digital radio transmission and for FM radio as well. Being applied to all services, it creates 'the second line of defense' in the cooperative fight against the loudness war.

2.2 Other sources

Nowadays, the separation between creation and delivery is not as clear as it used to be. Many distributors are generating and broadcasting their own content by means of video-on-demand, locally-inserted advertisements, show channels and interactive set-top box applications. It is the aim to have all these sources at the same target level. For all file-based material, the approach can be the same as for a broadcast station. An application such as ad-insertion is by tradition susceptible to the creation of loudness jumps. However, as long as the average loudness of the station itself is consistent over time, switching over from that service to a local source becomes a simple task if loudness normalization in distribution is being applied. As the content to be inserted is also file-based, it can be normalized in advance to –23 LUFS so that it can be

inserted into any desired channel, including those that support Dolby Digital or HE-AAC. By combining service and file-based normalization, all content leaves the head-end with equal long-term loudness, which forms the ideal take-off point for further processing at home.

3. CONSUMER EQUIPMENT

Perhaps one would suppose that loudness issues are solved once the levels of all services have been equalized in the head-end, as we just discussed. Regrettably, this is not the case. Another link in the chain is involved. To reproduce the television and radio signals, consumer equipment is used, such as a set-top box, a television set and home theatre equipment. These devices have a number of connections and are able to process several kinds of signals. That is exactly where the problems arise.

Whenever signals, codecs and interfaces come together in an audio device, there is a risk that differences appear regarding loudness levels. If, for example, the set-top box is used as the source, loudness jumps of 11 LU or more can currently spoil playback by a connected home theatre receiver when switching from one television service to another. This sounds about twice as loud to our ears and it usually forces people to grab the remote control urgently. The outcome can differ between set-top box and iDTV[2] brands and models. Most set-top boxes also have a volume control that often affects signal levels on one codec, while there is no effect at all on others, if the sound is reproduced by a home theatre device or certain iDTVs. This behavior shows itself as yet another level uncertainty when switching between TV channels.

The combination of faults makes it very difficult to achieve consistency and it may look as if it is just a fact of life that loudness levels simply never appear correct. However, that was not acceptable for the PLOUD group. They therefore included these concerns into the scope of the project in an ultimate attempt to resolve them. Not just some, but all of them, and also in the shortest possible term. To counteract the consumer-equipment issues, Tech 3344 contains extensive guidelines while integrating maximum backward compatibility with the installed base of devices.

3.1 Cause of the trouble

There are roughly two reasons for the level chaos. First, there is often a lack of clarity about the alignment of codecs, systems and interfaces. This includes analogue as well as digital systems. Although some of these alignments have been written down in technical standards and recommendations, something essential was missing: a common loudness level. The change to –23 LUFS as a reference level offered the PLOUD group the opportunity to create a completely new overview from the physical output of the broadcast studio up to the physical output connections on a consumer playback device. A number of these overviews are included in Tech 3344 in the form of tables and graphs for each European television and radio transmission system. It includes set-top boxes, television sets, radios (FM and DAB), media players and professional equipment. To improve understanding, the correct level relationship can just be read out graphically, giving the engineers the answers at a glance as to 'what comes out if a certain level goes in' (see Figure E-1 in Appendix E for details).

The second cause is a lack of interoperability that stretches out beyond the territory of one codec, one system or one interface. Connections like HDMI were added to set-top boxes at some point, but apparently without awareness that the designs of the different systems that can connect to this interface – a television set or a home theatre device – require other levels because they are intended to treat audio in a different way. Television sets need to be designed to operate with loudness levels that are compatible with the built-in analogue tuner and analogue interfaces like the SCART connector. Home theatre devices on the other hand, are built to function in an environment of decoders that are designed to operate with the highest dynamic audio range ever to be found on digital media. Hence, their internal loudness reference is considerably lower than that of a television set, something that needs to be taken into account when making connections. To make matters even worse, AV receivers are manufactured in one of two different ways, depending on model and brand, which causes a mismatch of 4 dB that adds to or subtracts from the level variation of set-top boxes. This 4 dB originates from THX guidelines and causes a loudness level difference between MPEG and Dolby Digital encoded television services that have been aligned to –31 LUFS.

Certified home theatre systems can be recognized by the attached THX logo. However, some manufacturers also decided to design their non-certified AV receivers like that, which makes it hard to see from the outside to which of the two families a particular unit belongs. However, many of these receivers expose their true nature by displaying 'Dialnorm –4 dB' or 'D-norm –4 dB' if a source that supplies a –23 LUFS referenced Dolby stream is connected, where the number indicates the opposite offset to a level of –27 LUFS. It is again a hidden piece of the puzzle of the reproduced loudness level.

Television sets themselves also seem to appear in two design families, one accepting Dolby encoded streams on their HDMI input, while others explicitly do not. Where loudness is concerned, the variations that can be found in the design of consumer equipment make the expected behavior of a set of devices in the home appear like a lottery. Only thorough awareness about what is happening makes it possible to understand why the levels appear different depending on codecs, systems, settings, connections, dynamic range control and volume control behavior.

[2] iDTV = A television set with integrated digital decoder.

3.2 One fix for all

As the current HDMI specification seems to miss a mechanism to faultlessly distinguish all the use cases that lead to different loudness levels depending on the connected device, the approach to resolve that matter is reduced to just one basic principle. Tech 3344 works around the problems in consumer equipment in a modest way: just make the user able to tell the set-top box or the iDTV what device is connected. Subsequently, the device applies the correct levels. This is the quickest and simplest way to fix the issues for television broadcasting and includes a golden rule for implementing a change: backward compatibility. After all, we cannot just tell the consumers to replace all their television sets and home theatre equipment now that EBU R128 is being applied.

The approach for the SPDIF output is exactly the same. By implementing the adaptation in the set-top box, all currently installed television sets and home theatre equipment can be used in a consistent and best possible way. Set-top boxes and iDTVs are the most 'reachable' devices from outside and if the design allows implementation of adaptations by use of a firmware update, the opportunity is offered to solve or at least reduce the problems remotely. Set-top boxes also have the relatively shortest 'replacement time' which makes it possible to get rid of faulty designs as soon as possible.

3.3 What's on the menu?

Instead of sometimes rather cryptic settings to be found in the user menu of present set-top boxes and iDTVs, the way to describe which device is connected can be done in several user-friendly and intuitive ways. Although the manufacturer is fully free to design his own menu lay-out and content, some hints can be given. Figure 1 shows a possible way how this could look. The menu could also offer a 'wizard' or graphics to guide the user through the settings during installation, but as has been said, this is up to the manufacturer to choose. Figure 2 shows a simple example of a graphical layout.

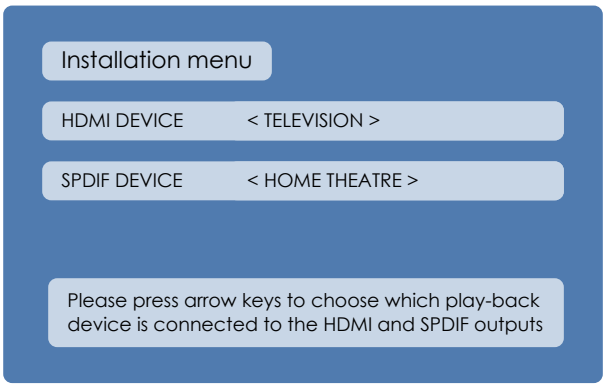

Figure 1: Logical and simple user settings in the installation menu that can be used to enable the set-top box to apply the right loudness levels.

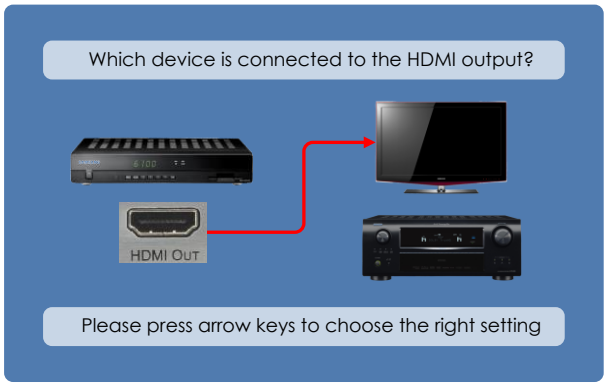

Figure 2: The same user setting as in the picture before, but now graphically displayed.

If the user connects a home theatre device to the HDMI output, we run into the complication that this device belongs to one of the two families that produce an undesired loudness difference of 4 dB when switching between codecs and services. Although settings like 'LOW and HIGH', 'LESS and MORE' or 'STANDARD and THX-STYLE' could help, a more intuitive way to assist the user in making the correct menu choice can be created if the set-top box or iDTV is able to play out a signal with an audible pink noise, sequentially and repeatedly coded as PCM and Dolby Digital. If the user does not hear a difference in loudness when the menu item appears, the setting is correct for the connected device. If the set-top box or iDTV contains an internal media player, all the required components are already available to perform this job. Another way to identify a THX-style receiver is to periodically interrupt a Dolby test stream that carries –23 LUFS referenced metadata. When the set-top box or iDTV asks the user to make the right choice, most THX-style AV receivers (certified or not) will show 'Dialnorm –4 dB' or 'D-norm –4 dB' on its display according to the THX requirements.

The scenarios above are backwards compatible with currently installed televisions sets and home theatre equipment. As pointed out, there are also two design families of television sets, but the possible complications arising from that state of affairs are taken care of by a specific setting included in the Tech 3344 approach: If the user selects the default setting for a connected TV set by choosing 'HDMI DEVICE = TELEVISION' in the text-based menu variant, the output of the set-top box is restricted to PCM audio only. This option guarantees that the loudness level of the television broadcasters appears correct on ALL television sets with HDMI input, regardless of brand and model, new or old.

An alternative scenario is created if a home theatre device is connected to the television set, as shown in Figure 3. This scenario is also included as an option in the menu settings described in Tech 3344, but in this specific case the television set must also comply with that document to properly process the levels. For older television sets it is recommended to choose the default setting 'HDMI DEVICE = TELEVISION'.

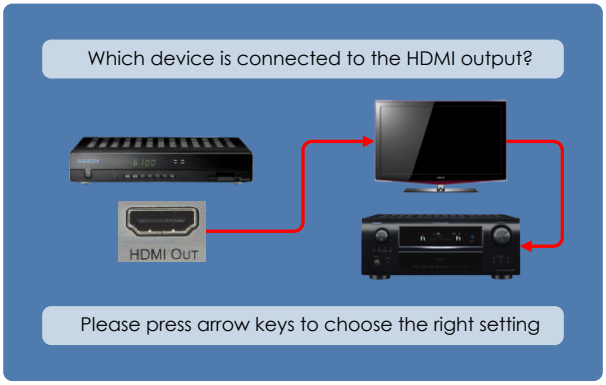

Figure 3: Correct loudness processing on a television set which is compliant to Tech 3344 that enables pass-through via the television set using the HDMI Audio Return Channel (ARC).

3.4 Future solution

Although Tech 3344 offers solutions to work around the loudness-level issues in consumer equipment, it would be very attractive if this matter could be handled automatically without bothering the user. This is an open opportunity and could be arranged as part of an update of the HDMI specification. Next-generation set-top boxes, iDTVs and home theatre equipment would then be able to recognize each other and apply the correct loudness levels once the connection is made, without intervention and independently of brand and model. Following the golden rule, this should be done while being backwards compatible: a next-generation model should still offer the Tech 3344 solutions to be able to connect equipment which is based on older designs. Eventually, the issues will die out completely.

3. ANALOGUE SYSTEMS

In spite of the continuing trend to transmit broadcast signals digitally, analogue TV systems and FM radio will still be present in many countries for several years to come. In cable head-ends, the digital signals need to be converted by means of digital decoders and analogue modulators. A few matters complicate this step. If loudness levels differ on digital platforms, the line-up for analogue modulation also varies per channel, despite the attempt of the cable operator to decrease these variations manually. This also makes it very difficult to standardize the adjustments. The lack of calibration results in the situation where the loudness of, for example, BBC1 can be higher than ZDF in one place, but reversed in another cable network. It is also hard to avoid changes in loudness level after the equipment has been exchanged due to a defect. If you perceive a sudden change in average level from one day to another, it may well be that your cable operator has replaced some equipment.

In the light of R128, it would help if engineers knew how to adjust their systems consistently. EBU Tech 3344 therefore includes level overviews and tables for all European television and radio systems. Yet, this way to standardization is still dependent on the number of broadcast stations that comply with R128 without too much deviation. At this point, loudness normalization at the distribution's head-end offers the solution. As all content is normalized to the EBU target level, there is also only one lineup setting left for decoders and analogue modulators. This means that if normalization is being applied to the digital network and the systems for analogue modulation are fed by that same platform, all radio and TV equipment – including analogue ad-insertion – can be left on default settings. Time-consuming maintenance on determining the adjustments can be avoided. Attention, Directors and Managers of Cable TV Networks! Normalizing the loudness at the head-end is an investment that earns itself back.

3.1 What about radio?

Radio on FM and DAB (+) is fully included in Tech 3344. Cable networks can use the same normalization process as for television and benefit from the same advantages: equal loudness and the use of default settings on equipment. The alignment scheme can be used to achieve standardization for combined DAB (+) and FM (car) receivers. For FM radio, Tech 3344 introduces an unambiguous loudness reference, independent from stereo or mono modulation and irrespective of the amount of bandwidth used for additional signals in the FM multiplex. It offers a practical alternative for terrestrial planning standard ITU-R BS.412 for optimum use of the terrestrial FM band. MPX power limiters can in that case become obsolete in order to achieve a stable and consistent loudness level.

3.2 Other consumer equipment and more

Instead of just working around the complications by improving set-top boxes and iDTVs, Tech 3344 also addresses the weaknesses in other consumer equipment directly. Manufacturers can find recommendations for home theatre equipment, television sets, media players including DVD and Blu-ray, DAB and FM radio receivers.

4. CONCLUSIONS

A lot of complications, uncertainties and problems can be found along the route from the broadcast studios to the viewers and listeners at home. Annoyances – because of level jumps and quality losses – are the result. By extending the scope of the EBU loudness normalization project up to and including distribution and consumer equipment, these issues have been investigated and identified. Subsequently, solutions for these troubles have been designed. Together with ready-to-use instructions, overviews, hints and suggestions for improvements, this has led to a document that engineers and designers could never have found before: the EBU Tech 3344 guidelines for broadcast distribution and consumer equipment.

Although some of the issues are related to shortcomings in the design of billions of consumer

devices already in use at home, Tech 3344 offers the possibility of working around them by adapting the most accessible devices: the set-top box and the iDTV. Engineers and technical management of distribution companies, broadcasters, legislators and manufacturers of head-end equipment and consumer devices are strongly recommended and very much invited to dive into this work and to follow up the recommendations and instructions at the first opportunity. The industry is also advised to extend the HDMI specification with reliable identification that can be used to automatically solve the loudness issues at a later stage. After all, that is what we demand in this decennium: connect and play, without having to deal with loudness-level issues.

The technical guidelines or specifications contained in this document refer to broadcast recommendations and standards developed by standard-settings organisations, in particular:

[1] EBU R128 (2011) Loudness normalisation and permitted maximum level of audio signals.
[2] ITU-R BS.1770-3 (2012) Algorithms to measure audio programme loudness and true-peak audio level.
[3] EBU Tech 3344 (2011) Practical guidelines for distribution systems in accordance with EBU R128

APPENDIX A: Loudness levels from the distribution head-end to the home

A hypothetical example of the loudness levels of seventy different TV channels at the output of the distribution head-end after applying normalization is shown in Figure A-1; all broadcast signals have around the same long-term loudness level. For the comparison, both a program reference level (loudness indicating metadata) and a reproduction target level of –23 LUFS is assumed in case of Dolby Digital (Plus) and HE-AAC audio streams.

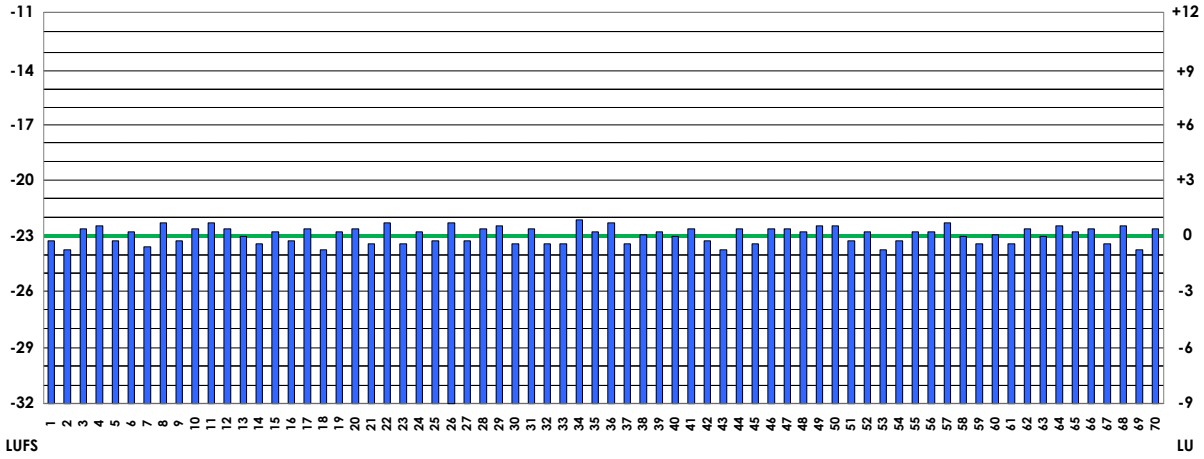

Figure A-1: Loudness levels at the output of a head-end. The services are listed on the x-axis.
The left y-axis shows the absolute loudness level (in LUFS), and the right y-axis displays the loudness on the relative scale (in LU).

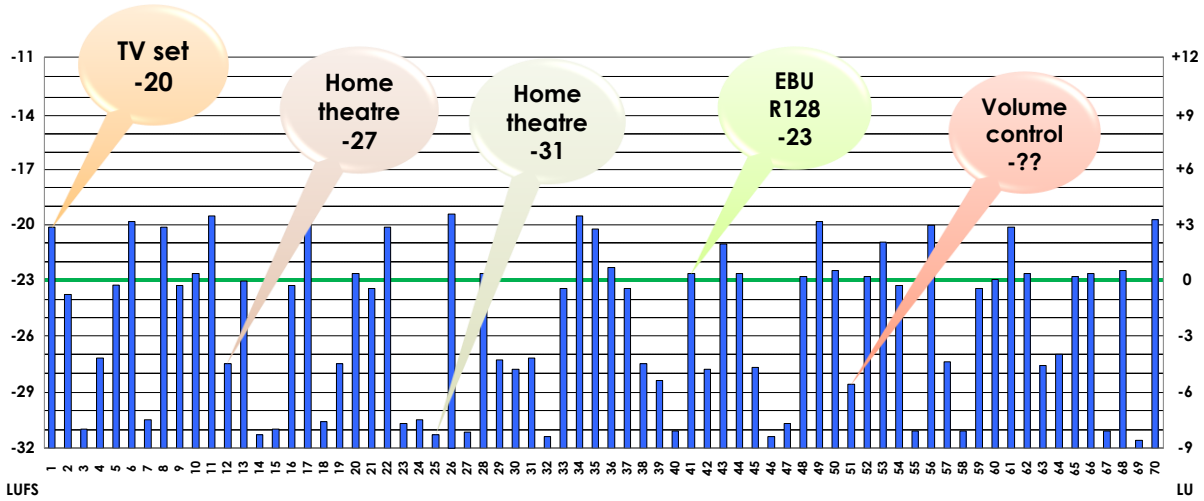

Figure A-2: The reproduced loudness level differs, depending on the equipment and the position of the volume control on the set-top box.

Based on the current situation, the reproduced loudness at home can be very different, in spite of the fact that the head-end levels are indeed equal. The reproduced loudness level depends on both the consumer equipment and the listening situation. For example, in Figure A-2:

- -20 LUFS => A set-top box or IDTV that uses its internal Dolby Digital (Plus) RF Mode decoder.
- -23 LUFS => A set-top box or IDTV that decodes a TV channel with MPEG-1 Layer II audio.
- -27 LUFS => A set-top box or IDTV delivering a Dolby Digital (Plus) bitstream to a THX style home theatre system.
- -31 LUFS => A set-top box or IDTV delivering a Dolby Digital (Plus) bitstream to a regular home theatre system.
- ?? LUFS => A set-top box or IDTV that decodes a TV channel with MPEG-1 Layer II audio or that uses its internal Dolby Digital RF Mode decoder, while the volume control is affecting the output level.

The problem is that a switch from one TV channel to another can cause a level jump between any two of the levels shown in this list, i.e. up to 11 dB when the user is unlucky. If Tech 3344 is being applied in the set-top box or iDTV, the

levels are adjusted for all codec systems to the reference level of all currently-installed TV sets and home theatre systems.

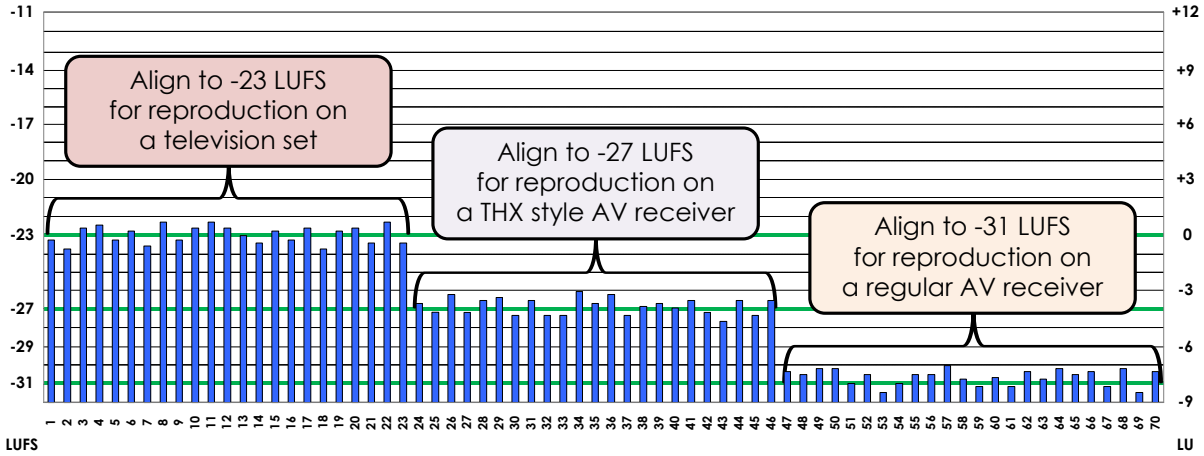

Figure A-3: The set-top box or iDTV aligns the loudness level correctly for the particular playback device to which the consumer is listening. Consequently, all codec systems of all services appear at the same loudness level within that box. By the way: the loudness levels in the graph are relative and indicate the jump caused by incorrect processing. Ultimately, the loudness playback level is defined by the consumer, who makes use of the volume control of the equipment chosen to reproduce the sound.

APPENDIX B: Different use cases for consumer equipment

This section shows an overview of the levels of a set-top box, a television set and a home theatre device. The level attenuation stated refers to PCM signals only. Native bitstreams like Dolby Digital (Plus) or HE-AAC are transferred over the interface without any change. Figure B-1 shows the level overview where the HDMI connector is used to feed the TV, while an AV receiver is connected to the SPDIF output. It means that the set-top box sends PCM audio to the TV set at a level of –23 LUFS, while it attenuates the signal to –27 or –31 LUFS on the SPDIF if a home theatre device is connected, depending on the type of AV receiver. The level adaptation inside a settop box is not meant to act like an automatic gain control; it only sets the level for the relevant device by using the proper and fixed attenuation. Figure B-2 displays the situation where the AV receiver is fed from the set-top box directly via its HDMI input while the TV set is connected to the HDMI output of that home theatre device. It means that the set-top box attenuates PCM audio to –27 or –31 LUFS depending on the type of AV receiver. In case the viewer wishes to listen to the loudspeakers of the TV itself, the AV receiver forwards the audio at a level of –23 LUFS. Figure B-3 shows the variant where the TV set is connected to the set-top box while the audio is forwarded to the AV receiver. The set-top box sends PCM signals to the TV at a level of –23 LUFS. The TV set attenuates PCM levels to –27 or –31 LUFS, depending on the type of AV receiver.

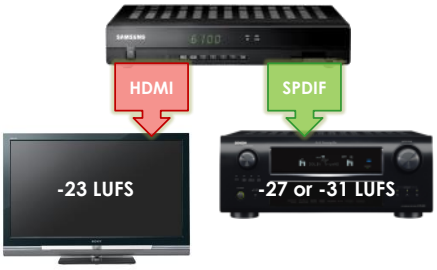

Figure B-1: TV connected to the HDMI, and the AVR connected to the S/PDIF, of a set-top box

Figure B-2: AVR connected to the HDMI of the set-top box, TV connected to the HDMI of the AVR

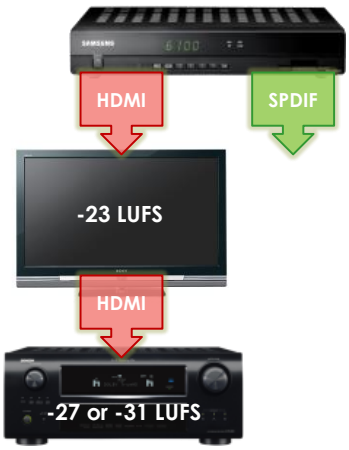

Figure B-3: TV connected to the HDMI of the set-top box while the AVR is connected to the HDMI of the TV, making use of its Audio Return Channel