

L2—UltraMaximizer

Software audio processor

User's Guide



Chapter 1 - About the L2

The L2-Ultramaximizer is the third-generation software processor, combining an advanced peak limiter, level maximizer and a high performance re-quantizer and dithering system called IDR (Increased Digital Resolution). In addition, the L2 software processor includes Waves' ARC Auto-Release Control system, and ninth-order noise-shaping in the IDR implementation. Processing in the L2 is 48-bit (double precision) processing, offering improvements for all resolutions, and the opportunity to dither to a 24-bit output for archives and DVD. The L2 offers superb requantization for all bit-depths, including 24, 22, 20, 18, and 16-bit outputs. The development of L2's lookahead peak limiter provides the mastering engineer with the capability to increase sound file resolution and production master levels with precise control and dithering options.

Waves IDR™ process brings more choices, greater control, and unmatched compatibility to the mastering environment, whether for high-resolution CD or low-resolution multimedia. IDR offers two dither types and three noise-shaping curves for optimal processing in a wide variety of applications and with a wide choice of source material.

While the operation of conventional limiters is well understood, the limiter section of the L2-Ultramaximizer is capable of a very fast, overshoot-free response, and once the limiter threshold has been set, the user can then go on to define the actual peak level that the processed signal will reach. Once set, limiting and level re-scaling becomes a one-shot process. For mastering purposes, the peak level of the processed signal would normally be set to 0dB, or just below 0dB. Because a typical digital audio file of music contains many high intensity, short duration peaks, simple normalization of the file may still result in a low average signal level. Using the L2-Ultramaximizer however, it is generally possible to significantly increase the average signal level of a typical audio file without introducing any audible side effects. In the event that a deliberately limited sound is required however, there is more than adequate range of the limiter parameters to recreate 'vintage' effects such as level pumping or severely limited dynamic range. The UltraMaximizer is specifically designed for mastering, digital editing, multimedia, and any application that requires limiting and/or requantization of the digital signal with highest quality. In order to ensure the maximum possible resolution of a processed signal, it is very highly recommended that the L2 be used last in the processing chain. Failure to observe this will not prevent L2 from working, but you should be aware that both the absolute brickwall limiting AND the benefits of IDR requantization will be compromised and will need re-limiting to maintain the original level.

The L2 is available as a stereo processor only. It will not show under mono inserts and to process mono data you will have to make it stereo or bus it to a stereo Auxiliary input.

Chapter 2 - Some secrets of digital audio

In order to make the best use of the L2-Ultramaximizer, it is important that we explain some of the less obvious implications of digital audio. Once these have been explained, you will understand why Waves thought a product such as L2 was necessary and you will be in a better position to make use of its powerful features. The operation of the L2 breaks down into two main areas:

- 1 - maximum level of the digital signal through proprietary peak control.
- 2 - maximum resolution of the signal through dithering and noise-shaping.

ABOUT MAXIMUM LEVEL

The maximum level of a digital signal is governed by the highest peak in the file. Simple normalization finds the highest peak, then raises the entire signal so that this peak is at the maximum value. However, many of these peaks may be of very short duration and can usually be reduced in level by several dBs with minimal audible side effects. Those familiar with digital editing systems may even have proved this for themselves by 'redrawing' some trouble-some peaks by hand. By transparently controlling these peaks, the entire level of the file can be raised several more dB than by simple normalization resulting in a higher average signal level.

The L2-Ultramaximizer avoids the possibility of overshoot by utilizing a lookahead technique that allows the system to anticipate and reshape signal peaks in a way that produces the bare minimum of audible artifacts. Because there is no possibility of overshoot, L2 can be used with absolute confidence in situations where brickwall limiting is important.

ABOUT MAXIMUM RESOLUTION

Any digital signal processing that alters the original digital data (mixing, gain changes, EQ, dynamic processing, etc.) generally increases the number of bits required to represent the signal. Conventional truncation results in a loss of signal-resolution each time the signal is processed. The human ear uses this low level information to construct a mental image of the stereo soundstage, so any compromise in this area manifests itself as a loss of spaciousness and transparency. Waves' IDR prevents this loss of critical low-level detail. Even when processing 16-bit signals, it is normal to process with at least 24 bits resolution, or more (as in the L2, which is now double-precision). However, as soon as the resolution is pulled back down to 16 bits by rounding or truncation (by removing the bottom 8 bits), the resulting rounding error produces an audible distortion at low signal levels, and a permanent loss of digital resolution that can never be recovered. If the audio signal is repeatedly processed and truncated back to 16 bits, the losses accumulate, causing a significant loss of fidelity, most evident as a loss of the tonal subtleties of low-level sounds within a mix. The human ear uses such low level information to construct a mental image of the stereo soundstage, so any compromise in this area manifests itself as a loss of spaciousness and transparency.

The solution is to properly dither and noise-shape a signal each time the wordlength is increased then reduced (such as nearly every digital signal process will require).

WHY USE THEM, AND WHAT ARE THEY?

Proper dithering is simply this: before the requantization (reduction of the wordlength), a precisely controlled amount of noise (termed 'dither') is added to the signal. This can convert the low-level nonlinear distortion caused by truncation into a simple steady hiss, thereby removing all traces of low-level non-linearity, but at the expense of a very slightly increased background noise. Obviously, increased noise levels are not ideal in high quality audio applications, but fortunately, the perceived level of this dither noise can be greatly reduced by **'shaping' the noise in such a way that it falls into an area of the audio spectrum where the human ear is least sensitive.**

The main point of maximum resolution is simple: to 'capture' the best possible quality into a shorter wordlength (smaller bit-depth) from a longer wordlength (higher resolution).

A BIT ABOUT IDR™

IDR is a Waves proprietary noise shaping dither system developed by the late Michael Gerzon and Waves, and signifies a major advance in preserving and actually increasing the perceived resolution of the digital signal being processed. You can utilize IDR each time during subsequent processes (for 16-bit storage), or once at the end of a high-resolution chain (say, 24-bit), in order to ensure that the final signal has the maximum resolution possible. IDR is of particular benefit when data is deliberately requantized from 48-bit to 24-bit, 24-bit to 20-bit and so on.

Dithered outputs of 24, 22, 20, 18, and 16-bit are available. The L2 features double-precision resolution, that is, for TDM systems, all internal limiting and gains calculated with 48-bit fixed precision. Dithering back to 24-bit output is now possible for the new DVD and other delivery media, as well as for archives of masters.

By using the IDR implemented within the L2-Ultramaximizer, optimum results may be achieved during final file preparation, mastering, and quantization or requantization. The greatest possible implementation of IDR available is in the L2-hardware limiter and this software L2 plug-in from Waves, both featuring 9th-order noise shaping (the original L1 software has 2nd-order noise-shaping).

Chapter 3 - Using the L2



THE PEAK LIMITER SECTION

1 - Open a soundfile in your host application. Select part or all of the soundfile you wish to process with the L2. Pass audio through the L2; for more information, see your host application system manual.

2 - Listen to the output of the L2. Set the threshold of the limiter by dragging the left triangle (Threshold) down. Leave the Out Ceiling setting at the maximum value: 0.0 dB.

3 - When the threshold is exceeded by the signal, you will see the Gain Reduction indicated on the single Attenuation meter to the right. Set the threshold about 4 to 6 dB lower than the peaks in the Threshold meters (input). You will see that you have from 4 to 6 dB of peak limiting indicated on the Attenuation meter (output).

4 - You'll notice that as you pull the Threshold down, the output level goes up. Leave the Out Ceiling setting at the maximum value: 0.0 dB. This is your maximum peak output.


Notice that you have significantly increased the output level. If your threshold is at -12dB, then you have pushed the signal up 12 dB (not that this is recommended)! With moderate gain reduction, the maximum level of a file can be significantly increased with minimal audible effect.

Only the signal above the threshold is limited; all signal below the threshold has a constant gain change that is controlled by the difference between the Threshold and the Out Ceiling. It is this function of the L2 that allows you to maximize the level with the amount of headroom desired.



5 - Leave the ARC™ (Auto-Release Control) engaged at all times unless you have a very specific type of limiter behavior, such as pumping or distortion effects. ARC calculates the release time every sample for optimum level with minimum artifacts. For nearly every possible source, ARC will outperform a fixed release time.

THE IDR SECTION



After reading this overview, make sure to read the chapter - **Important IDR info** - which will offer you complete IDR technical information so that you may be able to better choose from the IDR options in L2. Take the time to read about this new technology so that you may achieve the best processing for your application.

In the L2, there are two types of Increased Digital Resolution dither technology: **type1** and **type2**. The differences between them will likely influence your choice for your project.

- *type1* gives no nonlinear distortion.
- *type2* exhibits lower dither level.

It is up to you to choose no distortion or lower dither level for your soundfile. More information is in the section called **IDR Dither Type & Noise-shaping combinations**.

BASIC CONTROL OF IDR

- Select Quantize level for desired output (24, 22, 20, 18, or 16-bit) by clicking the button to toggle through the choices. Alternately, if you option-click-hold on the Quantize button, a popup menu will allow direct selection of bitdepth.
- Select Dither (type1, type2, non) by clicking the button.
- Select the type of noise-shaping (Moderate, Normal, Ultra, non) by clicking the Shaping button. (The following chapter has extensive information on the combination of these 3 options.)

For optimum results, level maximization (and at least normalization) should also be done! This is why the L2 includes both advanced peak controlling and IDR together: one step maximizes both.

TYPE1

This is the 'purist' technology. It is designed for no nonlinear distortion or modulation noise at low levels and combines optimal dither noise with psychoacoustic noise shaping. When signals are subjected to several stages of higher-resolution processing and re-quantizing back to lower bit-depths, the design of resolution-enhancement must satisfy different requirements than a technology designed just for one-stage use. If applied several times in succession, a digital resolution enhancement technology optimized for one-stage CD mastering can produce unwanted side-effects.

Waves **type1** technology, however, is the first optimized for use at every processing stage, allowing for the effects of cascading and subsequent signal processing. Type1 is also optimized to cause minimal side effects when used with stereo signals.

Type1 is the recommended choice for use with 20 and 16 bit files processing and other high quality mastering applications. By combining level maximization (peak controlling) and IDR processing, 16 bit files created from 20 or 24 bit masters this way can have an apparent resolution of 19 bits — an 18dB improvement.

TYPE2

Type2 also uses dither with a similar noise-shaping curve, but the dither is of a unique kind designed to minimize the amount of noise added, thereby giving a lower noise level than the IDR type1 process, but at the expense of some low level distortion.

Type2 does have some advantages for high quality mastering as well, and it is purely your choice whether the ultimate in low distortion of type1 is preferred, or the additional reduction in dither level of type2. Type2 is "auto-blackening" with no input signal, in other words, if the input signal is digital black (no signal at all) there is no Type2 dither added to the output signal.

NOISE-SHAPING OPTIONS

Another way to decrease the perceived amount of noise and increase perceived resolution is to 'shape' the frequency content of the noise so it matches the ear's sensitivity curves. In basic terms, noise-shaping shifts the energy of the noise to the frequency ranges where we hear it the least.

The three options of Noise-shaping provided in L2's IDR section push more of the noise energy to higher frequencies above 15 kHz, where our ears are least sensitive, and reducing the noise energy of lower frequencies. The three Noise-shaping options differ in the amount of this 'shifting action'.

Moderate is the lightest noise-shaping curve.

Normal is the recommended option for use under most conditions for all bit-depths.

Ultra is a very high-quality setting, suitable only for use at the very last stage of mastering high-resolution files (16-bit and higher) targeted for high-quality digital media. It is theoretically possible that the relatively high amount of high frequency energy could cause undesirable side effects *if the signal is going to be processed or digitally edited again*. Therefore, it is best that Ultra is used in the last stage of file preparation. However, with many thousands of L2- processed products, no such situations have been reported. Since it is *theoretically* possible, we wish to inform you of it. These theoretical side effects might cause clicks at later edit points if 'Ultra' noise shaping was used, *if played back on poorly designed D/A converters*. However, state-of-the-art designs rarely have these troubles anymore.

Of course, the effect of Noise-shaping is even greater when used with type1 or type2 dithering, since Noise-shaping reduces the audibility of the added dither noise. Now try the full effect of IDR technology by listening to the same material, with both IDR types and different kinds of noise shaping. The most obvious places to examine are notes or reverb during the end of the sound, or 'tail'; it is during this time that quantization error is most audible, although it is present on all low-level signals (such as elements that are soft in a mix, etc). Since the entire issue of dithering is a very subtle one, we recommend you listen to a rather long piece of audio (2-3 minutes) of high-quality, say 20-bits if possible, with a good dynamic range. Jazz and classical recordings are ideal. If you don't feel you fully understand the tradeoffs between IDR and Noise-shaping settings, the option that will generally work well for CD-mastering is type1 with Normal noise-shaping. For minimum noise with 16-bit and greater files, type2 Ultra; maximum resolution, type1 Ultra.

ARC (AUTO RELEASE CONTROL)

Waves ARC technology is part of the system that enables greater maximum levels, and like the lookahead technique, does so with a minimum of artifacts (unwanted distortion). The ARC (Auto Release Control) algorithm is designed to dynamically choose the optimum limiter release value for a wide-ranging input. ARC reacts much like a human ear, and can produce significantly increased RMS (average) levels with excellent audio clarity. The release time is simply how fast the limiter stops limiting, and the ARC system controls this release time.

In many uses of limiters (and other dynamic processes), the exact choice of time constants (attack and release) is set as a careful choice of reaction to RMS and peak transients. (Peak transients are sudden sharp sounds, such as drum attacks or any sound with a sharp beginning; RMS means root-mean-square, which represents the average level). To allow more gain reduction with fewer artifacts, the release time must change to control different parts of

the signal. In ARC, similar to the human ear, RMS and peak transients are analyzed and reacted to differently. In general, the release is faster for peak transients and slower for the overall RMS level. Fast release times with certain signals (such as bass signals, or high RMS levels) can produce unwanted distortion artifacts, so the Waves ARC system automatically updates the right release time constantly for the best possible sound, no matter what the input is.

In the case of the L2, fast control of transients and slower release times for sustained energy (RMS) above the threshold are equally essential to achieve maximum level and minimum distortion or artifacts. The ARC technology can be switched in and out, as some users will need the L2 only for light limiting; however, we recommend leaving the ARC auto-release controls "in" at all times. This ensures that even for light limiting, the release time is as fast as possible with minimum artifacts. Additionally, for some of today's demanding music genre and production "values", the ARC allows very heavy limiting with astoundingly clear sound.

Chapter 4 - In-depth IDR Information

ABOUT DITHER AND NOISE SHAPING

Dithering and Noise shaping are two independent, but complementary, techniques to improve the perceived quality of sound after it has been re-quantized. As will be explained here in some detail, each technique is responsible for the improvement of a different subjective quality of the noise imposed by re-quantization. Therefore, each can be used separately to improve that specific quality.

Dithering is done in order to change the character of the quantization noise to more closely resemble analog hiss, rather than digital quantization noise. The main effect of dithering is to reduce (or, in case of type1, virtually eliminate) all correlation between the quantization noise and the original signal, thus reducing (eliminating) non-linear distortion typical of digital quantization noise. The dithering process 'exchanges' these distortions for a steadier analog-hiss quality signal.

Noise shaping is done in order to optimize the distribution of overall noise energy across the spectrum. This optimization is according to the ear's sensitivity. This means a decrease in noise (whether distortion or hiss) in the ear's sensitive areas (1 to 6kHz), is 'exchanged' for an increase of noise in less sensitive areas (above 15kHz, toward Nyquist).

Hopefully this has helped you see that in both techniques, the issue is about 'exchanging' the character and frequency content of noise (hiss & distortion) according to subjective criteria.

How do these processes help 'capture' 3 more bits of detail? The easiest analogy is to point to dithering in graphics, which is exactly the same process, and exactly the same type of psycho-perceptual model. The brain is capable of perceiving detail that is lower than a noise floor (in this case, dither).

However, quantization noise is highly correlated to the signal; in other words, it is related to and governed by the signal. Dithering makes this noise become uncorrelated (as dither is a random signal), therefore allowing the brain to perceive the detail. The noise-shaping then helps to shift the energy of the noise to a less sensitive area of hearing.

DITHER

1. No dither <off>. This is not quite normal truncation but rounding to the 24th bit without adding any dither. It gives a high degree of nonlinear distortion at low levels. There is actually no reason to turn dither Off in normal use. The singular reason would be to provide a 24bit transparent (perfect clone) output from the input; it's easier to just use the Bypass button! Even if you are using an external dithering system to achieve your final desired wordlength, you should use the L2's IDR to dither to a 24bit output.
2. IDR type1 dither. This dither is a wide-band dither. type1 adds a certain amount of noise causing a 5dB increase in background noise compared to no dither. It completely eliminates all low-level distortion and signal-dependent modulation effects. The result is a very transparent and clean low-level sound with a high resolution, most resembling the steady low-level

hiss of an excellent quality analog system, and in lieu of digital quantization noise. This is the "purist" technology. It is designed for no nonlinear distortion or modulation noise at low levels, and combines optimal dither noise with psychoacoustic noise shaping. type1 is also optimized to cause minimal side effects when used with stereo signals. type1 is the recommended choice for use when processing high quality mastering applications. By combining level maximization (peak controlling) and IDR processing, 16bit audio created from 20 or 24bit masters this way can have an apparent resolution of 19 bits, more than an 18dB improvement! When signals might be subjected to more than one stage of processing and quantization back to 16bits, the design of resolution-enhancement must satisfy additional requirements than a technology designed just for one-stage use. If applied several times in succession, a digital resolution enhancement technology optimized for one-stage CD mastering can produce unwanted side-effects. Waves type1 technology, however, is the first optimized for use at every processing stage, allowing for the effects of cascading and subsequent signal processing, if needed.

3. IDR type2 dither. This dither is a narrow-band dither, adds virtually no audible noise, and so is nearly 5dB quieter than type1, but with some low-level distortion. However, this distortion is generally much lower than with no dither at all. type2 is of a unique kind designed to minimize the amount of noise added, thereby giving a lower noise level than the IDR type1 process, but at the expense of some low level distortion. type2 does have some advantages for high quality mastering as well, and it is purely your choice whether the ultimate in low distortion of type1 is preferred, or the additional reduction in noise of type2.

NOISE SHAPING

Another way to decrease the perceived level of noise is to "shape" the frequency content of the noise so it matches the ear's sensitivity curves. In basic terms, noise shaping shifts the noise to the frequency ranges where we hear it the least. The three options of Noise shaping provided on L2 push more of the noise energy to higher frequencies above 15kHz and toward Nyquist, where our ears are least sensitive, and reducing the noise energy at lower frequencies. The three Noise shaping options progressively differ in the amount of this "shifting action". The L2 features ninth-order noise shaping for optimal wordlength reduction quality.

1. **Off.** No noise shaping, resulting in more audible noise, (and distortion if dither is not used). The result has equal noise (distortion) levels at all frequencies, which is not optimal from a psychoacoustic point of view.
2. **Moderate.** This typically reduces perceived hiss (or distortion if dither is not used) by around 6dB. The HF noise gain is about 9dB for 44.1kHz.
3. **Normal.** This typically reduces perceived hiss (or distortion if dither is not used) by around 12dB. The HF noise gain is about 15dB for 44.1kHz. (In addition to being very suitable for creation of Production Masters, using Normal with type1 dithering was also designed to be excellent for masters that would be processed again for any reason, including consecutive re-dithering, with an accumulation characteristic that is optimized to be minimal.)

4. **Ultra.** This gives the greatest perceived hiss/distortion reduction, typically 18dB. The HF noise gain is about 23dB for 44.1kHz. Ultra is a very high-quality setting, suitable only for use at the very last stage of mastering high-resolution audio (16bit and longer wordlengths) targeted for high-quality digital media. It is best to use Ultra in the last stage of audio preparation (Production Master). Due to the HF gain, it is theoretically possible that the relatively high amount of high frequency energy could cause undesirable side effects if the signal is going to be processed or digitally edited again. However, with many thousands of L1-processed masters (L1 is the software plug-in that led to the development of the L2 hardware), no such situations have been reported or observed. Since it is theoretically possible, we wish to inform you of it.

You can hear the effect of Noise shaping by itself by setting Dither type to Off and selecting one of the Noise shaping options while monitoring the output. Of course, the effect of Noise shaping is even greater when used with type1 or Type-2 dithering, since Noise shaping reduces the audibility of the added dither noise.

Now try the full effect of IDR technology by listening to the same material, with both IDR types and different kinds of noise shaping. The most obvious places to examine are notes or reverb during the end of the sound, or "tail". It is during this time that quantization error is most audible, although it is present on all low-level signals (such as elements that are soft in a mix, etc.)

If you don't feel you fully understand the tradeoffs between IDR and Noise shaping settings, the option that will generally work well for CD-mastering is type1 with Normal noise shaping. For minimum noise with 16bit and greater sources, type2 Ultra. For maximum resolution, type1 Ultra.

The noise reduction figures given here apply to sampling rates of 44.1 or 48kHz. They are even better for doubled sampling rates. If audibility of noise were the only factor, the choice would almost always be to use Ultra noise shaping, but in some situations, heavy noise shaping of the Ultra kind can theoretically have some disadvantages, and the alternate settings, such as Normal or Moderate may be better.

For 16bit applications, Ultra shaping should be avoided in the following situations:

(a) Subsequent digital editing, when the signal is subjected to later editing. At the edit points, an extreme noise shaping might cause low-level, yet audible "clicks" in rare cases when played on very inexpensive CD players. An example application in which you would avoid use of Ultra shaping would be on CD's with production music or sound effects libraries that would certainly be subject to further digital editing.

(b) Poor Error Correction, i.e. when a signal is destined for a carrier medium with poor error correction, such as CDs pressed in pressing plants with poor quality control. When errors that are not properly corrected occur, the Ultra setting, like all forms of heavy noise shaping and other resolution enhancement technologies, tends to cause audible background crackles, especially on very cheap CD players. While these effects generally don't occur on the majority of mid- or hi-fi CD players, they can be noticeable on very cheap products. The amount of such crackles on poor pressings is

greatly reduced by the Normal noise shaping. Of course it can be argued that listeners with very cheap players may not be interested in high-quality reproduction at all!

(c) Subsequent equalization before duplication (not by the user!), when heavy treble boost equalization is subsequently employed. (Again, we're referring to EQ by a pre-mastering person before duplication, not if the user boosts the treble on their system. For the user to do so is fine, and is up to them.) This can cause the strongly boosted higher frequencies used by Ultra noise shaping to become so high in level that they might feed excessive noise energy into loudspeakers. Therefore Ultra shaping is best avoided in situations where subsequent equalization may be used in pre-mastering, such as in a compilation. However, if the mastering is done correctly the first time, large treble boosts would be quite unnecessary, and is somewhat moot in this context. Much less HF gain is used with the Normal and Moderate noise shapers (all of which could be observed on a real-time analyzer, such as Waves PAZ psychoacoustic Analyzer).

Chapter 5 - 16-bit (and higher) mastering

Here are the basic steps of using L2 in a 16-bit, 44.1/48kHz application.

These steps also apply to 24, 22 and 20-bit mastering.

- All processing, EQ, sample rate conversion, dynamic changes, etc. **MUST** be done before L2 processing. The L2-Ultramaximizer should be the last processing of the file. Ideally, dithering occurs only once.
- Using a 16 or higher bit input file, set the Threshold for desired peak limiting. For suggestions on how much limiting to do for certain applications, see the Specific Applications chapter. In general, set Threshold for about 4-6 dB of Gain Reduction in the Attenuation meter.
- Now take the Output Ceiling up to the maximum peak output you desire. You can take this Output all the way to 0.0 dB without any clipping. For CD's, a recommended setting is -0.3dB; for more information read the Peak Clipping section. Factory Presets already set the Out Ceiling control to the recommended value.
- Leave the ARC (auto-release) button engaged all the time.
- Set Quantize output for **16-bit** (for CD/DAT; or 22, 20, 18 for higher archival or mastering medium if your hardware supports the transfer of 16+ bits).
- Set Dither type (type1 or type2). **IDR type1** is recommended for most high-resolution applications.
- Set Shaping (Moderate, Normal, Ultra, none). **Ultra** and **Normal** are recommended for most high-resolution applications.

Chapter 6 - Use the L2 last

It is recommended that L2 be used as the final process after all dynamic and EQ adjustments have been made.

Only when all these processes are finalized should the question of peak level be addressed. Instinctively, it might seem appropriate to Normalize the file once all other processing has taken place, but in practice, it may be better to set the peak levels to around 1dB below clipping using L2. For an explanation of this reasoning, see the notes on digital clipping.

The choice of IDR setting depends on the final use to which the file will be put. Type1 or 2, Normal, is recommended for most work. Type1 or 2, Ultra, is considered best for final mass production of 16bit and greater masters, and for producing a complete disc master which will undergo no further edits. An example of this would be a production master CD run off from a hard disk editing system in a single pass and where no further editing is anticipated (all timing is finalized). This PMCD would then be transferred unchanged through to the glass mastering process.

When using programs like MasterList or MasterList CD which in effect string together separate regions of the soundfile, you must perform destructive processing with the L2. Since MasterList/CD is a form of editing (playing regions/songs with silence or splices between audio regions/songs), it is recommended to use Normal noise-shaping with either Dither type1 or type2. However, thousands of IDR-processed masters have been produced with type1 Ultra with no unusual problems whatsoever.

If you must process or EQ any previously L2-processed file, you will need to create headroom by lowering input on those later processes, then probably re-limit to restore the average level.

Chapter 7 - Important notes on digital clipping

The digital words representing an audio signal at each moment have a maximum possible positive value and a minimum possible negative value defined by the bit depth of the file format. Any attempt to force an audio signal beyond these maximum permitted values, for example, by applying excessive gain, will result in the audio signal being clipped. Clipping distortion generally sounds quite unpleasant and is to be avoided.

However, there are other ways in which a signal can become clipped, and some of these are far from obvious.

PEAK-NORMALIZED SIGNALS

A '**Normalize**' process allows a file to be processed in such a way that the maximum peak level within the file just reaches (but does not exceed) the digital zero or clipping point. This is obviously desirable as it means that the file is as 'loud' as is possible without clipping, and in turn, this maintains the best signal-to-noise ratio, especially at low bit word lengths.

In situations where a higher average sound level is required, L2's peak limiter allows the typical level of signals to be even further increased by gently pulling down the gain, without audible nonlinear distortion, of waveform peaks. L2 can simultaneously re-scale the audio data so that the limited peak signals approach or just reach digital zero.

However, by storing soundfiles at the maximum possible level, there is a risk that any subsequent processing may take these peak levels too high, resulting in clipping distortion. Peak limiting to 0dB, by whatever means, leaves no margin for any subsequent increase in peak level. Intuitively, you might think that simple gain reduction could be applied without incurring the risk of clipping, and equally that any increase in gain would be sure to cause clipping. This is true. You might also think that applying an EQ boost at any frequency might result in clipping depending on the peak energy level within the band being equalized. Again this is true.

What is far less obvious is that applying an **EQ cut** also runs the risk of causing clipping. To prove this would take a lot of math, but the following description should help get the point across.

At any instant, the peak level of a signal may be the result of several components at different frequencies and at different phases relative to each other. Some components will add while others will subtract, but what happens if you 'EQ out' a frequency that would otherwise be subtracting from the peak level by virtue of its phase? The peak is now *higher* than it was. For most audio material, this effect will be relatively small, typically increasing peak levels of the order of 0.3 dB, but is possible that under unfavorable circumstances or with non-typical signals, the peak level could increase significantly more than this.

If an effective peak limiter like the L2 is used, because it forces the signal to skim the peak level more often, this likelihood of clipping in this way is further increased.

In practice, filters attenuating mid or high frequency components tend to cause the kind of increase in peak level described, but high pass filters that attenuate the bass can sometimes cause much larger increases of peaks of

the order of several dBs on heavily limited signals. The phase response of certain high or low-pass filter types can also increase peak levels by up to 4 dB or thereabouts.

Bearing in mind what's just been said, it might seem logical to keep the signal peaking a few dB below digital zero until all processing has been carried out. After that, you can safely normalize the signal—or can you?

A related problem with peak clipping can arise when a normalized soundfile or signal is converted to a new sampling rate. The reason has to do with the sample-rate conversion process itself, and during sample rate reduction, the signal is effectively being filtered; the available audio frequency range is smaller at lower sampling rates. Such filtering can increase peak sound levels in exactly the same way as attenuating equalizers can. But even when increasing sampling rate, an increase of peak level can occur. This is because the continuous-time audio waveform is represented in the digital domain only by its values at the sampling instants. It is perfectly possible for the peak value of the continuous-time audio waveform to occur at instants lying between two sampling instants, and thus to be higher than the peak value at any of the sampling instants. When changing the sampling rate, new sampling instants are chosen for the continuous-time audio waveform, and these new sampling instants may coincide with an increased peak lying between the original sampling instants. This is especially likely to occur with signals with a lot of high frequencies, since such signal waveforms change more rapidly between the sampling instants.

Though artificially contrived signals can be created to really show up this problem, in real life an attenuation of at least 0.3 dB or so prior to conversion should provide adequate protection against clipping. You might expect sample-rate-converter designers to account for the possibility by designing in a small amount of attenuation, and the cheap ones generally do not.

But can you safely normalize a file that you know is at the final sample rate? Unfortunately not, because many compact disc players (and some other digital consumer equipment) use over-sampling digital-to-analog converters (DACs) to produce the analog signal fed to the amplifier. Such over-sampling converters involve a sampling rate conversion process which can (and does!) cause audible peak clipping. Once again, some designers appear to have overlooked this problem, although not as widely as they did in earlier DAC designs.

Chapter 8 - Recommended settings

RECOMMENDED IDR SETTINGS

Any combination of dither and noise shaping can be used, but the following settings are particularly recommended for different applications.

- General Purpose high-quality use, including material liable to be edited, EQ'd, and re-dithered: type1 – Normal.
- Lowest Noise: type2 – Ultra.
- Low Noise/Highest quality (final production masters): type1 – Ultra.
- Low noise while allowing editing/EQ: type2 – Normal.
- High Quality, with lowest risks of spurious noises on edits or cheap CD players: type1 – Moderate.
- Low noise, with lowest risk of spurious noises on edits or cheap CD players: type2 – Moderate.

RECOMMENDED PEAK LIMITER SETTINGS

The recommended settings for the L2 Peak limiter would be highly related to the processed program material. A general note we can offer is to use ARC for the optimal release time and the smoothest sounding results. The Peak-Hold numeric field under the Attenuation meter can also provide some important feedback, specially if you see a number under -10dB down to -30dB . In this case you will know that major attenuation is happening and you should be aware where in the program this happens and make sure it sounds good to your ears!

For commercial Audio CD, some mastering engineers recommend a ceiling of -0.3 dB to be safe from distortion due to playback system deficiencies.