CPARIS Web Site
Mastering strategies

Contents

1 What it is

Mastering is the process of 'sweetening' a collection of two-channel stereo titles so that it feels musically connected. Mastering is not meant to reinvent the sound of a mix or interfere with its character. However, mastering is likely to improve mixes that are perceived as dull, harsh, lifeless, 'muddy' or noisy.

1.1 Mastering

Compare mastering to the task of a book editor who turns a raw manuscript into a book.

The book editor must understand grammar, syntax and writing styles as well as binding techniques, colour separations and the printing processes.

Likewise, the mastering engineer must turn an audio project into a cohesive sounding product optimised for a particular broadcasting medium: CD, DVD, film, internet, multimedia, radio, television, theatre and/or video.

1.2 Equipment

The very best equalisation is to apply no equalisation at all.

To produce analogue masters from digital mixes, the only devices we've ever used were digital to analogue converters and analogue dynamic processors.

In 2013, we parted with our 1959 Fairchild 670 (stereo valve limiters) and our 1996 Crane Song STC-8 (solid state compressors).

We used to enjoy tweaking their potentiometers and maintaining their components. The French National Audiovisual Institute in Paris still has plenty of spares from the ORTF days.

The more time we spent with our analogue instruments, the more we learned to respect their imperfections and the more we enjoyed their 'warm' company.

In 2014, we went back to the source and wrote our 'dream' specifications.

We then commissioned a Swiss audio hardware manufacturer to design our mastering setup with the following specifications:

- Analogue signal processors with 1 µs attack times.
- BBC-type PPMs.
- Fixed processor ratios.
- Harmonic controllers.
- Infrasound analyser.
- No 'auto' functions.
- No dithering.
- No equalisation.
- No limiters.
- No phase shifts.
- Threshold detents in cB.

It took a year to build.

Many software companies have modelled the behaviour of analogue signal processors into digital emulations but we still find it more exciting and rewarding to operate in the analogue domain.

The actual design is a little more complex than what is presented here but we won't be disclosing any schematics to encourage reverse engineering pirates to work it out for themselves.

1.3 Loudness

Our masters may not sound quite as loud as the commercial big time products (one decibel softer on average) but our clients have no problem with that for the following three reasons.

1. They have no intention to compete with the big names.

2. They would rather hear the full dynamic range when their music is played on a CD player.

3. When they create a streaming file, the conversion to AAC or mp3 will do the trick for them (the additional distortion in the conversion providing the extra decibel without even trying).

1.4 Documentation

We cannot stress enough how important it is to document a work in progress. Given that we often work on four projects at a time, it is paramount to recall what we last did, why we did it and how we did it, especially when the clients asks for changes on a complex project.

1.5 Note

This document does not cover the administrative aspects of mastering, such as sequencing the titles in the order selected by the artists, including metadata (Universal Product Code and International Sound Recording Code) and setting up the spaces between the titles.

2 How it works

Our strategy consists of dividing the analogue signal into fourteen adjacent frequency regions in order to process their signal independently from each other.

There is no rule as to which region(s) should be processed and the values listed in the table below are our own preferences, which may not work for you.

2.1 Workflow

1. The digital stereo mixes are converted to analogue through a 72-bit digital to analogue converters (Nagra HD DAC). Two analogue stereo outputs are available: direct and delayed.

2. The direct analogue signal feeds the fourteen passive band pass filters, whose outputs drive the Analogue Stereo Level Processor (ASLP) side chains (they are not part of the audio chain).

3. The delayed analogue signal feeds the fourteen ASLPs.

4. Each ASLP has three analogue stereo components in series:
   - solid state downward expander (fixed 1:10 ratio).
   - valve upward compressor (fixed 2:1 ratio).
   - valve harmonic controller.

5. The fourteen ASLP outputs feed their fourteen Analogue to Digital Converters (ADC).

6. The ADC outputs are combined through an ultra high speed digital stereo multiplexer.

7. The stereo signal can then be converted into the digital stereo format required by our client.

This workflow shows that our mastering signal path is not 100% analogue but 99% analogue and 1% digital magic.

2.2 Listening

The very first time we listen to a mix is purely for enjoyment (and there's nothing wrong with that).

2.3 Expanding

Then, we listen again and identify the audibility threshold (usually around −72 dB) where the decay of a sustained instrument no longer needs to be heard.

This threshold is only adjusted on those regions that require it, an easier task with dance or electronic music mixes than with classical or jazz music mixes.

2.4 Compressing

Next, we focus on the timbre.

Identifying and selecting the correct frequency regions to process is made difficult by the fact that octave relationships, especially in the low part of the spectrum, are easily confused by the brain.

We adjust the thresholds of the compressors only for the regions that are identified as containing excessive spectral components.

This is similar to mixing levels although we balance timbres.

The threshold detents are calibrated in centibels (10 cB = 1 dB) so we have 31 steps to play with (from 0 cB to −30 cB).

Last, if necessary, we adjust the release times of the compressors that require it (the attack times are all internally preset to 1 µs).

2.5 Harmonising

The fourteen harmonic controllers enable us to add a mild amount of distortion to the transients by intensifying their existing harmonics and/or by artificially creating new ones.

There is a common school of thought that says that even-order harmonics are 'warm' and musical and odd-order harmonics are 'powerful' but unmusical.

When manipulating harmonic controllers, it is not the level of the epicentres that will be affected but that of their even- and odd-order harmonics.

Enhancing the second harmonic of the 'presence', for example, will affect the region centered on 1398 Hz, not on 699 Hz.

2.6 Epicentres

In this table, the regions we most often reduce are indicated by the letter C (for compress) and those we most often enhance by the letter H (for harmonise).

But again, this is only our rule.

Character	f_C (Hz)	
Sizzle	12,543	C
Dust	7,752	C
Anger	4,791	C
Click	2,961	C
Definition	1,830	H
Bite	1,131	C
Presence	699	H
Nose	432	C
Middle	267	C
Warmth	165	H
Mud	102	C
Boom	63	C
Bass	39	C
Weight	24	C

Table 1: The fourteen epicentres.

2.7 Stereo anyone?

Finally, we switch our monitoring system from mono to stereo.

We never tamper with the space and perspective the artists and mixing engineer have created.

Interestingly though, as the overall sound becomes clearer during the mastering process, the stereo image becomes slightly wider while maintaining the original proportions intended by the mixing engineer.

Incidentally, a mixing technique we particularly enjoy is when the mixing engineer has increased the contrast between the verses and the choruses as follows.

VERSES: Fewer active tracks and/or lower levels and/or more low frequency cuts and/or more high frequency cuts and/or low compression ratios and/or narrower sound stage.

CHORUSES: More active tracks and/or higher levels and/or less low frequency cuts and/or less high frequency cuts and/or high compression ratios and/or wider sound stage.

At this stage, we also check the wanderings of the polarity, just in case our ears missed some stray low frequencies that cancelled each other out while we were monitoring in monophony.

Stray low frequencies are a nightmare when cutting vinyl records.

2.8 Examples

The examples on our web site are mp3 files (internet oblige), which somewhat defeats the purpose of mastering.

However, despite the mp3 artefacts, it should still be possible to detect most of the ‘before’ and ‘after’ differences.

If you monitor these files on a laptop computer, be aware that, because of their loudspeaker size, the frequency spectrum will not be fully represented.

The headroom was matched so that you can focus on the timbre rather than the level differences.

The examples are presented in a repeating sequence, starting with the unmastered mix followed by the analogue master.

The switch takes place every ten seconds (watch the elapsed time if you get lost) with a short (unmusical) crossfade.

Unless the client specifically requested to ’make it louder’, don’t expect vast differences.

As mentioned in the introduction, mastering is not meant to reinvent the sound of a mix or interfere with its character.

3 Why it works

People are always surprised that our ‘corner shop’ is still active despite the fact that our rates have been the same for 25 years and despite the number of $0.99 mastering plug-ins available.

Regardless of the artificial intelligence of the software, it is not about to replace the human element in people relationships.

Most engineers still prefer to have an outsider master their mixes as they value a fresh perspective on their work.

This may change in a hundred years but by then, there won’t be any more man-made music to mix, let alone to master.

Don’t waste you time trying to connect with people who know people who know people (repetition intended) that might lead you to work with the big names in the industry.

The big names only trust people with experience, which no one has at the start of their career.

Starting at the bottom, we decided to only produce high quality masters even if it meant turning down clients. The first ten years were not easy but, as a result of our persistence, artists started to ask us to master their next album or passed on our name to their friends.

Eventually, it is their word of mouth that did our advertising.

This is called ‘the long tail’, also known as ‘niche marketing’.

Over 25 years, we’ve mastered the works of about 900 artists.

4 How we monitor it

When monitoring audio, the combination of head, shoulders and auricles provide the brain with directional and spectral data. Wearing headphones prevents the brain from accessing this information.

Consequently, we don't recommend using headphones (even the electrostatic types). Don't fret though because whatever sounds good on loudspeakers also sounds good on headphones.

4.1 Monitoring by ear

4.1.1 Low levels

During a mastering session, we keep our monitoring level constant at approximately 40 dB(A) on a set of Earthworks Sigma 6.2 wide range loudspeakers.

It is only at the end of the session that we check the mastered title at a level of around 90 dB(A).

The last adjustment we might then make would be to tame some frequencies below 20 Hz that might otherwise interfere with the key of the music.

4.1.2 Mastering suite

Our mastering suite is insulated from the rest of the building. However, we have selected not to treat its internal acoustics.

The reason is that the acoustics of the space becomes insignificant when we monitor audio at around 40 dB(A).

At those soft levels, the difference between the 40-phon contour (top curve in the figure) and the 90-phon contour (bottom curve) is approximately 40 dB at 30 Hz, a level difference we have learned to estimate by ear.

4.1.3 Advantages

Monitoring at a very low level highlights what represents the heart and soul of the mix: the balance between the background and (1) the 'message carrier' (melody or solo), (2) the 'high frequency beat carrier' (snare drum or equivalent) and (3) the 'low frequency beat carrier' (bass drum or equivalent).

It is also at this level that the non-linear characteristics of our own ears, our amplifiers and our loudspeakers are at their lowest.

It is easier, for example, to detect the tiniest digital distortion clicks at this level (just try it!).

But to us, the greatest advantage is that monitoring at a very low level works regardless of the music genre we are mastering.

4.1.4 Monophony

If the title is destined to be broadcast on the internet, radio or television, to cut a vinyl record, to encode AAC or mp3 streaming files or to be played in a live venue (since most of them don't bother about stereo), the master must contain a large proportion of the mono signal.

So we monitor in monophony, which also enables us to notice when instruments drop in level (or vanish) from the mix and to ensure mono compatibility.

4.2 Monitoring by eye

Twenty five years ago, we settled on the ballistics of the BBC-type Peak Programme Meters (PPM) whose response is the closest to that of our ears.

Apart for this meter, at which we glance for a quick second opinion, we turn off the 'light shows' of our equipment to avoid visual distractions from the audio.

We don't need to watch spectrograms as we've learned to identify by ear the major constituent frequencies of a sound (Spectral Solfège training).

The only exceptions are when we need to 'track' a singer's sibilants (they love playing moving targets) and when we need to identify the frequencies of instruments whose harmonics may conflict when combined.

We've also trained ourselves to gauge by ear the approximate amount of compression a sound may have been subjected to.

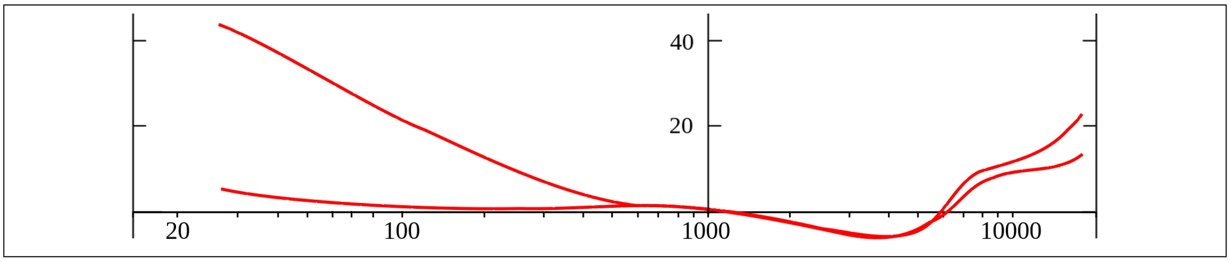

Figure 1: Equal-loudness contours vertically offset to match at 1 kHz.

5 Summary of our 2015 blog

The aim of this blog was to discuss the features of our mastering setup. The blog was closed (see 2013 blog).

5.1 Managing levels

Compression provides additional headroom by reducing the impact of transients.

The rest of the signal can then be allowed to rise, which we perceive as the body of the music being louder.

At the same time however, since the transients are reduced, the loudspeaker excursion (the extent of the back and forth cone motion) is narrower, the cone does not move as far as before and there is less thrust from the motion of the air that carries the sound from the loudspeaker to our ears.

The music has less depth, impact, openness and punch.

As a result, when we compress, we need to compromise between the perceived loudness and the dynamic contrast.

For this section, we use the compressor as an example but the same principles apply to any time-dependent processors, such as expanders or effects.

5.1.1 Ratio

The ratio of a compressor is the amount of compression, expressed as the proportion between its input and its output.

In our compressors, the ratios are fixed at 2:1 by design, which means that we don't use limiters.

5.1.2 Threshold

We adjust the thresholds (the levels above which the compressors are active) to never exceed 2 dB of gain reduction.

In our system however, the signals shaped by the band pass filters drive the signal chains of the compressors at the regions precisely selected by the filters.

We also use 'average follower' compressors which apply compression to the signal's RMS values (average energy), which is how we perceive loudness.

5.1.3 Attack time

The attack time is the time taken by the compressor to activate.

Fast attack times trigger the compressor to quickly respond to incoming transients. On percussive instruments for example, this quickly reduces the impact of a transient.

Slow attack times trigger the compressor after the transient has passed, letting the transient through unaffected.

In practice, we start with a slow attack time then shorten it until the transients are tamed but not suppressed, a subjective decision based on the music genre.

5.1.4 Release time

The release time is the time taken by the compressor to cease to be active.

Fast release times trigger the compressor to quickly return to its inactive state. On percussion instruments for example, this results in a pumping sound.

Slow release times triggers the compressor to gradually return to its inactive state. However, it keep compressing the softer levels that follow a transient.

In practice, we start with a fast release time then lengthen it until the pumping is inaudible, another subjective decision based on the music genre.

5.1.5 Make-up gain

The make-up gain is a convenient feature that makes sure that any gain reduction is compensated for by an equal amount of extra output gain, ensuring an equal output level whether compression is applied or not.

5.1.6 Delay

The period of a waveform is the time it takes for a signal to complete a cycle.

It is inversely proportional to its frequency:

$$T = 1 \div f$$

A 20 Hz signal, for example, has a period of 50 ms.

Disturbing the formation of a waveform's envelope creates distortion and this happens whenever the attack (or release) time of a compressor is shorter than the period of the signal.

In the previous example, applying an attack (or release) time shorter than 50 ms will distort the signal.

What this means is that reducing the impact of a transient with a compressor at that particular frequency will not be possible without introducing distortion.

The solution is to feed the signal to the side chain of the compressor (as a control signal) and add a delay to the audio signal.

The side chain can then 'see' the signal ahead of time. Our fear of distortion may be put to sleep.

6 Summary of our 2014 blog

The aim of this blog was to discuss the features of our mastering setup. The blog was closed (see 2013 blog).

6.1 Managing dither

Dithering is the process of adding a low level noise to the signal to minimise the distortion due to the quantisation noise that is inseparable from digital files at low levels.

Dithering is only relevant when the quantisation noise is likely to be perceived such as near the end of fade outs.

The dithering modes are designed to act in the regions where we are the most sensitive: 2 to 5 kHz (needed for speech intelligibility) and 12 to 15 kHz (needed for spatial localisation).

If you work with 24-bit (or higher) files, there is no need to apply dithering since the low levels are too low for us to hear.

If you work with 16-bit mixes, you will need to use one of the dithering modes available on your system prior to mastering.

Don't delve into the technical differences between the various modes. This is an area for people who love debating. Skip the debate, test a few modes by ear, select the one whose workings you actually don't hear and move on.

1. Listen to the last seconds of the end of a mix at around 90 dB(A).

2. Don't turn up the monitoring level to hear the effect of dithering during the fade out as this would defeat the purpose.

3. If the background noise quality does not change, then the dithering mode you have selected is just fine.

The other solution is to switch to 24-bits from the start, not dither anything and pass the dilemma to the mastering engineer.

6.2 Managing timbres

Our band pass filters are as simple as they come, the oldest analogue design ever created.

To work on timbres, we divide the frequency spectrum into adjacent frequency regions with passive band pass filters.

It enables us to subdivide the spectrum into a number of adjacent regions to be processed independently from each other.

The three parameters that govern the operations of a passive band pass filter are the centre frequency, the attenuation and the bandwidth.

1. The centre frequency f_C is the centre of the selected frequency region:

$$f_C = \sqrt{f_L \times f_H}$$

2. The attenuation is the decibel decrease at the centre frequency (a passive filter cannot amplify).

3. The bandwidth B is the width of the selected frequency region:

$$B = f_H - f_L$$

Another way to describe the bandwidth is to refer to the quality factor Q, a dimensionless value characterising the filter:

$$Q = f_C \div B$$

This can be rewritten as:

$$f_C = Q \times B$$

This relation shows that Q and B are inversely proportional: a high Q corresponds to a low B (a narrow bandwidth) and a low Q to a high B (a broad bandwidth).

Contrary to popular belief, we do perceive phase shifts but our brains use them to localise sounds in space, not to analyse their spectral qualities.

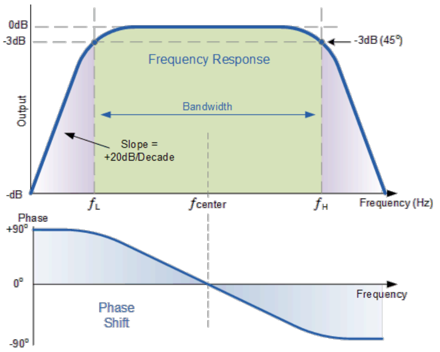

Figure 2: Band pass filter - Level and phase.

7 Summary of our 2013 blog

The aim of this blog was to present a list of what audio engineers consider as definite no-nos in a mastering context. The list rapidly became extremely long and included comments that were irrelevant to the initial aim of the blog. We ended up having to close it. However, we've decided to republish here the most useful safeguards after discarding the pointless, useless and unsuitable comments.

7.1 Managing headaches

To avoid insanity when mastering, resist anything from the following list of temptations.

1. Automatic make-up gain: Make-up gain was described before. The problem here is the 'automatic' part as it has its own attack and release times which will interfere with the compressor's attack and release time. So turn the 'automatic' off.

2. Actually, turn any form of 'automatism' off: don't let machines tell you what to feel.

3. Digital distortion: the only way out is to cut out the clipped waveform and not tell anyone.

4. Brick-wall bulldozers, finalizers, loudness maximizers, terminators and other robotic bullies. They are not your friends.

5. Frequently changing your monitoring level: this is the best way to lose your balance and/or confuse your own references.

6. Meter watching: a mesmerising distraction counterproductive to focusing on the audio.

7. MS techniques: exciting if you enjoy unpredictable rough edges but your clients won't really share your enthusiasm.

8. Reference material: do you really want to sound like everybody else? Believe in yourself.

9. Stray low frequencies: difficult to filter with precision without destroying that beautiful bass tone you spent ages polishing.

10. Compilation album: the most demanding challenge for any mastering facility as it involves mastering different artists, different sounds mixed in different recording studios. Making the album sound consistent is one thing. Making every artist on the compilation happy is another. But, if you succeed in both categories, the artists will come back to you with their own albums to master.