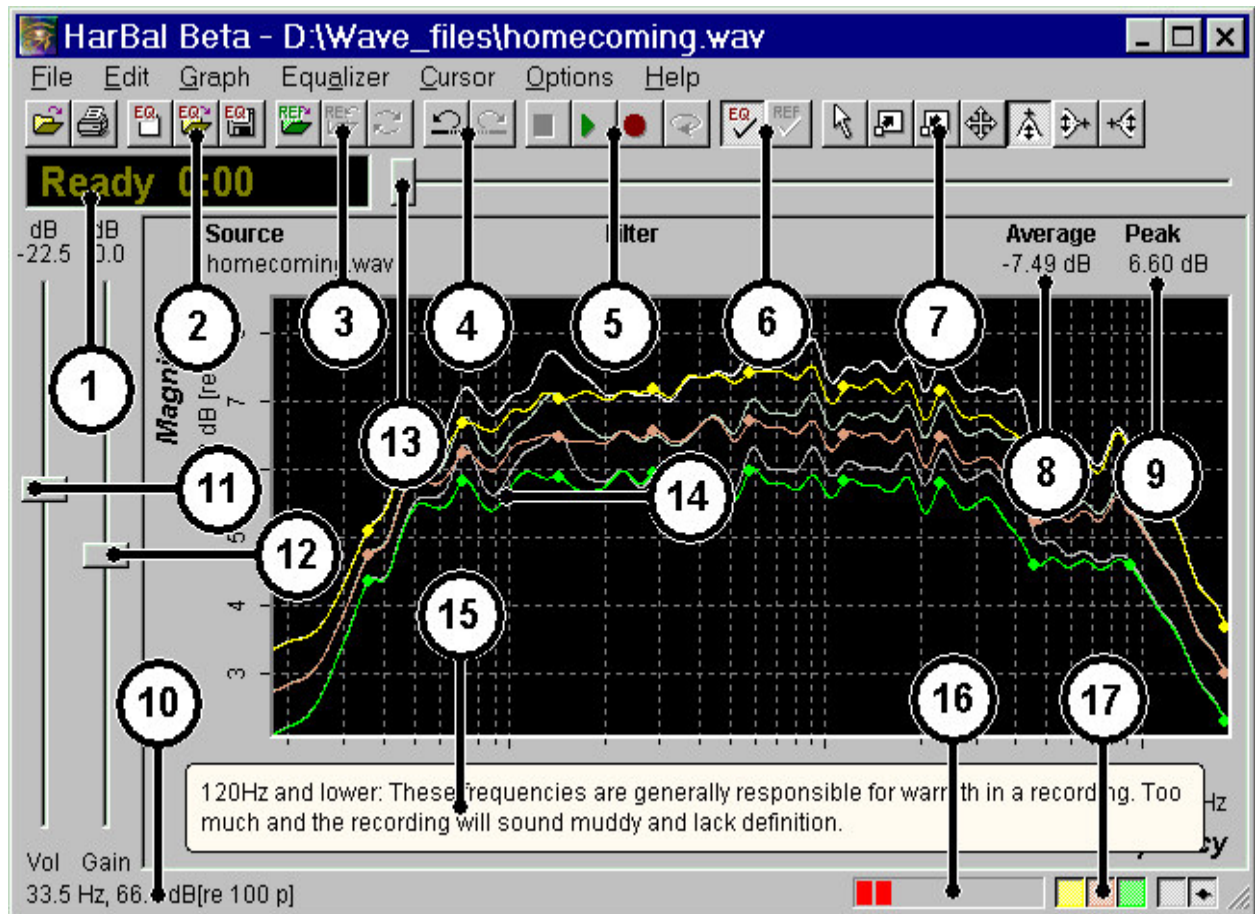


HarBal Tutorials

HarBal is the software that introduces the concept of Harmonic Balancing and the means of achieving it. In this section we introduce the software and the process of Harmonic Balancing as implemented through HarBal. For an introduction to the concept, its origins and its benefits over other mastering processes see [Harmonic Balancing](#).

HarBal Layout

The essential elements of the HarBal environment are shown below:



- 1. Status and time display**
This display shows the live time in seconds of the track being previewed and also display the current state of HarBal (Ready, Playback, Record)
- 2. EQ Filter related toolbar buttons**
These toolbar buttons allow you to open, save and create new EQ filters for the track being harmonically balanced.
- 3. Reference related toolbar buttons**
These toolbar buttons allow you to open, close a reference track or reference track analysis and swap a playable reference with the current track being harmonically balanced.

- 4. Filter design undo/redo toolbar buttons**

These toolbar buttons allow you to undo and redo EQ filter design changes made using the gain cursor.
- 5. Playback/Record control related toolbar buttons**

These toolbar buttons allow you to playback and/or record the current track with or without the EQ filter inline.
- 6. Filter and reference control toolbar buttons**

These toolbar buttons allow you to switch the EQ filter in and out and to switch the reference track in and out of playback.
- 7. Spectrum display control toolbar buttons**

These toolbar buttons let you control the spectrum display and design the EQ filter.
- 8. Average loudness figure of merit indicator**

This indicator displays the A-weighted figure of merit of the average power spectrum for the current track. This can be useful in determining gains to use to normalise the loudness of a number of tracks in a compilation.
- 9. Peak loudness figure of merit indicator**

This indicator displays the A-weighted figure of merit of the peak power spectrum for the current track. This can be useful in determining gains to use to normalise the loudness of a number of tracks in a compilation.
- 10. Status bar message area**

This region of the status bar is used to report context sensitive help on menu items and toolbar buttons as well as general data values from the graphical display.
- 11. Volume control slider**

The volume control slider can be used to change the playback monitoring level. This has no effect on recorded levels.
- 12. Gain control slider**

This gain control slider allows you to apply a gain the playback and/or recording of the current and/or reference track. This is typically used for normalising tracks.
- 13. Cue slider control**

This slider control allows you to cue the current and/or reference track to a particular location.
- 14. Spectrum display**

The spectrum display shows a graphical representation of the overall energy content of the current and/or reference track. The spectrum display shows the average spectrum, the peak spectrum and the geometric mean of the two. We make use of the spectrum display and the gain cursor to design EQ Filter.
- 15. Tips popup**

The tips popup gives spectrum context sensitive clues to designing appropriate equalisation filters.
- 16. Saturation/Overload indicator**

The saturation/overload indicator gives a graphical indication of the amount of saturation/overload occurring during playback. We can use this in conjunction with the gain control to maximise the loudness of the current track without audibly distorting it.
- 17. View control buttons**

The view control buttons can be used to change the visibility of various aspects of the spectrum display.

Where to Start?

HarBal does its magic on sound files only, so a good place to start is with collecting the source material to harmonically balance. For a list of supported file types consult the [Open Command](#) topic.

One notable exception from the list is the mp3 format. We intend to support mp3 format in the future but it is also worthwhile noting that mp3 is generally not a good format to use for source material that has yet to be harmonically balanced and with good reason. The mp3 format uses lossy compression built around perceptual encoding of the recorded sound. What that means is that if the model believes that you cannot hear a particular part of the spectrum in a recording (because it is masked by an adjacent sound) then it will drop it altogether. This is fine if the recording is already harmonically balanced but if you have a poorly balanced recording and harmonically balance it, normally, much of what is masked becomes readily audible and hence, the improvement in clarity obtained using harmonic balancing. But if you have encoded the source using mp3 before equalisation the mp3 encoder may drop a significant proportion of content that can never be recovered, with or without harmonic balancing. For that reason we only recommend using mp3 post harmonic balancing.

Ideally the source material should be arranged such that each track is in a different file. You could, if you wish, perform harmonic balancing on a single file corresponding to an entire album but the outcomes that you will end up with is likely to be less than optimal. This is due to each track in a compilation having its own specific characteristics which become lost in the other tracks if lumped together.

Given a set of tracks that you wish to master we are now ready to start. The first thing to do is to open one of your tracks using the [Open Command](#). The first time you do this on any given track HarBal will proceed to analyse the spectrum content. This may take some time depending upon the length of the track and the speed of your computer but once done the result is saved to disk for re-use. If you then re-open the same track, provided it has not been modified in any way, HarBal will open it immediately by reading the analysis file of the track and using it to initialise the spectrum display. Alternatively, if you have a compilation of many tracks that you are going to work on you can pre-analyze those tracks in the background while you work on another file in HarBal by using the [File | Batch Analysis menu command](#).


For the sake of this tutorial I shall be demonstrating Harmonic Balancing of a CD compilation using the *CD Human Frailty by Hunters & Collectors CD 53205*. Note that we do not condone illegal pirating of copyright protected recorded material in any way. We have a legitimate copy of this CD and are re-mastering it for demonstration purposes only.

Beginning a Session

Using the [Open Command](#) I will open the track *Everythings on fire* first (I like the song but have trouble with the mix). This is the first time I have opened this track so the analysis progress dialog pops up as shown below.



When the analysis runs to completion the spectrum display shows the following result.



The spectrum display has three traces; the lower trace is the average energy content of the track at different frequencies, the top trace is the peak energy content of the track and the middle trace is the mean of the two.


Looking closely at the spectrum I would argue that it is generally a good recording but there is a significant hole in the spectrum between 150Hz and 500Hz which, to my ears, makes the recording sound thin and without body. When you consider the vintage of the recording and the career stage the band was in at the time, I would guess that the large and sharp peak at 190Hz was probably due to a studio resonance and that it was EQ'd out during the recording process. But clearly the characteristics of the EQ filter did not match that of the resonance very well, and hence, much of the content of the track was inadvertently attenuated for the sake of making the resonance tolerable. This in essence, is the problem with traditional mastering : for difficult recordings it is extremely difficult to do justice to the mastering process because we cannot clearly see the problem.

Improving the Mix


Having identified the problem in the track how do we go about correcting it? This is where the process of harmonic balancing is like no other mastering process. All we need to do is design a filter to match the recording characteristics, bringing out what is now masked by that resonance. And here is how simple it really is.

Start off by selecting the [Gain Cursor](#) from the toolbar or from the corresponding menu item. We shall now use it to design the EQ filter by restructuring the recordings spectrum. A good place to start is with the 190Hz peak which is far too pronounced. To reduce it we position the gain cursor over the peak and hold the left mouse button down. Then by dragging the mouse down we, in effect, dial up a gain cut in a parametric EQ. Likewise dragging the mouse up dials up a gain boost and dragging the mouse left and right alters the Q. If you want the Q to remain fixed while adjusting the gain press and hold down the SHIFT key. The centre frequency of the EQ is the frequency at which you first depressed the mouse button. While you are doing this a preview of the change tracks the mouse movement to give a clear indication of the change you are about to make. If you want to cancel the


edit press the ESCAPE key prior to releasing the mouse button. You can also switch the adjustment from one trace to another (say from the mean to the peak spectrum) by pressing the tab key while in drag mode. The first change I made is as shown.



Now that we have controlled the peak, we need to deal with the hole in the spectrum. This is done in a similar manner by clicking and dragging the spectrum at the middle of the hole and adjusting the gain and Q to best match. Note that we can apply as many of these changes as we like and not degrade the performance of the EQ filter because we are only using these *parametric EQ changes* to design the magnitude frequency response of the filter, which is only realized once when needed (during playback or recording). The realised filter has either a linear phase or minimum phase response depending upon the configuration in the [Equalizer Options dialog](#). If we consider the alternative approach in traditional methods, the more changes we make the more generational noise is added to each equalisation operation. To remove the hole in the spectrum I'll choose a centre frequency at around 235Hz because it corresponds to a minor peak in the hole. I'll talk more about why later.




Now the spectrum is looking a bit more uniform but we still have a significant dip at around 150Hz. Again using the gain cursor to match the hole in spectrum we obtain,



The spectrum is now looking much more uniform than before and now is a good time to validate our changes by listening to the result of our EQ filter. To do so you will obviously need a sound card and a decent pair of speakers to monitor through. Also, should your sound card have some sort of loudness built into the speaker output it would be advisable to disable it if you can, and if you can't, wire the monitors through the line out and an external amplifier. By default HarBal uses the default Windows multimedia output device driver for playback. However, you can specify the driver of choice from the [Multimedia](#), [Direct X](#) or [ASIO](#) interfaces in the [Driver Options dialog](#). To open this dialog select the [Options | Driver menu command](#).

Assuming that you are setup to monitor the playback, start by pressing the [play](#) button or the [play](#) menu item. Alternatively you can press the SPACEBAR to toggle playback on and off. You should hear the track being played and see the time display and the cue control update during playback. If the [EQ In button or menu item](#) is checked then you will be listening to the EQ'd version, otherwise you will be listening to the original. You may toggle between the two modes during playback to A/B test your equalisation filter. When I do so, my ears tell me that although the recording has much more body than before, the bass guitar sounds a little too prominent, generally muddying the vocals. Looking at the spectrum it is reasonable to assume that the broad peaking around 100Hz is responsible so we use the gain cursor again to make another change. We get,




Note that I can make the change during playback and I immediately hear the change to the equalisation. On listening to this version there is a significant improvement in clarity of the EQ'd version over the original but I believe it can still be improved upon. On careful listening I can pick up on significant resonances in the vocal of the track. At a guess I would say the the vocal has been recorded in an ambient studio environment giving the recording a sense of space that is pleasing. On the other hand, there are some resonances in the vocal that impact upon the intelligibility of the

vocalist. This is the point at which experience in using HarBal is important to make effective use of the software.

When presented with audible resonances it is worth noting that more often than not they tend to stand out in the peak spectrum trace. Bearing this in mind I would suggest that the problem is confined to three main areas; the sharp peak at 240Hz, and the broader peaks at 660Hz and 1.3kHz. You might think that there are many more possible problem areas like the large peaks around 9kHz and the other general bumps and dips in the spectrum but beware of excessive massaging of the spectrum. If you try and completely flatten the spectrum you will almost certainly extract the life from your recording.


The peaks around 9kHz are perfectly legitimate and correspond to the natural resonances of the cymbals and hi-hat in the recording itself. If you were to flatten this out it is highly likely that you will destroy the quality of the reproduction of these instruments. This is generally true of any instrument as each will have parts of the spectrum that will be strong and others weak. You must bare in mind that the spectrum display is smoothed to 1/6th octave resolution, principally to make adjustment easier, but in doing so if a region shows a strong dip it could be because there is little sound in that region. If you were to boost it too much you may well find that all you are doing is amplifying background noise. Similarly, anyone with experience can tell you that if you record in an ambient environment this will impose peaks and dips in the sound radiating from instruments. These peaks and dips actually define the ambient sound of the room, so in my case, If I were to remove all these peaks I would be attempting to remove the room from the equation. That's certainly not what is intended here. What we really want to do is to improve the acoustics of the room electronically by reducing the effect of dominant peaks. So if there is one general rule to live by using HarBal it is that you should aim to rid yourself of dominant peaks but be wary of eliminating certain peaks such as those generated by hi-hats, cymbals and brass.

On instigating these changes we obtain,



The recording is now sounding clear but is a little strong around 3kHz and lacks the mid range fullness that goes so well with a live sound. As a final change I will reduce the amplitude around 3kHz with a fairly low Q gain change and boost the mid-range with a low Q gain change at around 700Hz. The final result is phenomenal when compared with the original. It sounds full, spacey and perfectly clear with absolutely no sense of fatigue.

Now before showing you the final result lets save the filter using the [Filter Save](#) menu command or [save filter](#) button and open the origin file as a [reference](#) so that we can see how our new spectrum compares to the original. Opening the *Everything's on fire* track without a filter as a reference we obtain the following spectrum display.




The differences are clearly obvious but, as you can also clearly see, the essential defining character of the recording has been preserved. This is exactly what is required of a mastering system. To preserve the intentions of the artist and producer whilst making the track as audibly pleasing as possible.

Tone and Loudness Matching


You might think that just being able to harmonically balance a recording is enough but there is one issue that is difficult to master without the appropriate tools for the job. Harmonic balancing is fine if all you have is a single track to master but what if you have an entire album? You might find that if you simply harmonically balance each individually, that although in isolation each sound great, when put together in a compilation the tonality and loudness varies considerably from track to track. For the listener this can be disruptive to the overall experience because as soon as your mind adjusts to the tonal frame of one track it suddenly changes in the next, setting your mind off balance.

In HarBal we can manage this issue without much difficulty using referencing. In the example given, now that we have harmonically balanced *Everythings on fire*, we can move on to balancing the remaining tracks on the Album. So that we maintain the correct tonality we can use the track *Everythings on fire* and the filter we design for it as our reference using the Open Reference command. As an illustration let's pick another track off the album to re-master. This time we'll use *Dog*. The spectrum for this track overlaid on our reference is shown below.



This track shows similar problems to the previous one but in this case we have a reference spectrum to guide us in our adjustments. Now provided that both songs follow a similar genre then matching the spectra (in a loose sense) will give good tonal matching. If the genre is different then you will have to make some judgements for yourself as to how to best adjust the spectrum. This will typically be an issue at the low and high frequency extremes of the spectrum. As an example, if this second track didn't have cymbals and hi-hat in it you could expect that it would have significantly less energy around 10kHz than the reference and in such a case it would be wrong to boost that region to match the reference.

In this case, because the tracks are similar in style the adjustment process is simple and we end up with the following result.



But that is only half the story. How loud it sounds is determined not only by the shape of the spectrum but how high the recording level was during mix down. Ideally we want to make all tracks as loud as possible and as loud as one another without audible distortion. This is commonly referred to as normalisation, although the most common normalisation technique is to merely increase the output level to the level that is highest and without saturation/overload. However, with this simplistic approach you end up with tracks whose mastered level is determined by the largest peak level in the track rather than how loud it actually sounds. The net effect is that it is unlikely that you will have a consistent perceived loudness throughout the compilation.

With HarBal we take a different approach. HarBal provides an estimate of the loudness figure of merit of a given track based on both the A weighted peak and average spectrums. If you choose to match the average power loudness figure of merit for the current track and the reference then you will have a reasonably accurate normalisation. You can use the [Equalizer | Match Loudness command](#) to do this gain matching in one go. But as always, fully the automated approach will not work flawlessly all the time, particularly if you have a compilation with many styles or genres. In these cases you can perform [A/B loudness comparisons](#) by listening to the track and toggling between it and the reference while making gain adjustments to get the matching just right. Note that when the reference is playing the gain control adjusts the reference gain and when the current track is playing the gain control adjusts the current track gain. This still leaves the question of how much gain can be safely applied without compromising the quality of the recording through saturation. This issue is handled through the Saturation/Overload indicator.

The saturation/overload indicator is not at all like a VU level meter. In this application the VU level is essentially irrelevant. What we really want to know is when and by how much the recording is saturating. In a traditional sense you would judge this by how much time the VU meter is spending in the red. Rather than relying on human judgement we have taken the liberty to create a meter that actually measures how long the signal is spending in a saturation state.

The saturation/overload indicator is a log scale meter. If the meter is 100% red then it means that the track is in saturation 100% of the time. The range of saturation that it displays is shown in the following table.

Meter reading % full scale	% time in saturation
0%	< 0.1%
10%	0.2%
20%	0.4%
30%	0.8%
40%	1.6%
50%	3.2%
60%	6.4%
70%	12.6%
80%	25.1%
90%	50.1%
100%	100%

Now whatever level you should choose to be permissible saturation is a matter of opinion although, in my opinion, it is difficult to hear the distortion for readings less than 30% full scale. Generally speaking, I usually adjust the output level so that the saturation indicator lights up only occasionally. If you work by this rule then you should have adequate headroom to accommodate the different spectral content of the tracks that you are trying to normalise. In the case of our example I am able to boost the level by 5.5dB on both the reference and current tracks without significant saturation. Also, in performing A/B testing during playback I find that both tracks have comparable perception loudness. Alternatively, if you have access to a limiter the best approach would be to adjust the gain so that you never have saturation, render to 24 bit files and then use your limiter to bring the final tracks up to the desired level.

For a complete compilation you will have to optimise the level for all tracks. The easiest way that I have found for doing this is by setting the reference track gain to the desired level and then leaving it “*as is*” while using the same reference for all other tracks on the compilation. Then as you harmonically balance each track in turn you can perform the normalisation at the same time.

Conclusions

There are a couple of other points worth noting that haven't been touched upon. Firstly, although in this case the reference and session tracks had similar recorded levels, in many instances you will find that they do not. In such cases the reference and current spectra usually don't lie one on-top of the other. This may seem to present a problem in ppectrum matching but it is easily overcome. All you need to do to fix this issue is to use the gain cursor to amplify or attenuate (whichever is appropriate) the current spectrum using a very low Q filter. So low in fact, that it is just like a simple gain control. You can then overlay the current spectrum with the reference spectrum, thus making the job of tonal matching that much easier. You might think that this will in turn amplify the track so that it might now saturate but this is not the case. HarBal determines the overall gain of the filter such that the A-weighted average spectrum after filtering has the same power as the unfiltered version. In effect this means that the vertical position of the spectrum is ignored. The only thing HarBal doesn't ignore is its shape.

Finally, when doing tonal matching you quite often find that you can come up with a more pleasing balance in the current track than that of the reference and you would like to re-adjust the reference to match the current (rather than the other way round). Normally this would involved the annoying task

of closing the current reference and the current track and re-opening the new one so that you can perform the adjustments that you would like to make to your reference. To save time HarBal has a single button operation that will swap the reference with the current track, including the track cue position and gain settings for each. To do so use the [Reference Swap](#) command. After making your adjustments you can the restore your reference at a click of a button.

That concludes our tutorial on using HarBal. In time and with experience you will discover the best way to use it. Should you have any ideas or comments concerning HarBal we would certainly be interested in hearing them. You can email us on support@har-bal.com.