



Har-Bal Harmonic Balancer: An Introduction to Version 3

Har-Bal™ new thinking, new directions...

Limitations and Motivations

The first generation of Har-Bal introduced high resolution filtering coupled with spectrum analysis and a novel *"loudness compensated"* filter designing user interface to improve the capabilities of mastering engineers in tackling troublesome equalisation issues. As good as the fundamental design is it has shortcomings when dealing with certain classes of mastering problems.

Non-stationary Tonal Balance

Prior to the era of hyper-compression and hyper-limiting many recordings that presented tonal imbalance issues could satisfactorily be treated with static equalisation. That is, one EQ setting applied to an entire track. In the pursuit of absolute loudness new factors are introduced as a result of an interaction between compression/limiting and the track dynamics : the forte passages within tracks often have accentuation of the upper mid range as a side effect of dynamics processing leading to a harshness of tone in those passages. If treated with static equalisation the harshness can be counted but only at the expense of the tonality of the quieter passages that become coloured by the equalisation. With tracks of this nature what is required is equalisation that is tailored to specific parts of the track.

The contemporary approach to *"dynamic"* equalisation issues is the application of multi-band compression to a track. This is a less than ideal approach for a number of reasons. Firstly, the application of equalisation through multi-band compression is a rather indirect approach to equalisation that is difficult to accurately target to a specific track. Secondly, attacking what is essentially an equalisation issue with compression necessarily alters the dynamics of a track because the *"effective"* equalisation is being constantly adjusted. Thirdly, because of the constant adjustment, processing artifacts from multi-band compression can occur.

The Har-Bal approach to such equalisation issues is to instead split a track up into its structural components (intro, verse, chorus, verse, chorus, bridge etc.) and have a custom filter for each. Within each segment the filter characteristic is fixed so the dynamics are essentially un-altered, only the relative dynamics between sections are altered. Because the filter characteristic is fixed there is no dynamic adjustment side effect except at the transition from one section to another. However, if the transition point is chosen carefully, significant change in equalisation without audible side effect is possible because the timbre of the track naturally changes from, say the verse to the chorus, and through our listening experience we expect to hear a corresponding change in timbre. The net effect is that the potential side effect of our change is masked in the track structure. With careful filter design artifact free dynamic equalisation is possible and the new user interface within Har-Bal 3 provides many aids to getting the design right.

Channel Specific Issues and Problems in Stereo Image

Sometimes a pre-master track has been inappropriately mixed and has imaging problems. Perhaps the track needs to be prepared for vinyl mastering and requires monophonic bass or perhaps some parts in the track are panned too widely or too narrowly and for whatever reason re-mixing is not a possibility. Perhaps you would like to demonstrate to the mix engineer the imaging you believe the track requires so he/she can have something to reference against to achieve that end.

Har-Bal 3 addresses these issues by providing separate mid and side filtering capabilities coupled with full mid and side spectrum analysis. From this analysis the spread of the imaging can be readily seen from the relative strengths of the mid and side spectrums and readily altered by independent filtering applied only to the side component. For completeness, Har-Bal also provides independent left and right filtering to attack channel specific issues should the need arise.

Defining a Coherent Space

Sometimes one is presented with a recording that lacks a sense of immediacy, of not being present in the performance because of a lack of recorded ambience. Sometimes ambience is present in some parts but completely absent in others (lead vocal for example), leading to a sense of disconnection between parts.

To bring some cohesion to these types of issues Har-Bal 3 introduces Haas zone cross-coupled ambience processing. This new Har-Bal Air implements an idealised ambience tail with a reverberation time of less than 50 milliseconds (Haas zone) and less ambience at higher frequencies to give a natural warmth to the ambience.

Being Haas zone ambience, there is little coloration to the natural timbre of the recorded music except when used in the extreme. The new Har-Bal air simply provides a space, particularly when the space is absent in the recording. In mono recordings Har-Bal air can also be used to simulate stereo effectively and naturally.

Controlling the Dynamic Range

Pre-masters of popular music usually have more dynamic range than is ideal for the genre. To obtain a commercially loud master some compression combined with limiting is usually required to achieve suitably high levels.

For such pre-masters it was not generally possible to obtain suitably loud masters in Har-Bal alone because it had not means of reducing dynamic range except through limiting, which if used exclusively to obtain loudness often has delaterious side effects.

To provide for such cases, Har-Bal now includes a means of altering the dynamic range through mapping the average track level to a different one. This approach is more closely comparable to volume riding than traditional compression and works by mapping the volume envelope obtained in analysis to a different one specified through a novel interface. This approach leads to a compressed sound that is quite unlike traditional compression.

The Inner Workings of Har-Bal 3

Track Analysis

To provide for the ability to perform dynamic filter and the mid and side processing a new approach to analysis is required.

In traditional Har-Bal the analysis is an average for the entire track. In the latest generation we need support for track splitting with arbitrary split points that we cannot know a-priori. To cope with such a requirement in a timely fashion we need to store a partial analysis of instantaneous spectrum results over the entire track time so that we can quickly average them over a track segment to give the segment analysis result.

This is achieved by taking overlapping time windows at nominally 50 millisecond intervals and obtaining a spectrum estimate through Fourier transformation, binning the result into 1/12 octave bins and storing the result to disk in an analysis file (file extension .hba).

The spectrum estimate for a stereo track is actually four fold because we need analyses for mid (left + right), side (left - right), left and right components. We also record the average and peak power levels for each of those components per time step. This strategy is summarised in Figure 1 below.

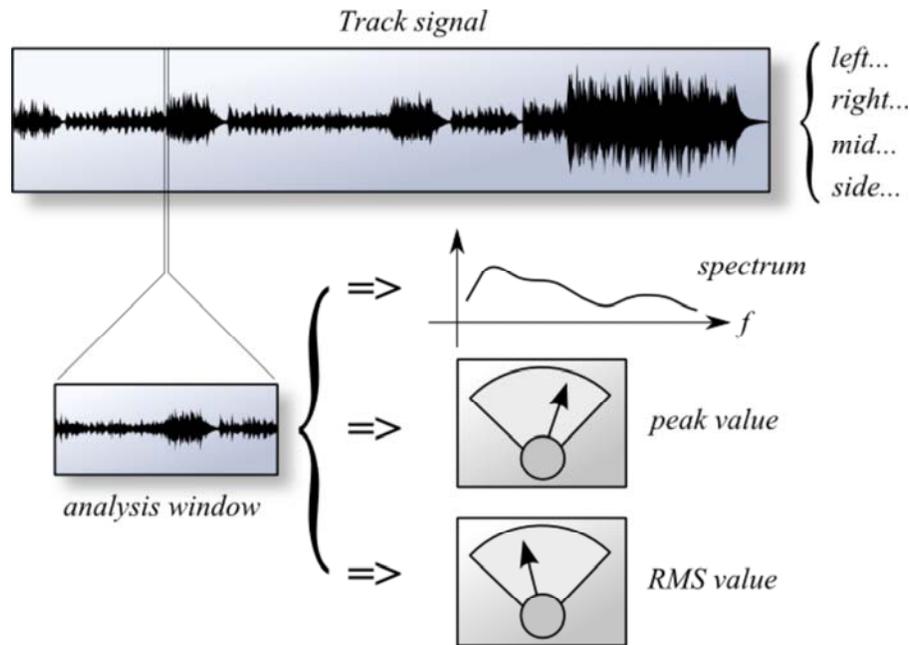


Figure 1: The comprehensive analysis scheme implemented within Har-Bal

This degree of analysis has a number of side effects. Firstly, an analysis file must contain a lot of information because we record four spectrum estimates at each 50 millisecond interval for the entire track. Typically for a stereo PCM encoded wave file the corresponding analysis file (.hba) is larger than the original source. This is unavoidable if you want to have access to the power that this degree of analysis gives.

Secondly, as a side effect of this detail, a real-time spectrum can be displayed in the Har-Bal analysis graph. Having access to this real-time information is invaluable in determining which parts of the spectrum are responsible for uncomfortable sounds in a given recording.

Thirdly, the average power component for each 50 millisecond interval can be combined to form a time series that accurately represents the track volume envelope, which in turn can be used as a basis for implementing dynamic range processing.

Filter Structure

To cater for the requirements of independent filtering of left, right, mid and side components and to provide Haas zone ambiance, a new filter structure is required. Figure 2 illustrates the new structure of the Har-Bal filter block. Keep in mind that each component of this filter can have many different response functions owing to track splitting (for each segment in a Har-Bal filter there is one of these filter blocks).

Copyright © 2011
Har-Bal™. All rights reserved.

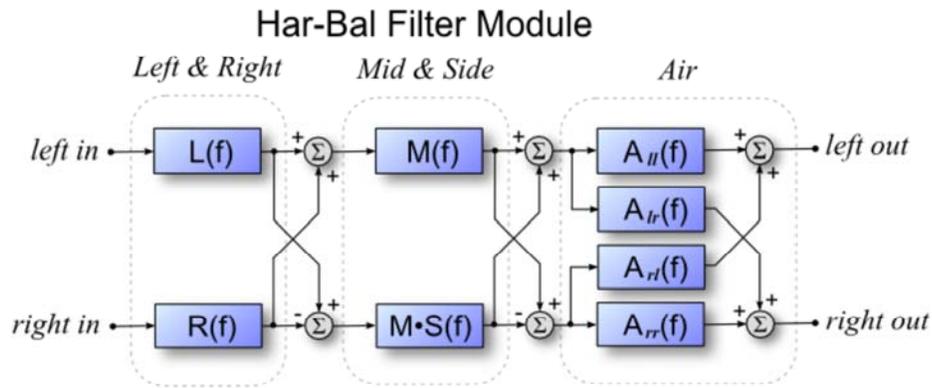


Figure 2: The basic Har-Bal filter module structure

Left and right filters come first, then after decoding to mid and side components, the mid and side filters follow. Note however, that the mid filter component is included in both the mid and side filters. This is deliberate by design. The reason being that the mid spectrum view is the view you will most often use to design equalisation schemes for tracks but when you apply equalisation in this view you do not want to alter the stereo image. For that to be guaranteed we must apply the same equalisation to the side component, hence its inclusion. However, the side component can be adjusted independently to allow us to alter the stereo image when we see fit to do so.

Following decoding from mid-side representation back into stereo we follow with Haas zone filtering (Har-Bal Air). This is composed of minimum phase cross coupled filters that introduce a log normal random deviate to an all pass (flat) response. The greater the deviate the more ambience. As the responses are nominally flat, they add little colouration. The synthesis process is illustrated in Figure 3.

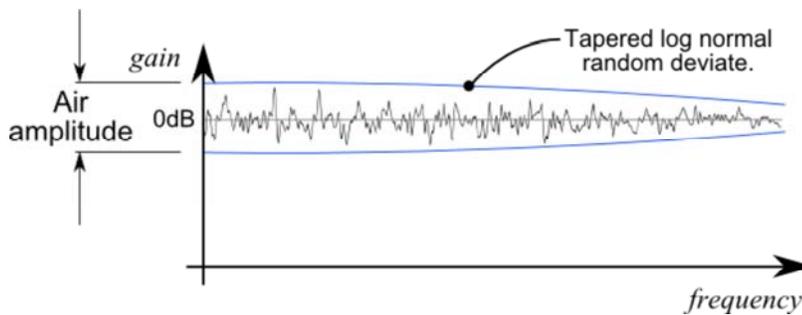


Figure 3: Frequency Domain Synthesis of Haas Zone Ambience

A flat frequency response is deviated by a log normal random deviate with characteristics suitable to confine the ambience tail within the Haas zone. The magnitude of the deviate is controlled by the Air setting. To create a natural and pleasing sounding ambience the magnitude of the deviate is tapered off at higher frequencies. This has the effect of making the reverberation time of low frequencies longer than that of high frequencies, a characteristic that is appropriate for music reproduction. This frequency domain representation is translated to the time domain by realising a minimum phase impulse response that matches the given frequency response. This results in an impulse response like that in Figure 4.

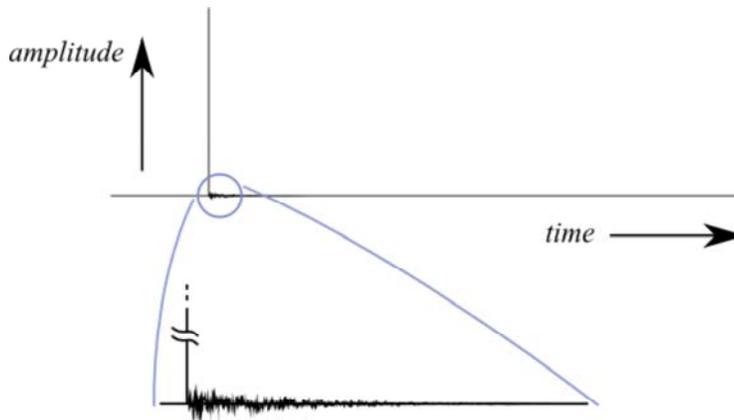


Figure 4: Resulting Time Domain Impulse response with Ambience Tail

This process is used to realise all four responses required to synthesise the complete Air filter network, except that the cross couple responses of $A_{lr}(f)$ and $A_{rl}(f)$ have no unit impulse component, just the ambiance tail. Each filter uses a unique uncorrelated random deviate in the synthesis process and it is this uniqueness of response which results in a spatialising effect.

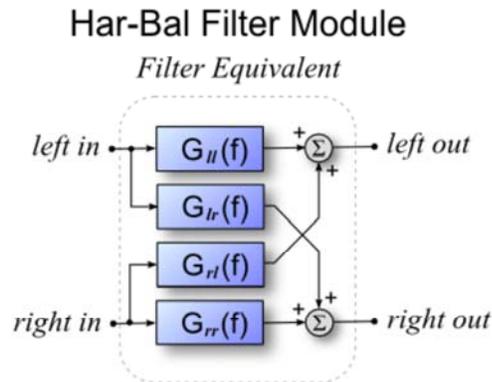


Figure 5: The basic filter module structure reduced to a structurally simpler equivalent to reduce signal processing demands

The Har-Bal filter module (as illustrated in Figure 2) contains eight filters. If implemented independently, more processing power will be used than is actually necessary. As an efficiency measure, the entire filter structure is replaced by the equivalent filter shown in Figure 5 and only requires four filters per module. This compares favourably with the first Har-Bal implementation which only has two filters (left and right, but with a common response).

What is Mid – Side Encoding?

The Har-Bal filter module makes use of Mid – Side encoding and decoding to provide extra control in mastering equalisation. For those unfamiliar with it, Mid – Side encoding is defined by the following equations:

$$\begin{aligned} \text{Mid} &= \text{Left} + \text{Right} \\ \text{Side} &= \text{Left} - \text{Right} \end{aligned}$$

You can see that with a mono source $\text{Side} = 0$ as $\text{Left} = \text{Right}$. So the level of the *Side* component relative the *Mid* component gives an indication of the stereo spread of your mix.

To reconstruct *Left* and *Right* from *Mid* and *Side* use the following equations:

$$\begin{aligned} \text{Left} &= (\text{Mid} + \text{Side}) / 2 \\ \text{Right} &= (\text{Mid} - \text{Side}) / 2 \end{aligned}$$

It can be noted that if the *Side* component is reduced in amplitude relative to the *Mid* component we will narrow the stereo image. This can be inferred by the fact that in the limit as the amplitude of the *Side* component is reduced to zero the mix tends to monophony. You can use this fact to alter the stereo image in Har-Bal by applying filtering to the side channel in a given Har-Bal filter module.

Segmented Filtering

To obtain artifact free processing a track should, when necessary, be split into segments based on musical piece structure. Figure 6 shows a typical segmentation and how many filter responses are involved.

In this case the track is split into eight segments and there are 9 filters : one for each segment plus one overarching filter for the entire track. Har-Bal provides this additional filter to make it easier to control the overall colour of the track. Without it, should you wish to change the overall colour you would have to clumsily apply the same filter response to each of the eight segments individually.

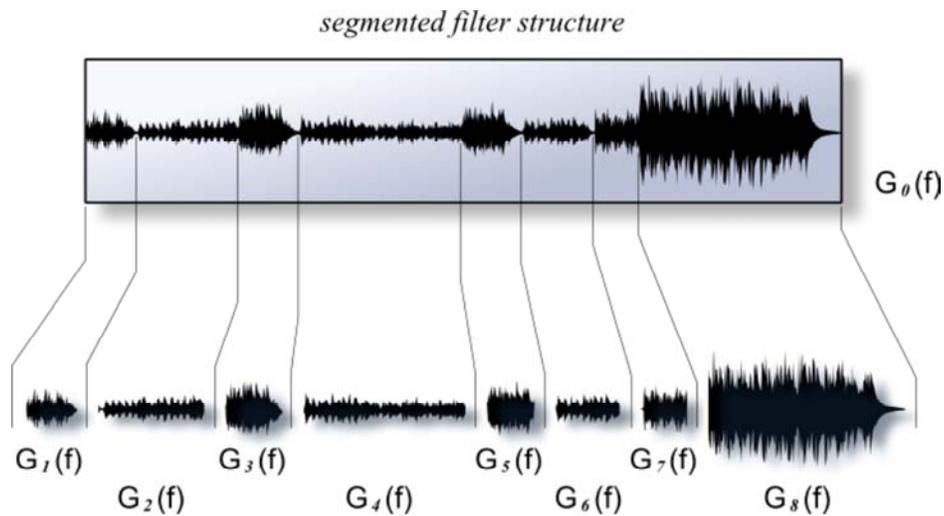


Figure 6: Building a hierarchical filter tailored to the track musical structure

Creating, deleting, modifying and navigating the segments and corresponding filters is handled through the time-line control in Har-Bal. Filters are selectable through the selection bars on the timeline control as indicated in Figure 7. This schematic representation naturally represents the actual filter hierarchy whilst also allowing easy navigation of the hierarchy.

The two traces in the time line show the average and peak power time series obtained from the track analysis and the extent of the time shown corresponds to the selected filter, which in this case corresponds to the outermost overarching filter.

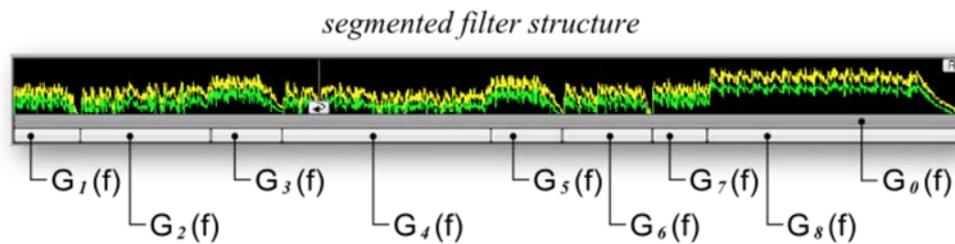


Figure 7: Representation of filter structure through the Har-Bal time line

For each selected filter Har-Bal displays the peak and average spectrum traces with and without filtering from the top of the hierarchy downward, as in Figure 8. In this case, as the outermost (top) filter is currently selected, the average spectrum shows the estimated spectrum that results from the application of all the filters applied to their corresponding segments plus the effect of the overarching filter.

If an individual segment were selected then the spectrum estimate would show only the average spectrum for the time range of that segment and the effect of that filter on that spectrum only. The reasoning here is that if you are adjusting a segment filter you only want to know the effect that the segment filter has on the segment average and not the effect of any filtering applied above the level of the segment. Hence, in this case the effect of the outermost filter is not included.



Figure 8: Har-Bal spectrum view

With the frequency response view selected we see the filter response of the selected filter, which in this case is the outermost one (see Figure 9). Whichever is selected, Har-Bal only shows the response of the selected filter and not contributions from other levels of the hierarchy. The logic here being that for frequency response view we wish to see the filter response contribution for the selection.

The frequency response view shows the change in level applied to the components of the spectrum at a particular frequency. It is analogous to the image cast by the sliders on a graphic equaliser and in this case shows a high pass response with a cut-off at 30Hz.

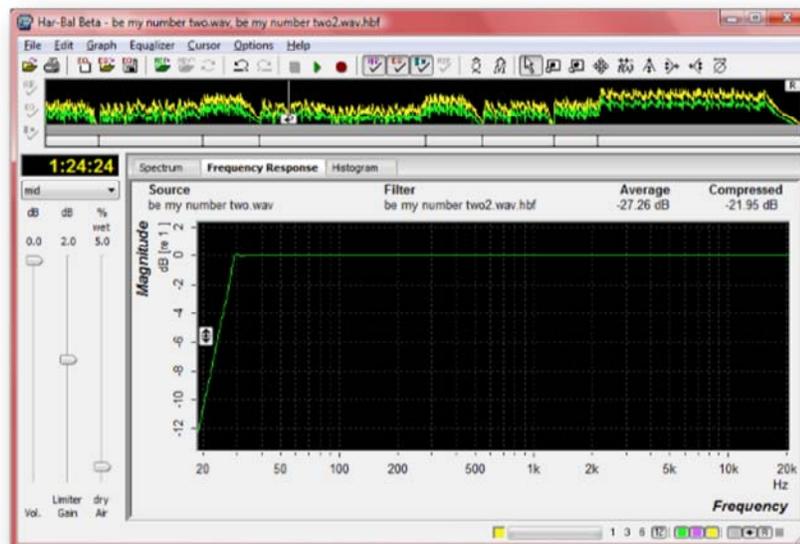


Figure 9: Har-Bal frequency response view

The final view Har-Bal provides is the histogram view, which corresponds to the histogram of timeline levels, binned in 1 dB wide bins, both peak and average, for the selected segment. This view is also used to design a dynamics processing scheme for the selected segment using dynamics node editing. Figure 10 shows the histogram view for the entire track and the effect of a simple compression scheme on the histogram.

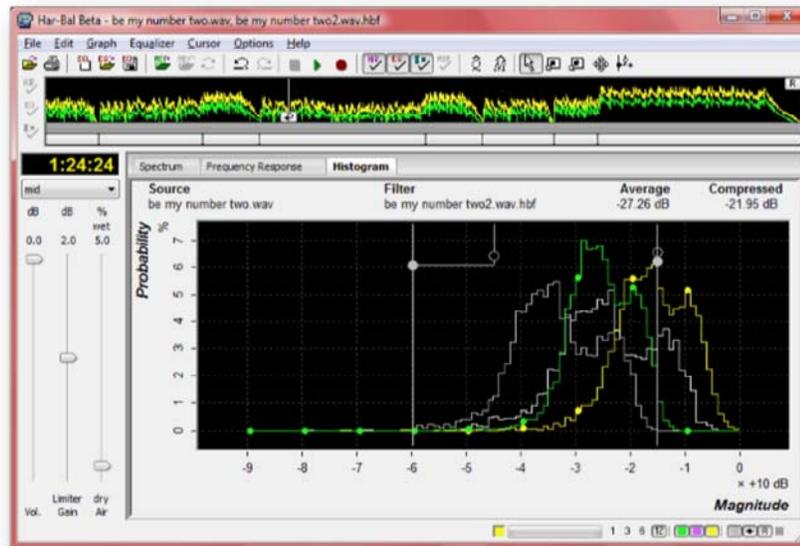


Figure 10: Har-Bal histogram response view

Histograms are used to quantify track dynamics because they provide a complete, convenient and quantitative snapshot of the track characteristics, unlike metering which can only ever show you a very localised snapshot. The histogram in Har-Bal is a graph representing the likelihood (probability) of the x- value (horizontal axis) occurring. In our case the horizontal axis is level in dB and it is binned in 1dB wide bins (hence the city line appearance of the graph). If you take any one of those tall rectangles, look at the dB value below it and the percent value corresponding to it's height, you can interpret it as follows. Lets say the column I'm looking at is centred on -5 dB and it's height is 7%. That is telling me that for 7% of the time the level is between -5.5dB and -4.5dB (note that the bin is 1dB wide and centred on -5dB). That could be either the peak level or the average level depending on whether you are looking at the green or the yellow trace.

If you happen to move the gain slider you'll note that the histogram marches up and down the horizontal axis by the amount of gain you apply. Now that you know what the histogram is you should see why that should be: because a level that was in the range from -5.5dB to -4.5dB after a 2 dB gain will now be in the range of $(-5.5 + 2)$ dB to $(-4.5 + 2)$ dB. However, you might be wondering why when you increase the gain the 0dB bin gets taller and taller. Well, if you think about it, the maximum level we can produce is 0dB. If something was in a -3dB bin and we apply a 6dB gain it can't go into the +3dB bin because there isn't one. What actually happens is the limiter cuts in and forces it into the 0dB bin. Thus, the more you push the gain higher the more the topmost bin in the histogram fills up as the values from other bins are forced into it.

This is a very useful thing in mastering because you can immediately see how over-zealous your limiting is. As an example, take any current generation master of popular music and look at the histogram of the peak level and you will almost certainly see a massive 0dB bin and little left in the other bins. That is what over-limiting does to the dynamics of your track. It squashes them into oblivion.

Dynamics Processing

The dynamics processing in Har-Bal is based entirely on the average level envelope that you see in the time-line. You create a transfer function to map the input average level to a new output average level and Har-Bal uses it to figure out what gain to apply dynamically prior to the limiter at any point in time. Unlike conventional dynamics processing, it doesn't use attack - decay envelope extraction. The envelope it already has from the analysis. As such the gain change that tracks the music is in sync and symmetrical about the content. By that I mean with a normal compressor you would typically have a fast attack time and slow release which means that the gain quickly drops at the start of a transient and slowly rises at the tail after the transient.

Now you'll probably think that makes no sense and will sound awful. However it doesn't sound awful, and here is the reason why. In traditional compression we principally have fast attack and slow release because a compressor cannot predict the future so it doesn't know when to pull down the level and if it does it slowly a very big transient will push through before it has a chance to bring it down. We have a slow release time because fast gain change will lead to a lot of distortion. Because Har-Bal knows the entire history of the track we don't need fast attack because we know when to pull the level down. The net result is that the dynamics processing in Har-Bal works more like the volume riding technique of recording engineers of days gone past where they would adjust down the level while the music is being recorded because they knew when the loud parts were coming. This point is clearly illustrated in Figure 11, which shows the difference between envelope detection used traditionally and in Har-Bal.

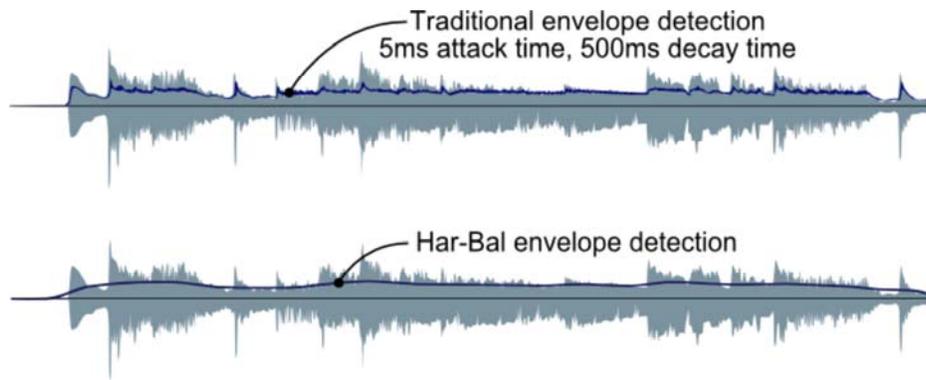


Figure 11: A comparison of envelope detection in traditional and Har-Bal dynamics processing

By using a slow moving envelope for dynamics processing we ensure that the short term transient material remains largely unaltered and we simply change the relativity between loud and soft passages of a track. This also ensures that very little distortion is introduced by dynamics processing. These two factors are responsible for the clean sound of Har-Bal's dynamics processing.

How do you actually use it. Well, as hinted at before you design an input-output transfer function using the dynamics node editing tool in the histogram view. Why the histogram view? Because the histogram view shows you a summary of the dynamics content and to stick with a theme in Har-Bal, it will show you the effect your dynamics processing has on that histogram.

The process works like this. A transfer characteristic is made up of nodes. You create a node by clicking the dynamics node tool anywhere on the histogram graph. The node displays as two circles connected by lines, one filled and one not. Think of the unfilled one as an "O" for output. It represents the level you are mapping the input to. The filled circle represents the input level. So a typical compression scheme would work like this. You look at the average level histogram (green trace) and look at the highest level it has (ie. where the hill comes down to the flat on the right hand side). At that point you put a node with input and output equal. Why? because we want the top of the histogram to stay put and just squash up the the bottom (quieter) parts. Its saying map the input level to the output level of, for arguments sake lets say it is -15dB .

Now look at the other side of the hill. Where the left hand slope comes down to the flat we want to map that level to a higher level (compressing it) so we press down the left mouse button (that sets the input level) then drag it up to the level we want to map it to and release the mouse button (that is the output level). For example, it might be an input of -60dB and an output of -45dB . You'll note that after releasing the mouse Har-Bal gives you an updated histogram showing the effect of the scheme on it and also the updated average track level figure of merit. You'll also note that if you happen to render those changes and re-analyse the result the prediction is pretty accurate. Figure 12 illustrates this example. The upper node mapping input level to output level is placed at -15dB and is at the upper most limit of the un-processed histogram (the black trace). The lower node maps the input level from -60dB to -45dB , the light grey trace shows the effect of the dynamics processing on the track histogram and the **processed** level shows the average track level after dynamics processing. In this case we see the level has gone from -27.26dB (un-processed) to -21.95dB (processed).

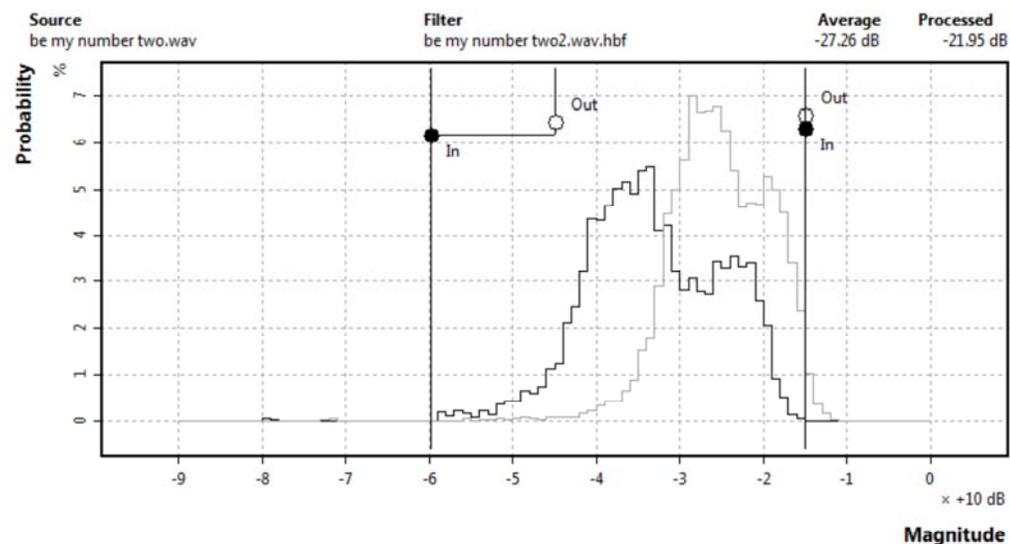


Figure 12: A compression processing example

If you are used to looking at compression in terms of transfer characteristics mapping input to output level and compression slopes, we can relate Har-Bal's histogram views and nodes to just that. This relationship is illustrated in Figure 13. The histogram on the left is obtained by transforming or projecting it through the input - output transfer characteristic (the solid dark line). When adjusting the output node you'll notice that the status line in Har-Bal displays information about the left and right slopes. These figures correspond to the transfer characteristic line slopes as indicated.

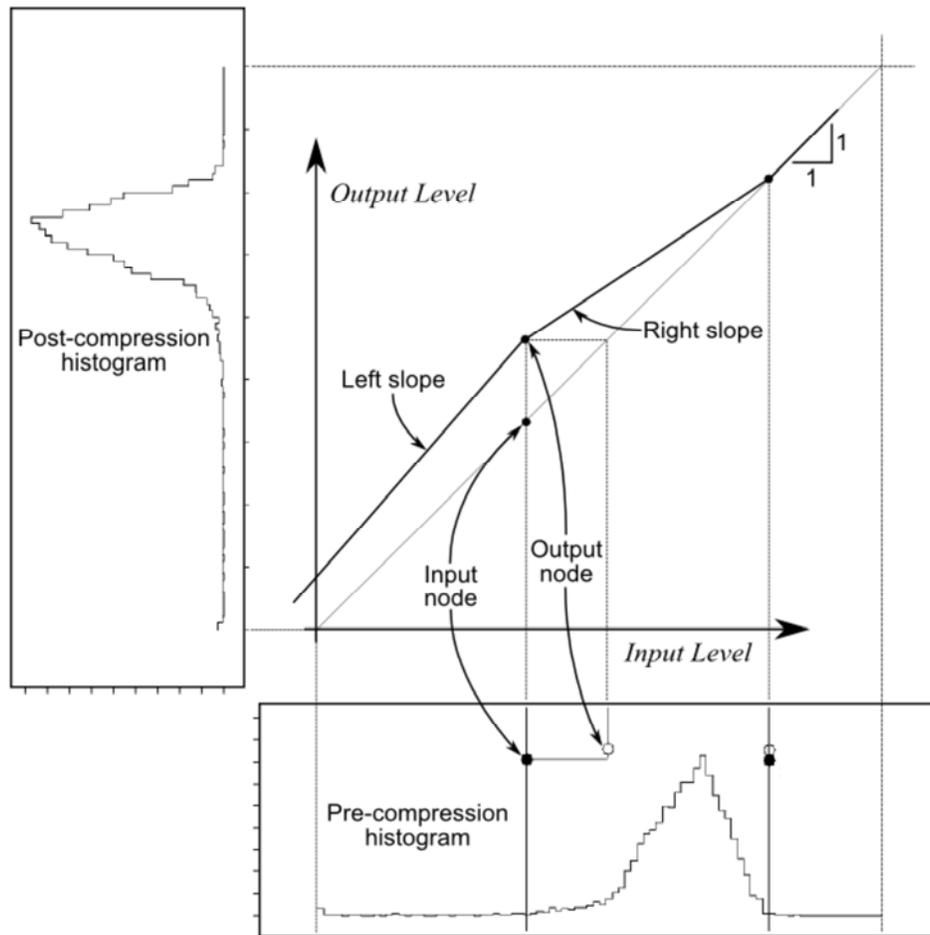


Figure 13: Relationship between the Transfer Characteristics and Dynamics Processing Nodes

A couple of observations can be made regarding histograms. When compressing, the width of the histogram is compressed but that squeezing makes the peak higher. This occurs because the area under the curve is constant and sums to 100% (ie. everything). You can think of it as a lump of plasticine you use to model a hill. If you squeeze the sides together then the plasticine pushes out at the top because the volume has to go somewhere. Similarly, if you stretch it, then it falls down (ie. dynamic range expansion).

Another thing to note is that this processing can be used for more than just compression. Another useful thing to use it for is noise gating. What I would typically do for that application is to provide a split at the end of the track where it fades away and apply dynamics processing to that split to gate the noise. The setting of the nodes is essentially exactly the same as with compression except you now map the input to a lower level on the left hand side to make the noise quieter. If you have additional dynamics processing applied to the overall track the two mappings are combined.

It all makes for a powerful and flexible architecture for managing the dynamic range of your track.

Example Filters

The best way to illustrate the power of the approach that Har-Bal 3 allows is through hearing the results obtained from well designed filters. This process is inherently objective and the resulting sound is heavily influenced by the interpretation and expectation of the person constructing the filter. As such I would like to point out at the outset that these filters are a product of my personal expectations for the tracks in question.

What I have done you may or may not agree with, though that is beside the point. Rather, in listening to and assessing the example filters given, I suggest that you adjust your mindset from "which version sounds best" to "how much can we transparently change the sound of a track through filtering in Har-Bal without damaging its integrity". In particular, you should pay attention to the filter transitions from segment to segment and try and pick them through listening alone. That is, without watching the timeline cursor during playback. Also note that most of the filters have a degree of gain cut applied globally to ensure that what you are hearing is not overly coloured by limiting upon limiting.

The examples given are deliberately commercial releases to avoid the criticism of the examples being contrived and not "real world" cases. They have already been mastered by professional engineers of some standing though I shall resist from naming them. Should you wish to know the necessary information should be discoverable on the net.

Finally, as they are commercial releases they must be paid for, so if you would like to hear how the filters perform you will need to part with a minor sum of money to the music download site of your choice, though if you choose to download from a location

other than the links provided below, be aware that you may have chosen a different mastering to the one I have used to design the filter. By downloading from the location given below you are guaranteed to be hearing what I hear.

Artist	Track	Filter
Coldplay	Fix You	04 - Coldplay - Fix You.wav.hbf
Paul Dempsey	Fast Friends	02 - Paul Dempsey - Fast Friends.wav.hbf
Foo Fighters	In your Honour	01 - Foo Fighters - In Your Honor.wav.hbf
Foo Fighters	Best Of You	03 - Foo Fighters - Best Of You.wav.hbf
Soundgarden	Room A Thousand Years Wide	08 - Soundgarden - Room A Thousand Years Wide.wav.hbf
U2	With Or Without You (Remastered)	03 - U2 - With Or Without You (Remastered).wav.hbf
U2	Bullet the Blue Sky (Remastered)	04 - U2 - Bullet the Blue Sky (Remastered).wav.hbf

More on Understanding Histograms

A histogram shows the frequency (as in how often something occurs) of an event. It comes from probability and statistics theory.

For a pure statistics explanation, lets say I have a bag of ping-pong balls with numbers written on them. The numbers are in the range of 1 to 10. Some of the numbers are the same. Lets say you empty them out and you see these numbers:

1,3,4,5,3,3,1,4,6,7,9,9,9,10,7.

A histogram showing the frequency of the numbers is constructed by counting how often each number occurs. So in the above case we have,

Number	Frequency
1	2
2	0
3	3
4	2
5	1
6	1
7	2
8	0
9	4
10	1

That is, there are two 1's, no 2's, three 3's, two 4's, one 5 and so on and so on.

A normalised histogram usually specifies things in percentages so to convert into a percentage we divide by the total number of balls in the set and multiply by 100. The frequency numbers therefore become,

$100 * (2 / 16)$, $100 * (0 / 16)$, $100 * (3 / 16)$,

because there are 16 balls in the set. That is what a histogram is. You may note that for normalised histograms the sum of all the frequencies is always 100% because that is the set of all balls in the bag.

Now in the context of Har-Bal's histogram, the balls are each of the sample points on the time line (which are spaced at a nominally 50ms interval) and the number on the balls is the average value for the average histogram and the peak value for the peak histogram. The main difference with the above example is that the time line is a real number (ie. fractional) so how do you group them. Well we group them in "bins" of 1dB width. What does that mean? Take a small subset of average values:

-0.3dB, -0.4dB, -0.7dB, -1.6dB, -2.4dB

The bins are centred on exact dB values of 0dB,-1dB,-2dB,-3dB The boundary from one bin to the next is the mid-point between centres. The top bin 0dB is special because it doesn't have an upper boundary.

bin	lower boundary	upper boundary
0dB	-0.5dB	
-1dB	-1.5dB	-0.5dB
-2dB	-1.5dB	-2.5dB
-3dB	-2.5dB	-3.5dB
.		
.		
.		

Going back to the sequence of average values, all we do is figure out which bin the number fits into and when we find that bin we add 1 to it because this corresponds to a count of 1 value fitting in that bin. -0.3dB fits in the 0dB bin, so does -0.4dB, -0.7dB fits into the -1dB bin, -1.6dB fits into the -2dB bin, -2.4dB fits into the -2dB bin and so on and so on. After counting all the values in the bins they are normalised by converting to percentages (100 times the bin count / total number of time line samples).

A histogram, like an average spectrum is much more useful for judging the dynamics of a track because it presents it in a static single image summary. Show me a histogram of the time line and I can tell you immediately whether it has high or low dynamic range, what the limits of the dynamic range are, whether it has a bi-modal behaviour (as in loud parts and quiet parts), whether it has been over limited and so on and so on. It is all there to be seen in the histogram.

If you find it hard to believe just open up a track you know to have high dynamic range and a track you know to have low dynamic range. Compare the histograms. Can you see the difference? The pattern is obvious and immediately revealing.

The histogram also contains within it the total average level within it. If you give me the histogram I can calculate the track average from the data it presents. It makes a great deal of mathematical sense to summarise dynamics with histograms.

Why is Har-Bal 3 slower at updating filters?

If you are a Har-Bal 2.3 user used to playing a track real-time and making and hearing Har-Bal filter changes in near real time you'll probably be wondering why Har-Bal 3 is so slow at updating. Well, it isn't because the code is slower but rather, what it has to realise is much greater. The detailed answer is as follows.

Har-Bal 2.3 realises only left and right filters with a common response (ie ganged). In the basic Har-Bal 3 filter block Har-Bal realises four filters. It has two extra responses to realise : the new Haas zone ambiance derived "air" which uses cross-couple filters. That is the first two fold factor difference but there is more.

To keep CPU usage during playback to a minimum Har-Bal uses an equivalent filter to realise the independent left, right, mid, side and air filters. By doing so, we reduce the complexity by half but it means that all those responses are lumped together, which consequently also means that if you want to have real-time switchable air and eq (toggle air, toggle eq) then we have to calculate four different combination's of filter responses, three that are non-trivial. That is,

```
no air + no eq, trivial
air + no eq
no air + eq
air + eq
```

So now we have a 6 fold increase in the amount of calculation to realise Har-Bal 3 filters over Har-Bal 2.3.

In the future that will increase further because a proposed feature is additional room compensation equalisation in Har-Bal 3. It's already coded for but there is no front end to add or indeed design the responses. With room EQ it then becomes,

```
no air + no eq + no room eq, trivial
air + no eq + no room eq
no air + eq + no room eq
air + eq + no room eq
no air + no eq + room eq
air + no eq + room eq
no air + eq + room eq
air + eq + room eq
```

So with that version it will grow to be a 14 fold increase in the amount of calculation required.

Given the added amount of calculation you'll probably be asking what's the point. The point is that with the more powerful analysis and the more powerful filtering, Har-Bal 3 can improve on tracks with relative ease that Har-Bal 2.3 can never make headway on. I know that with certainty from experience with using it. There is so much stuff around today that Har-Bal 2.3 just can't help with a great deal, especially when hyper-compression is involved.

For me personally, despite the slower speed at realising filters, I feel absolutely no desire to go back and use version 2.3 and if I'm forced to 2.3 I'm always looking for the Har-Bal 3 features it can't do and despairing. The near real time filter update in Har-Bal 2.3 is neither here or there to me. It is entirely in-consequential. If it is critical to you I can't offer a solution because there isn't one, bar not offering all the added processing power of Har-Bal 3 which is something retro-grade and of no value to me.

Real-time filter updates isn't even in the design use case as far as I'm concerned so the lack of it doesn't rate as a fault because it is not a design goal. If that is what you need and can't live without you'll either have to stick with Har-Bal 2.3 for as long as it will continue to work on current generation machines, or go with the myriad of other EQ tools you can find in plugin format, or as dedicated hardware units.

What is important in the Har-Bal design is the mid /side filtering capability, the track splitting and near instantaneous update to spectrum averages based on whatever your filters do, the dynamics processing and the predictive average level estimates, real-time spectrum, complex time selection and filtering to control rogue resonances (when that feature is implemented), room eq compensation (when that feature is implemented), improved air that even works on mono sources, playable track referencing to allow you to easily bring consistent sound to compilations and all the other aids Har-Bal provides.

I don't use Har-Bal in a realtime fashion, not even Har-Bal 2.3. Rarely have and I dare say, rarely will. That mode of operation has no value to me. If it were that important to me I can do what I pretty much always do, which is to do spectrum editing with nothing playing, then play and toggle EQ to hear the difference. Sound better, keep it, sound worse, undo it. It isn't a usability issue for me. Not having the power to do everything else that Har-Bal 3 offers is.

There are a few other points that should be mentioned regarding this version and it's speed. It is not compatible with older machines and by older machines I mean ones more than four years old. The reason being that current generation intel processors do floating point faster than fixed point so all the signal processing in Har-Bal 3 is now in floating point. Older machines are slower on floating point than fixed point so when run on an older machine Har-Bal 3 will seem very slow. Fast disk access is a necessity for Har-Bal 3 because it needs to read and write a lot of information to files to support the real-time spectrum and track splitting. As such, you should never open a file on an external drive, be that a network drive or a thumb drive. Doing so will almost certainly result in very slow response.



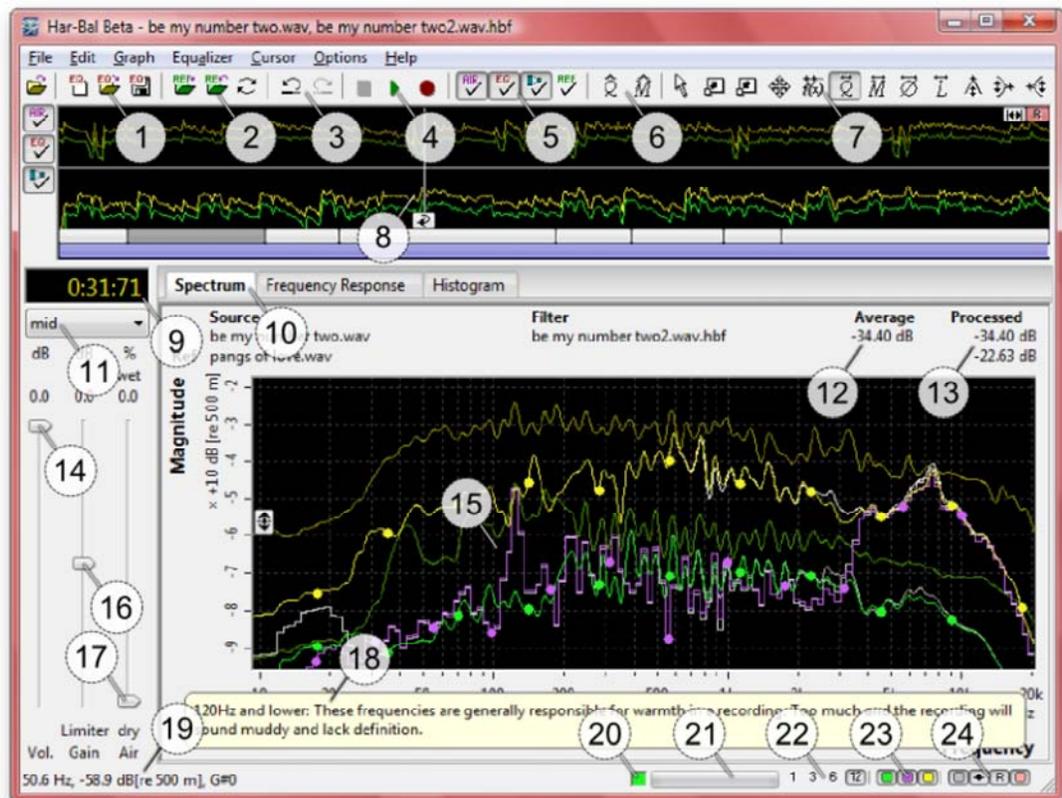
Har-Bal Harmonic Balancer

Har-Bal™ new thinking, new directions...

Har-Bal Tutorials

Har-Bal 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 Har-Bal. Before continuing and further, I thoroughly recommend that you first read the [Har-Bal introduction](#) as you will learn a great deal about the structure of the software, its capabilities and the thinking behind its implementation.

Har-Bal Layout



1. Filter related toolbar buttons

These toolbar buttons allow you to open, save and create new filters for the track being harmonically balanced.

2. 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.

3. Filter design undo/redo toolbar buttons

These toolbar buttons allow you to undo and redo filter design changes.

4. Playback/Record control related toolbar buttons

These toolbar buttons allow you to playback and/or record the current track with or without the filter in circuit.

5. HB Air, Filter, Dynamics and Reference control toolbar buttons

These toolbar buttons allow you to switch HB Air in and out, switch the filter equalisation in and out, switch dynamics processing in and out and switch the reference track playback in and out.

6. IntuitQ and IntuitMatch toolbar buttons

These toolbar buttons apply the IntuitQ and IntuitMatch algorithms to the spectrum content of your track, with the frequency range of application being determined programatically.

7. Filter editing and view control toolbar buttons

These toolbar buttons let you control the graphical view and design the Har-Bal filter.

8. Time line control

This control cues the track playback to the desired position, selects looped playback, creates and selects track segments, navigates through segments and selects segment references.

9. Time display

This display shows the live time in seconds of the track being played.

10. View Selection Tab

This tab controls is used to select between Spectrum, Frequency Response and Histograms views.

11. Spectrum selection dropdown list

This dropdown list allows you to select the spectrum to view and design a filter for. It is one of Mid, Side, Left or Right.

12. Average power figure of merit indicator

This indicator displays the un-weighted average power figure of merit for the current track without any processing. This value is relative to the power of a full-scale square wave.

13. Processed figure of merit indicator

This indicator displays the un-weighted average power figure of merit for the current track with dynamics processing and applied gain. This value is relative to the power of a full-scale square wave.

14. Volume control slider

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

15. Spectrum view

The spectrum view 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 real time spectrum for the track. We make use of the spectrum editing tools to design Filter equalisation schemes.

16 Limiter gain control slider

This limiter gain control slider allows you to apply a gain to the playback and/or recording of the current track. This is typically used for normalising track levels. This gain is on the input side of the integrated Har-Bal limiter, which limits the final output to -0.5dB relative to full scale.

17. HB Air control slider

This HB Air control slider allows you to apply the HB Air effect to the recorded and played back tracks. This is typically used when your track lacks a sense of space.

18. Tips popup

The tips popup gives spectrum context sensitive clues to designing appropriate equalisation filters.

19. 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 views.

20. Cursor focus indicator

This indicator informs you which spectrum trace is being edited when using the parametric and shelving cursors.

21. Limiter gain reduction meter

The limiter gain reduction meter gives a graphical indication of the amount of gain reduction being applied during playback by the limiter. The meter is a log scale meter (linear in dB) with full-scale corresponding to -15dB.

22. Spectrum resolution buttons

These buttons allow you to alter the resolution of the displayed spectrums. You can switch between 1/12th octave, 1/6th octave, 1/3rd octave and 1 octave resolutions.

23. Spectrum view toggle buttons

The spectrum view toggle buttons allow you to toggle on and off the average, peak and real time traces.

24. Spectrum view control buttons

The spectrum view control buttons allow you to toggle on and off the original spectrum, spectrum focus indicator, reference and segment reference traces.

Where to Start?

Har-Bal works with 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.

It is recommended that source material should be arranged such that each track in a given compilation is in a separate file. You could, if you wish, perform harmonic balancing on a single file corresponding to an entire album but the outcomes that will result is likely to be less than optimal. This is because each track in a compilation has its own specific characteristics, which when lumped together, become lost in the whole.

Copyright © 2011
Har-Bal™. All rights
reserved.

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 Har-Bal 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, Har-Bal will open it immediately by reading the analysis file of the track created on first opening it. Alternatively, if you have a compilation of many tracks that you are going to work on, you can pre-analyse those tracks in the background while you work on another file in Har-Bal by using the [File | Batch Analysis](#) menu command.

For the sake of this tutorial I shall be demonstrating Harmonic Balancing of a selection of tracks from the **U2** album, **The Joshua Tree**.

Beginning a Session

Assuming that you are mastering new material or re-mastering old material then the first thing you should do is thoroughly listen to the compilation of tracks to build a mental picture of the sound the producers are trying to achieve. Make note, mental or otherwise, of anything that concerns you.

After doing so pick the track from the compilation that has a broad spectrum (typically tracks with the most instruments and instruments that are played loudly) but also the best sound quality. The reason for this choice is that the first track we master will become a reference for the remaining tracks. If you were to choose a slow song it won't fill the entire spectrum and then you will have some difficulty drawing inferences from your reference if you were to then open a busy track.

Using the [Open Command](#) I open track 4, **Bullet the Blue Sky** as the first track as it has a full spectrum and a reasonable balance. When the analysis runs to completion the spectrum view shows the following result.

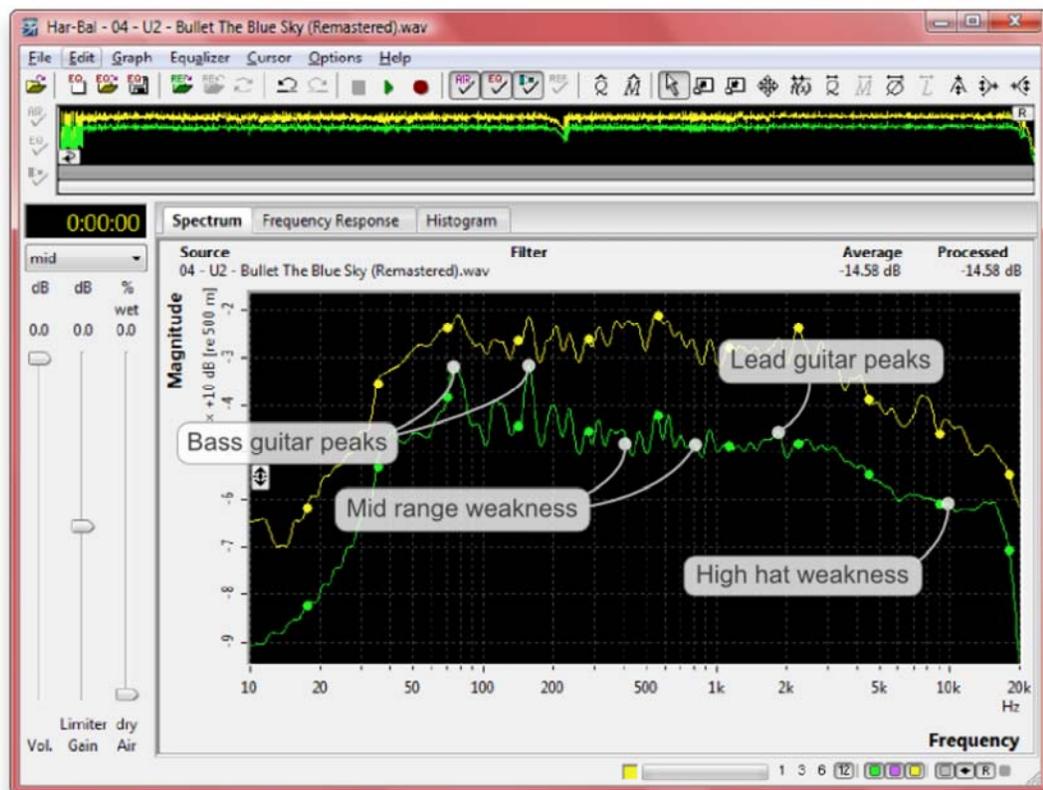


Figure 1: Mid-spectrum view of "Bullet the Blue Sky"

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 violet trace shows the real-time spectrum content.

On listening to the whole track the following observations are apparent. The track has a prominent bass guitar part that has a boxy character. Whilst the drum part is a prominent part of the track, the high hat sounds as though it is suffering from destructive interference arising from leakage of the high hat sound into other microphones on the drum kit (perhaps the overheads). The snare drum, though prominent, also sounds a tad wooden, lacking a sense of crispness. The lead guitar is largely well controlled but has occasional stridency in the upper mid-range (2 – 3kHz). Finally the lead vocal sounds brittle in places, though generally well controlled.

Looking at the spectrum we see that it extends from around 40Hz to 18kHz. There are two prominent peaks at around 75Hz and 150Hz corresponding to the bass guitar part, which are most likely responsible for the box character of the part. The region between 4 – 14kHz is somewhat subdued and corresponds to the weakness in the high hat part. There is possibly some weakness

in the mid range area in places. Rather than attempt to address these issues in this view we should instead split the track along song structure lines to obtain a better understanding of the sources of the imbalances and a better corrective filter for the track.

Splitting the track into its parts (intro, verse, chorus, verse, etc.) is achieved with the split shortcut command applied to dragging the time line cue point. I find a useful what of identifying the correct point to split at is to first zoom into the region where a transition from one part to the next occurs. To zoom mouse over the timeline, press the left mouse button down and drag the highlight over the region of interest.

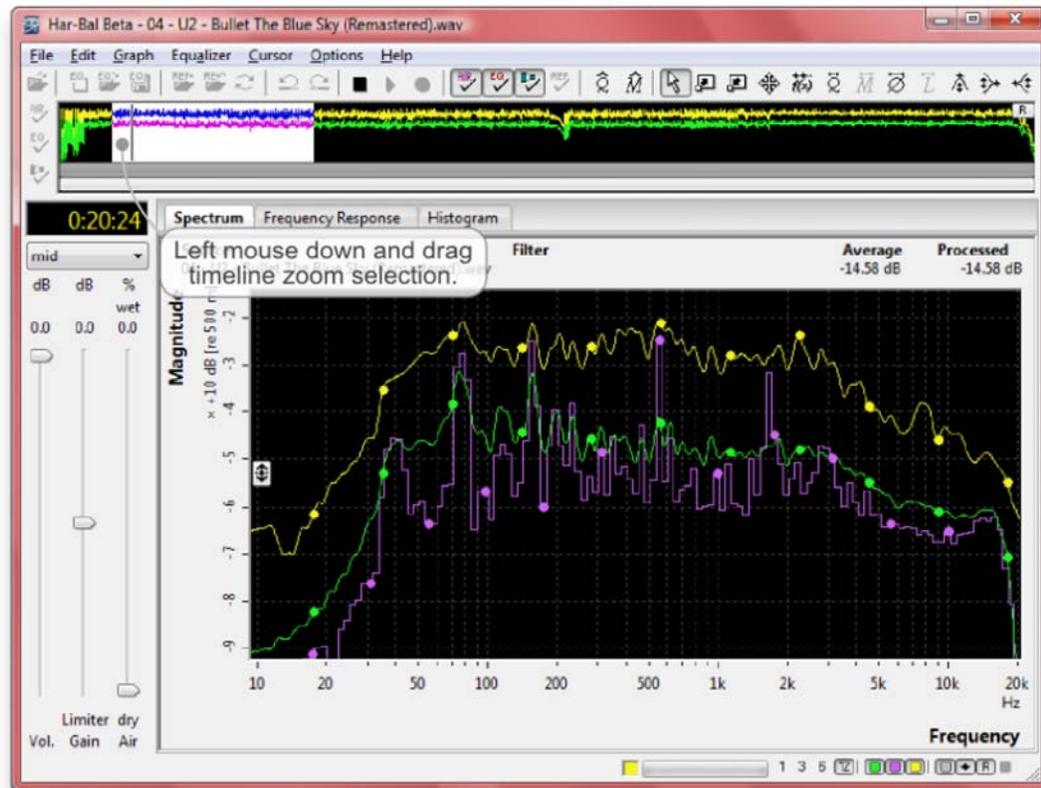


Figure 2: Zooming in to the time line of "Bullet the Blue Sky"

I find the best way to create a split is to zoom in to a close region either side of the transition point and set the track playing close to the point of transition from one part to the next. Then watch the time line track cursor and listen for the transition. When you think you know the general region where it occurs, then press down the left mouse button on that point and release it to resume playback from that point. If the location is correct you will hear the entire beginning of the section cleanly and if you chose wrong then it will sound discontinuous. If the latter move you location in the direction appropriate and try again until you have found the transition point. At that stage press and hold the left mouse button down on that position in the time line and press **Ctrl+S** to create the split point. This process is illustrated in Figure 3.

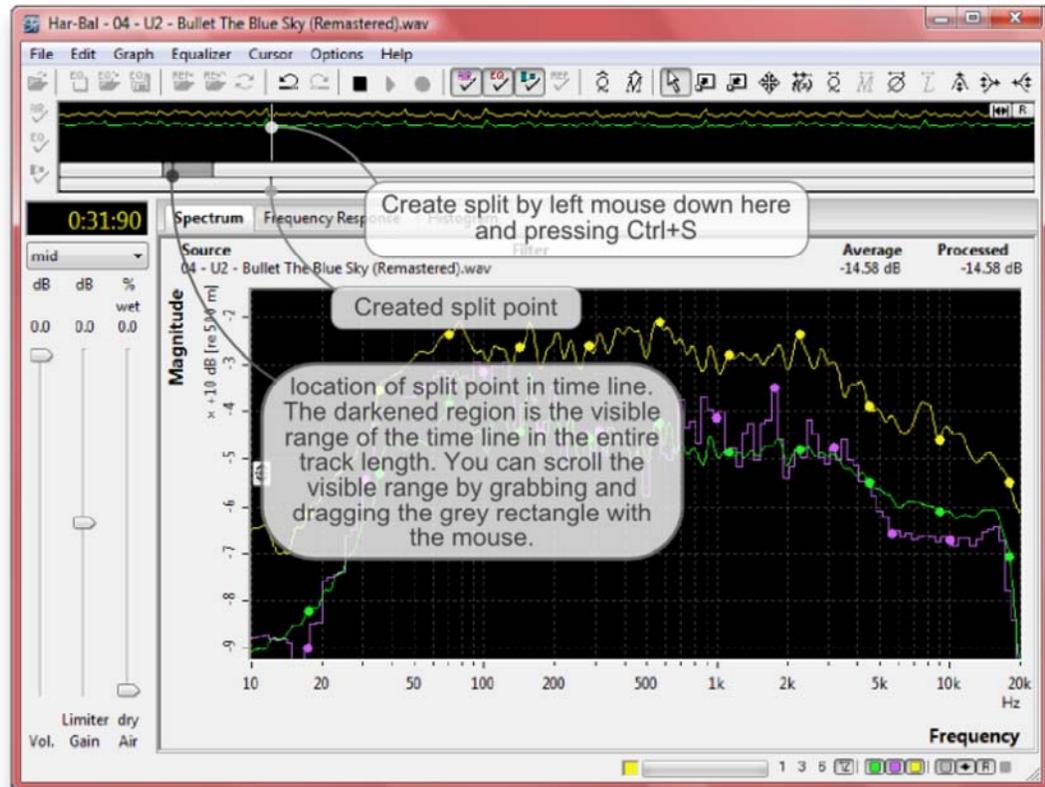


Figure 3: Creating split points

To avoid loosing work it is useful to occasionally save changes to a Har-Bal filter file. Selecting the **File | Filter | Save** Command, we save it to the file *"04 - U2 - Bullet The Blue Sky (Remastered).wav.hbf"*. Then we continue on adding split points to the track time line. We add 9 split points at the times outlined in Table 1.

Split number	Split Time	Description
	0:00:00	Track introduction
1	0:31:90	Track introduction, lead guitar opening
2	0:44:62	First verse
3	1:08:31	Chorus
4	1:31:95	Second verse
5	1:56:09	Chorus
6	2:20:48	Bridge
7	2:28:70	Third verse
8	3:17:50	Guitar solo
9	4:06:68	Guitar feedback exit

Table 1: Split Point Times

Having created the track segmenting we can now begin constructing an equalisation scheme for the track. Start by selecting the first segment. This is done by mousing over the first segment and clicking it with a left mouse button press.

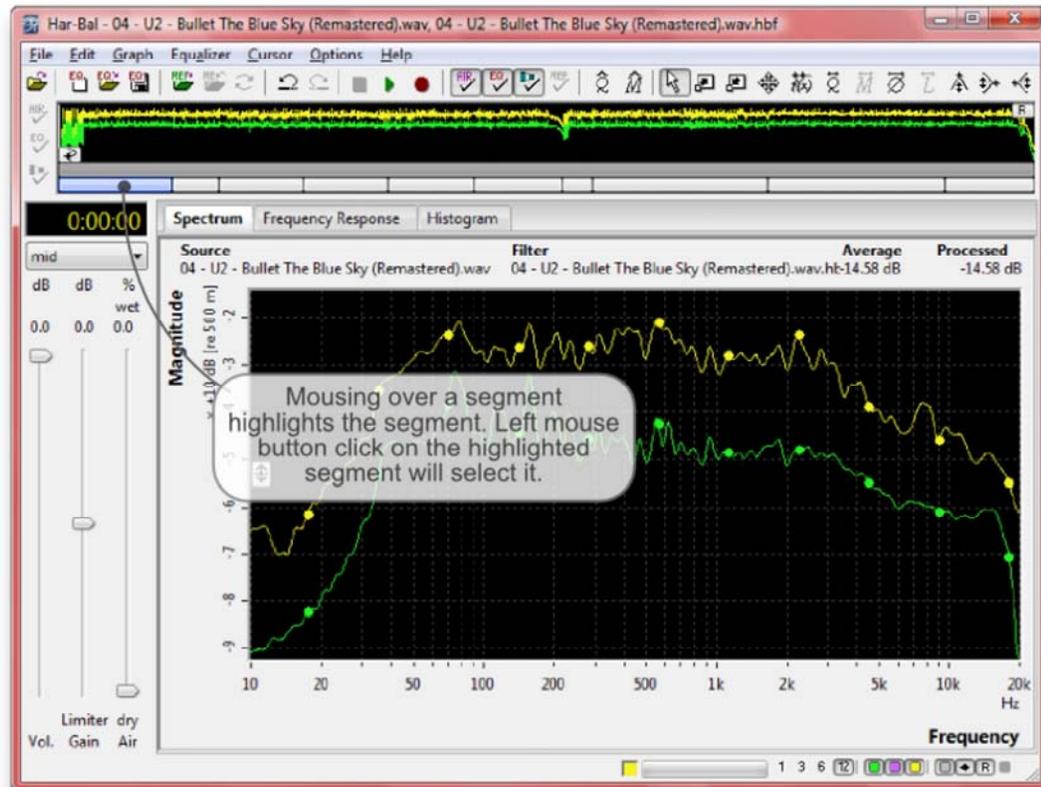


Figure 4: "Bullet the Blue Sky" track segmentation : selecting the first segment

Before continuing on with the design of the equalisation scheme for this track we need to introduce and discuss the technique of empathetic equalisation.

Empathic Equalisation : What is it and why does it work?

Empathetic Equalisation can be loosely defined as the process of high Q notching and peaking applied to the tracks average spectrum (in the 1/12th octave resolution view) in such a way as to produce a uniform average spectrum when viewed with 1/3rd octave resolution. It isn't a rigid prescription but rather, a recommended approach that needs to be tempered by what you hear. If a particular edit produces detrimental results then undo it, re-assess your approach and try a different variation.

Rule 1:

If you have a high peak that needs cutting then cut the adjacent troughs to the peak or cut at the mid-points between the peak and the adjacent troughs using high-Q parametric tool edits.

Rule 2:

If you have a low peak that needs boosting then boost the peak directly with high-Q parametric tool edits.

Rule 3:

If you have a broad and solid group of peaks with weak troughs that needs reducing in level then reduce the group by using high Q cutting of the weak troughs, aiming to generally maintain the trough depth.

Rule 4:

Similarly, If you have a broad and solid group of peaks with weak troughs that needs boosting in level then boost the group by using high Q boosting of the peaks.

Rule 5:

In general you should shape the average spectrum so that the resultant spectrum overlays the original spectrum but is smoother, particularly when viewed in 1/3 octave spectrum resolution. The reason for aiming for this type of shape is that it ensures you are respecting the mix and not attempting to drastically alter it, whilst reducing the masking of one instrument by another, thereby bringing out more detail in the recording. The assumption here is that mixing was carried out in a tonally neutral environment with neutral speakers and ears. This is not always the case so the case may arise where you do end up altering the mix composition to compensate for poor conditions. Rules 1 – 4 are illustrated in Figure 5.

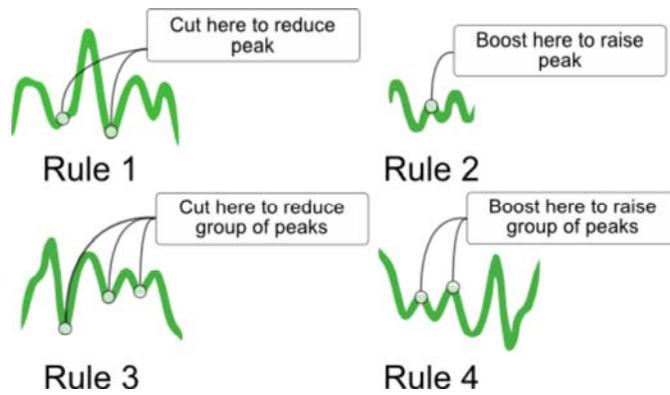


Figure 5: Rules 1 – 4 of the Empathetic Equalisation technique

So why use these rules? Is there any deeper reasoning as to why we should equalise using this approach? In fact, there is justifiable reasoning behind it. The general consensus you are likely to hear elsewhere is that you should only use low Q boosting and cutting and usually with only a small degree of boost or cut, yet we are suggesting essentially the complete opposite of the widely held view. Why so?

If you happen to look at the structure of musical notes played on typical instruments you will generally find that each played note (or percussive hit for that matter) is composed of a short lived transient component that generally has a broad spectrum and a much longer lived resonant component that is largely confined to a handful of frequencies, the harmonics as it were. This is clearly illustrated with a sonargraph of a single note plucked on a guitar string as shown in Figure 6. It is clearly visible that immediately after the string pluck a broad range of frequencies are generated but the high frequency components quickly die away to leave us with the long lived harmonics. Indeed, the initial broad band spectrum is essentially atonal and corresponds to the spectrum content of the impulsive event, in this case the act of plucking, giving rise to the string vibration. The components that live on are the ones that resonate in the guitar string.

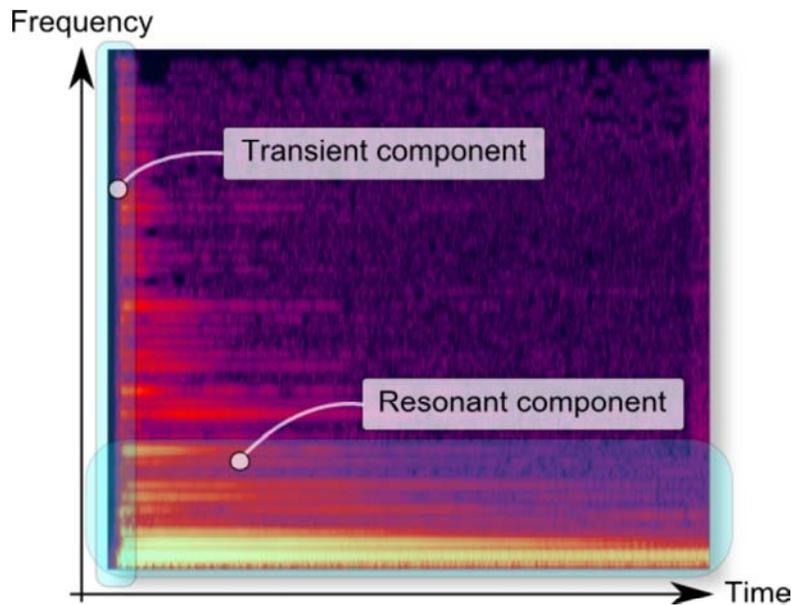


Figure 6: Sonargraph of a plucked guitar string. Brighter colours correspond to higher spectrum magnitude. Note that the transient component of the sound covers the entire frequency range and is short lived, whilst the resonant component is long lived but confined to resonant frequencies.

A sonargraph belongs to the domain of joint-time and frequency analysis as both time and frequency are involved in the presentation of the analysis. The Har-Bal average spectrum however, is purely a frequency domain phenomena. Time has been averaged out. This means that both the transient and the resonant components of the spectrum resulting from the many events making up a musical performance are present in the average spectrum result. The strong peaks correspond to the long lived resonances and the content in between is at least partly made up of the short lived transient components. Figure 7 shows a simplified representation of this superposition.

Now consider what happens to the relativity between the transient components and the the resonant components when we apply equalisation. If we were to use the traditional approach of cutting overbearing peaks by cutting the peak itself, then the harmonic component is reduced in magnitude more than the content in between the peaks, which is where the transient component lies. In effect we are reducing the resonant component and increasing the transient component. That may not sound like a bad thing given how often transient response is made mention of in hi-fi equipment reviews but it actually is very bad.

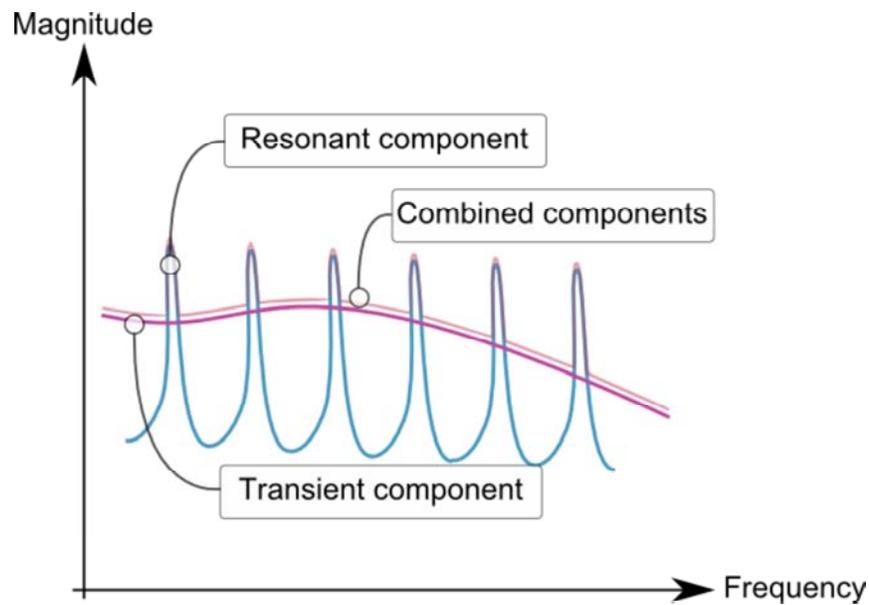


Figure 7: How transient and resonant spectrum components combine to form the average spectrum that Har-Bal sees. The transient component is relatively smooth and spans a wide bandwidth, whilst the resonant component is confined to discrete frequencies. Har-Bal sees the two components combined together.

Lets think back to our plucked guitar string example and consider how we might increase the relative amplitude of the transient component over the resonant component. All we need to do to achieve that is to kill the resonance which can easily be achieved by muting the string with a finger. Now think of how a muted guitar string sounds when plucked : a dull thud is an apt description! It may be an extreme example but the conclusion generally holds in the sphere of equalisation, which is that cutting peaks directly tends to a dull lifeless sound.

Now lets go back to the generally held belief that you should only apply low Q cutting and boosting and only in moderation. If we consider what happens in this case we see that because of the low Q filtering, any peak modification will have almost the same (but not quite) effect on the content in-between the peaks, thereby not damaging the sound quality much. But now consider the empathetic equalisation approach of cutting a peak by cutting in the adjacent troughs. In this case the greatest cut occurs in the troughs meaning that the resonant content is raised in level relative to the transient content. Thus, with high Q peaking and cutting filtering placed precisely, we can actually make radical changes to the spectrum without bringing about the dullness in sound that often results if trying the same thing using more conventional equalisation equipment. In short, though the low Q filtering recommendation made by mastering engineers does apply to regular equalisation equipment, it does not apply to Har-Bal because Har-Bal has the analysis power to be able to place the boosts and cuts in pricisely the right location to avoid the problems that can result from poorly target equalisation.

Here is another contrived example to illustrate this mechanism. We have a 440Hz sine wave mixed with pink noise and propose cutting its level by a nominal 5dB using normal peak cutting and empathetic equalisation trough cutting. The two different filter realisations are illustrated in Figure 8. If we apply these two filter realisations to the pure tone and the pink noise we obtain the results shown in Figures 9 & 10 (in the context of resonant and transient components, the pure tone represents the resonant component and the pink noise represents the transient component). Note that the pure tone spectrum measurement has a noise component mixed in to stop Har-Bal from introducing a large loudness compensation gain change.

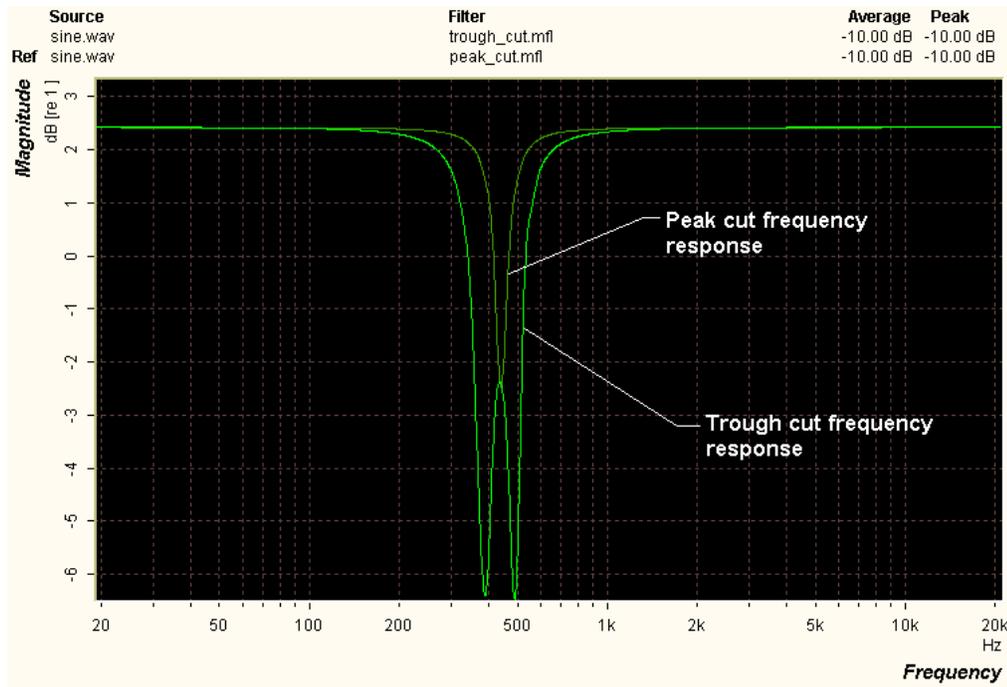


Figure 8: Conventional and Empathetic Equalisation Filter Realisations

In the pure tone case the 1/3 octave spectrum (measured using a separate spectrum analyser on the Har-Bal filtered wave file) amplitudes of the 440Hz peak are the same, which is what we would logically expect to see. However, in the case of the pink noise the amount of attenuation afforded to the pink noise within that same 1/3 octave band is significantly greater for the Empathetic Equalisation approach (see Figure 9).

More importantly we note that in the case of a conventional peak cut filter response the tone is attenuated by 4.74dB whilst the noise in that third octave band is attenuated by only 2.04dB. Thus we have a loss of separation between the noise and the tone of 2.7dB with this filter. In contrast, with the Empathetic Equalisation approach the tone is again attenuated by 4.74dB but in the noise in that third octave band is now attenuated by 5.83dB. Hence we have an increase of separation between the noise and the tone with the Empathetic EQ filter by 1.09dB!

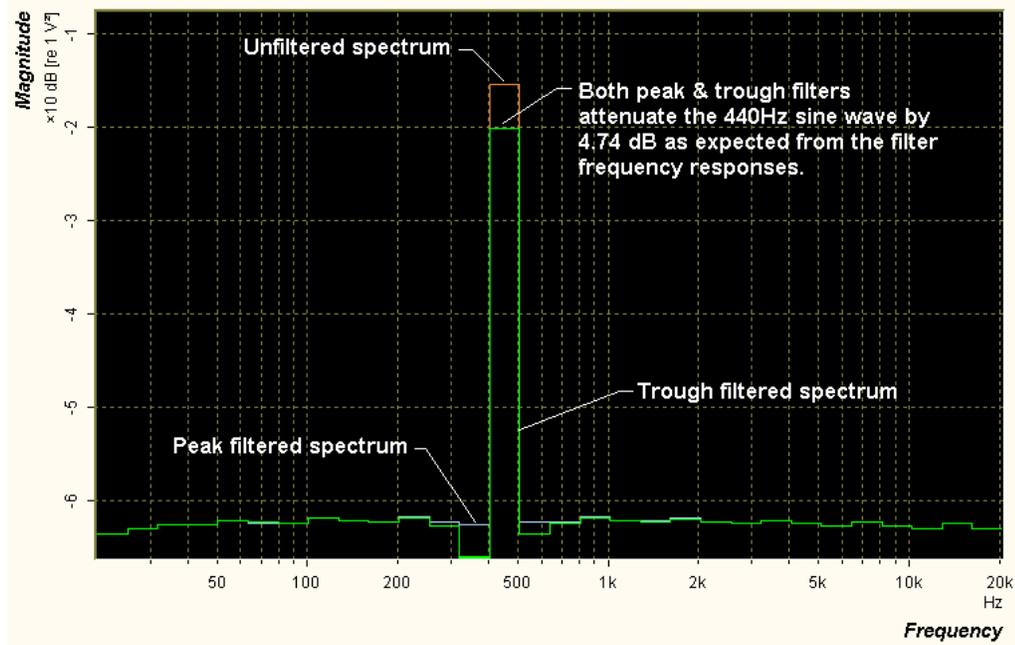


Figure 9: Attenuation of 440Hz Tone by Conventional and Empathetic Equalisation Filter Realisations

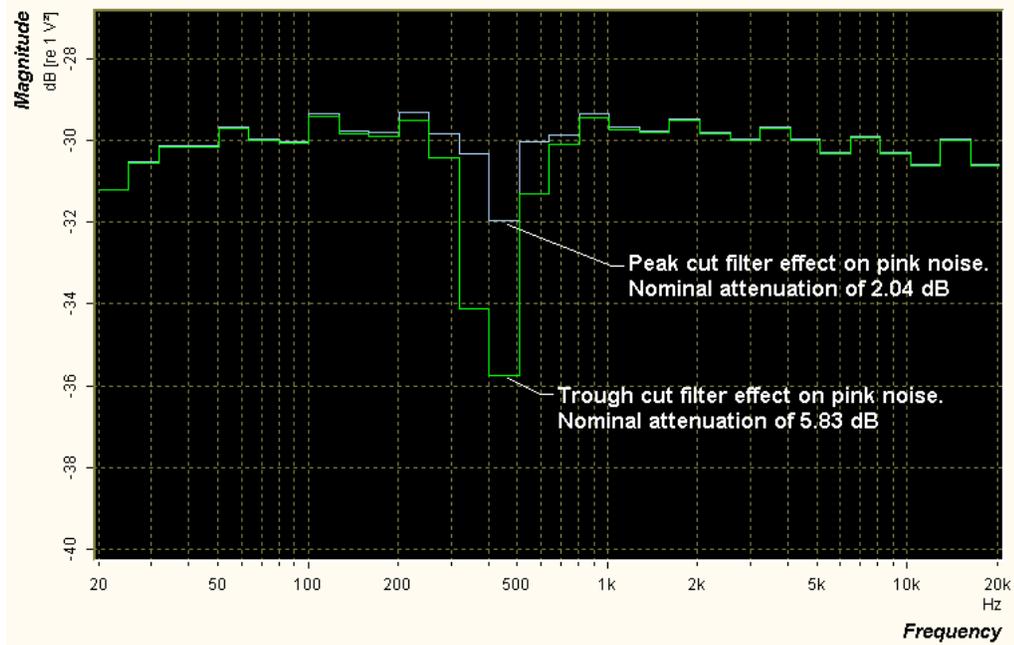


Figure 10: Attenuation of Pink Noise by Conventional and Empathetic Equalisation Filter Realisations

Bullet the Blue Sky Continued...

Having selected the first section of our partitioned track, let's begin by designing a filter for it. The first thing we should do is switch to 1/3 octave resolution so that we can get a better understanding on where the imbalances in the mix may lie. The importance of 1/3rd octave stems from the tuning of the basilar (the hairs which convert sound to nerve impulses) in human hearing, which approximately corresponds to 1/3rd octave bands (except at that low frequency end where the bands are narrower). 1/3rd octave view therefore gives an impression of the balance that is closely consistent to the way we hear.

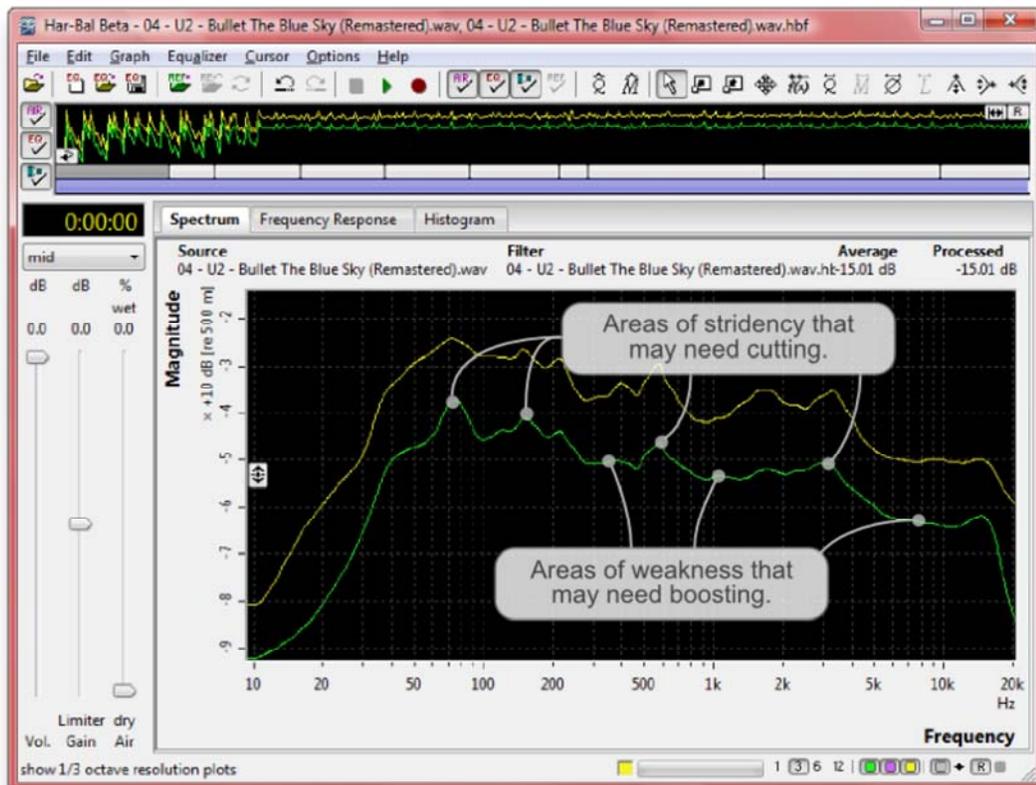


Figure 11: "Bullet the Blue Sky": the first segment in 1/3rd octave resolution

Looking at the first segment spectrum for "Bullet the Blue Sky" (see Figure 11) in 1/3rd octave resolution, we clearly see areas of weakness and strength over the entire spectrum. There is prominent peaking at 75Hz and 150Hz corresponding to the bass guitar output and prominent peaking at 600Hz and 3kHz corresponding to snare resonance and the lead guitar respectively. There is significant weakness around 300Hz, 1kHz and 7kHz areas.

These areas may or may not present auditory problems. The only way to determine if they do is by bringing in more uniformity and seeing if it actually sounds better. The question is, given our hypothesis to test and the empathetic equalisation recommendations on smoothing spectrums, how do we go about achieving this result? Generally speaking, we start by selecting the gain cursor and the 1/12th octave spectrum resolution and then introduce peaking and cutting at appropriate locations. So how do we do filtering with the gain cursor?

The gain cursor is equivalent to a parametric equaliser. The centre frequency is chosen by moving the cursor to the desired frequency on the spectrum plot and then pressing and holding down the left mouse button. The gain is controlled by moving the mouse up or down and the Q is controlled by moving the mouse left or right. During the dragging operation a preview of the filtered spectrum follows the mouse movement. To keep the Q fixed while making boost or cut, hold the SHIFT key down. To keep the Q at the maximum value press and hold the M key down. When you release the left mouse button the filter change is made and the spectrum plots updated. If you wish to cancel an edit press the ESCape key before releasing the left mouse button.

The cursor preview is attached to either the peak or the average spectrum plot and you can change the focus from one to the other by pressing the TAB key. Which plot holds focus is shown in the status bar by the focus indicator. The status bar also displays the parametric EQ centre frequency, gain and Q. The values are displayed in scientific notation, which means large and small values are prefixed with the appropriate letter indicating a base 10 multiplier as indicated in Table 2. As an example, 1.1kHz means 1100Hz. Figure 12 illustrates these aspects of Har-Bal.

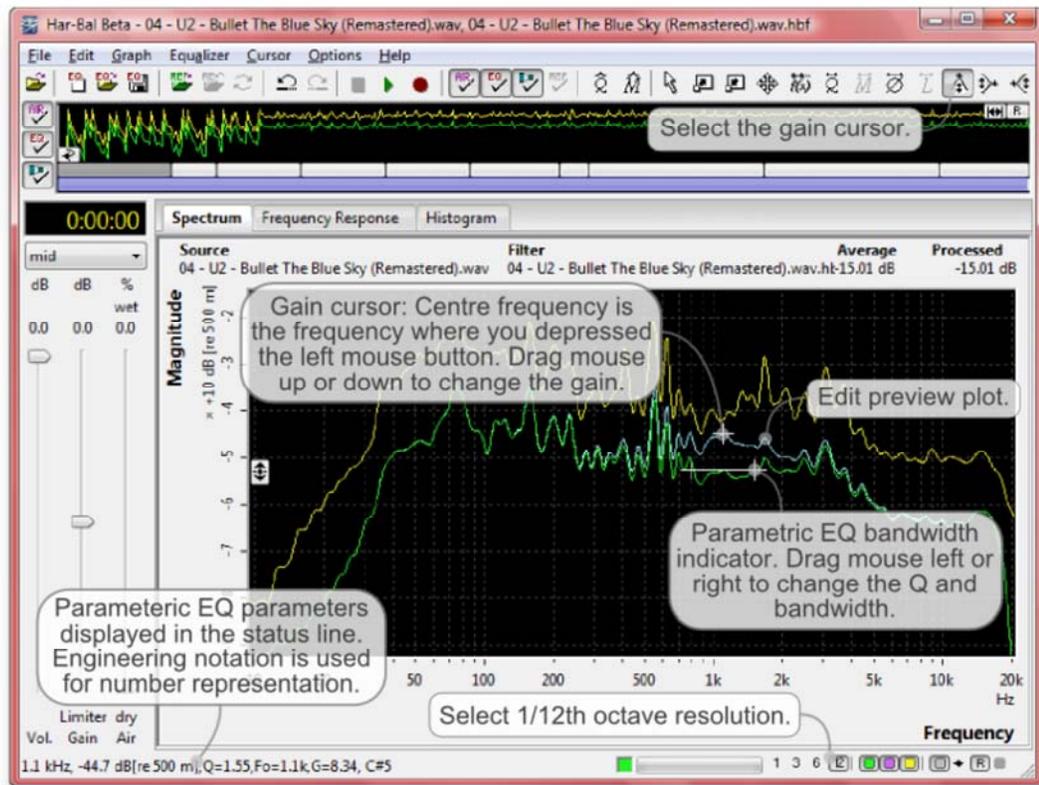


Figure 12: Aspects to Performing Parametric EQ Filter Edits

Engineering Prefix	Multiplier
T	1000000000000.0
G	10000000000.0
M	1000000.0
k	1000.0
m	0.001
u	0.000001
n	0.000000001
p	0.000000000001

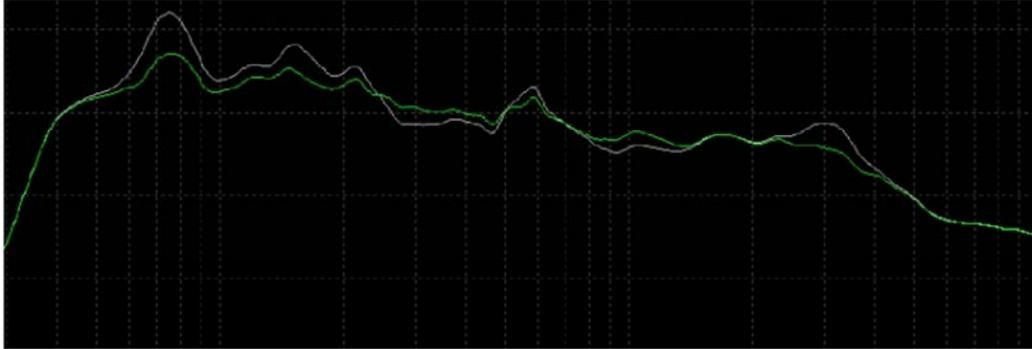
Table 2: Engineering Notation Prefixes and their Corresponding Multipliers

Continuing on, we introduce high Q peaking on the peaks and cutting in the troughs in such a way as to smooth out the average spectrum when viewed in 1/3rd octave resolution. We don't alter all peaks and troughs but rather just as much as is necessary to bring a level of smoothness to the spectrum.

A point to make in this regard is that the level of smoothing to aim for will depend on the density (in terms of number of instruments) of the recording. If your recording is of a solo instrument great care needs to be taken to not destroy the character of the instrument. With more instruments timbral alteration has much less influence and is generally encouraged (mixing generally involves using EQ to allow different instruments to blend in a more pleasing way) for the sake of allowing each instrument to be heard in a competing landscape.

Returning to the process, we smooth by applying edits in 1/12th octave resolution and then checking the 1/3rd octave resolution to see if we are smoothing in the right direction. After a bit of practice you become adept at recognising what is needed in 1/12th octave resolution and rarely need to switch to 1/3rd octave resolution, except for the final check and balance. Using this approach I came up with the following filtered spectrum for the first segment (see Figure 13).

1/3rd Octave



1/12th Octave

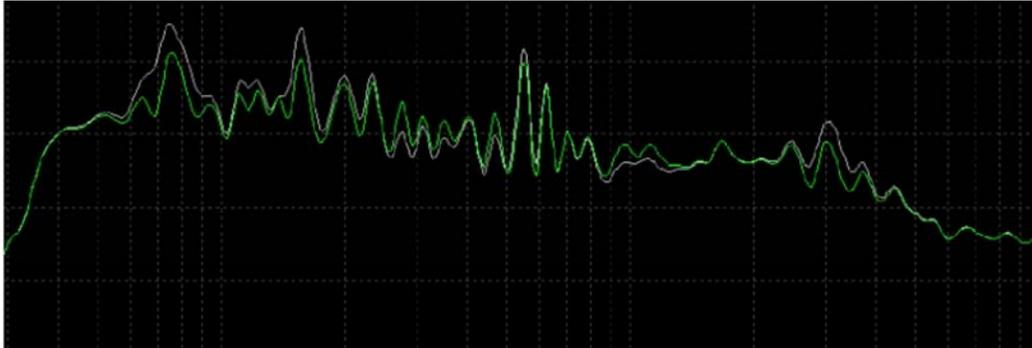


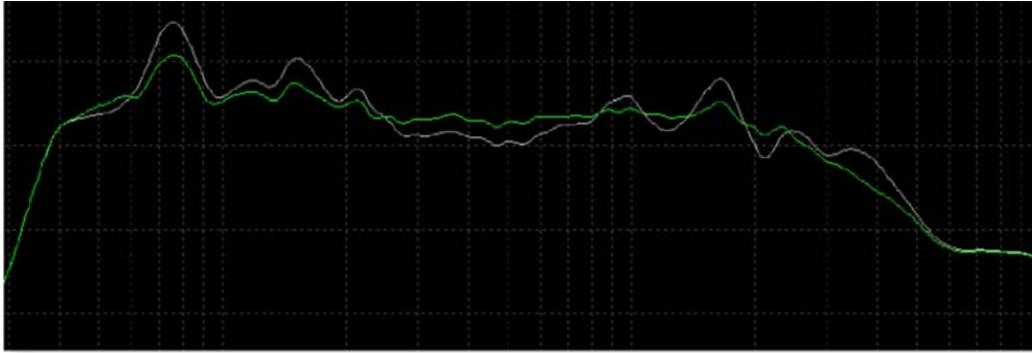
Figure 13: Close up of the first segment filtered average spectrum in 1/3rd and 1/12th octave resolution. Note how the filtered spectrum nominally overlays the original to preserve the overall mix intention and cuts are achieved by cutting troughs and boosts are achieved by boosting peaks.

The average spectrum is significantly smoother than the original and holds up well on playback. The snare drum shows more definition and the high-hat has extra clarity whilst the bass guitar and kick drum are still prominent parts of the mix. In 1/12th octave resolution you can see how the smoothing has been achieved through selective peaking and cutting.

Moving on, we select segment two and apply the same filtering approach to obtain the smoothing results of Figures 14. Likewise, the results of filter designs for segments three through to nine are shown in Figures 15 through 21. Studying these filters, the general approach of smoothing using the empathetic equalisation technique is readily seen.

What is readily apparent from this approach and which dramatically sets it apart from other methods of equalisation is the level of detail in the tailored filter frequency responses. It is clear from the 1/12th octave spectrums that to produce a similar response using a hardware equaliser would be difficult, if not impossible.

1/3rd Octave

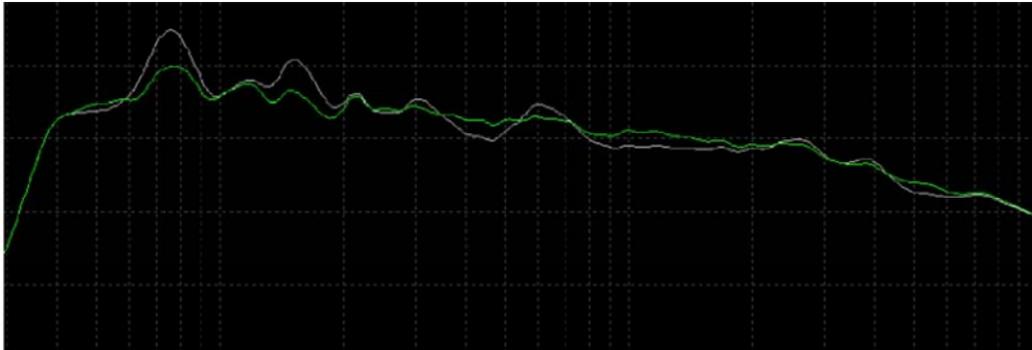


1/12th Octave



Figure 14: Close up of the second segment filtered average spectrum in 1/3rd and 1/12th octave resolution.

1/3rd Octave



1/12th Octave

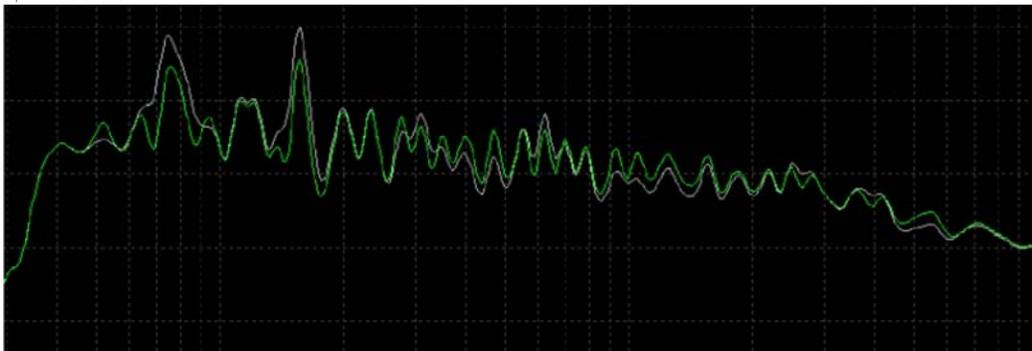


Figure 15: Close up of the third segment filtered average spectrum in 1/3rd and 1/12th octave resolution.

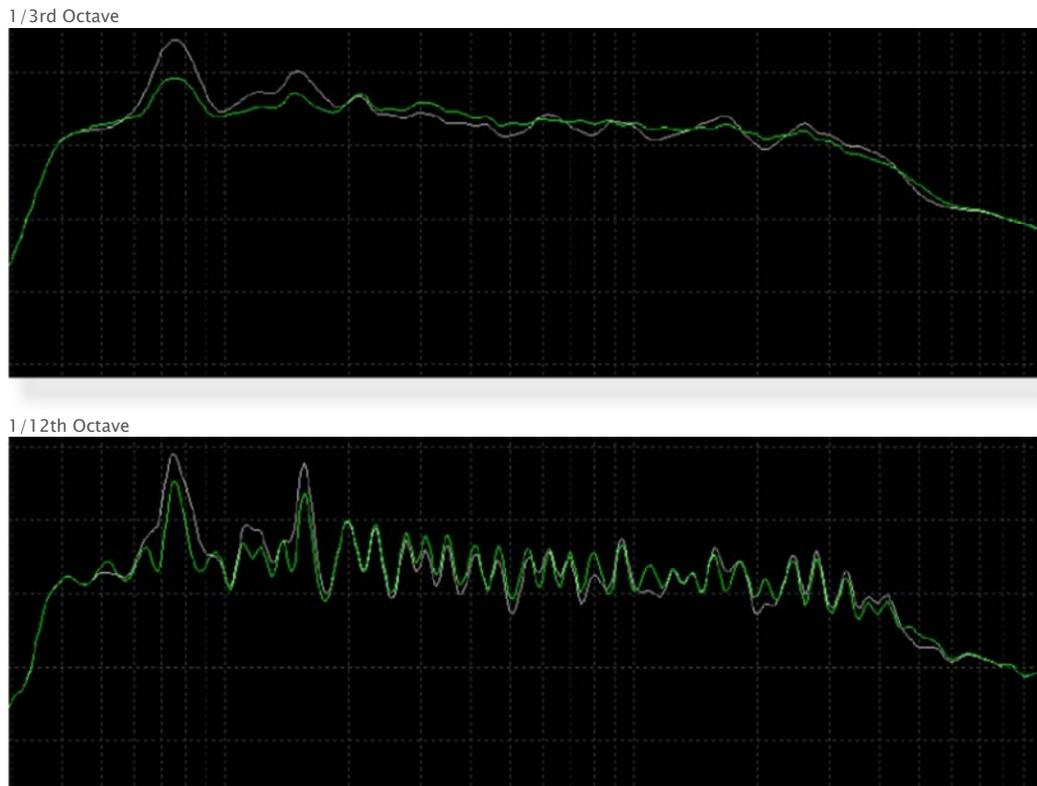


Figure 16: Close up of the fourth segment filtered average spectrum in 1/3rd and 1/12th octave resolution.

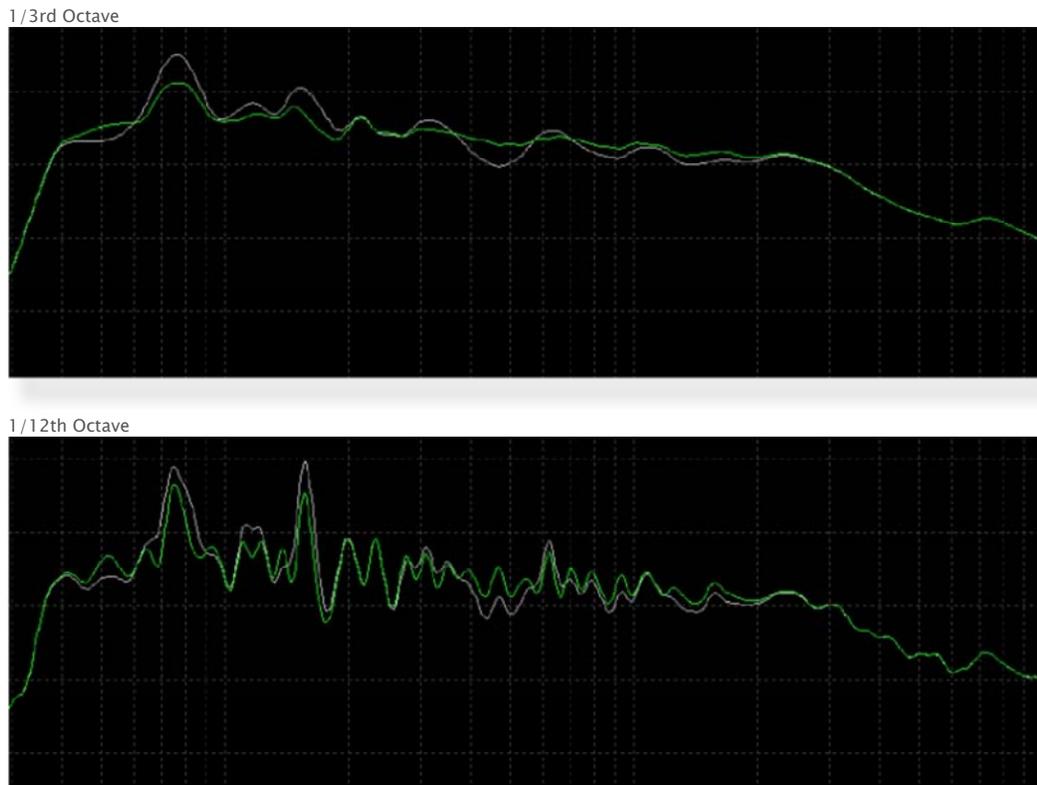


Figure 17: Close up of the fifth segment filtered average spectrum in 1/3rd and 1/12th octave resolution.

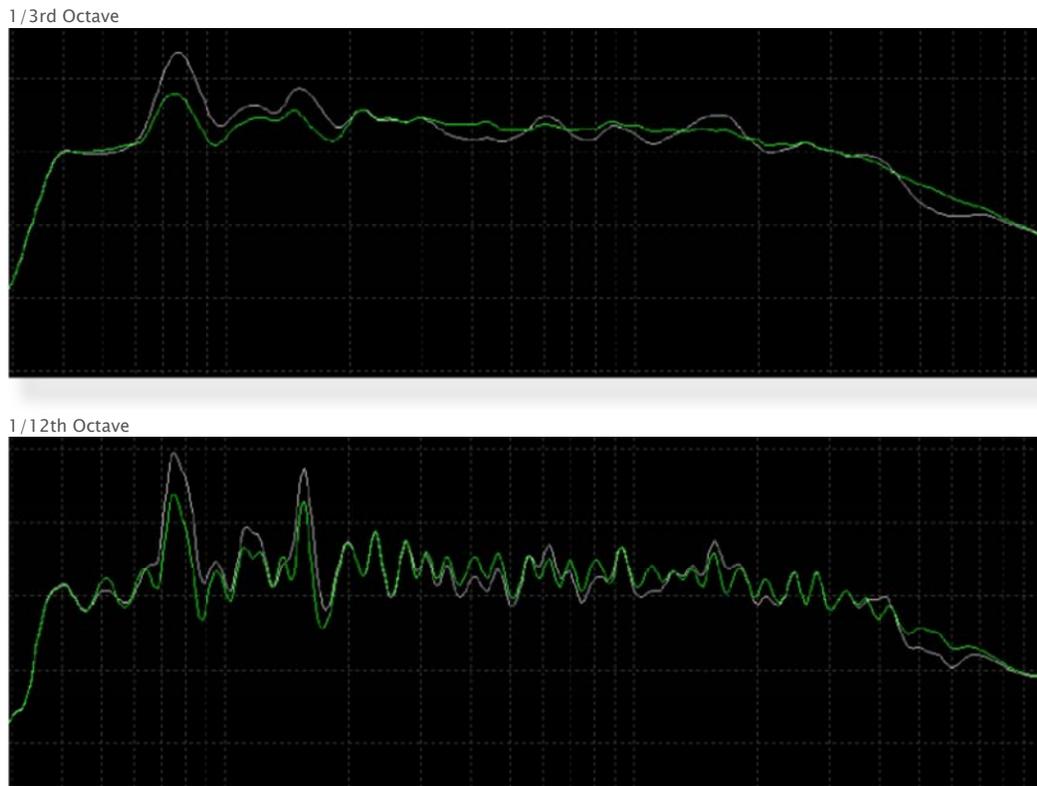


Figure 18: Close up of the sixth segment filtered average spectrum in 1/3rd and 1/12th octave resolution.

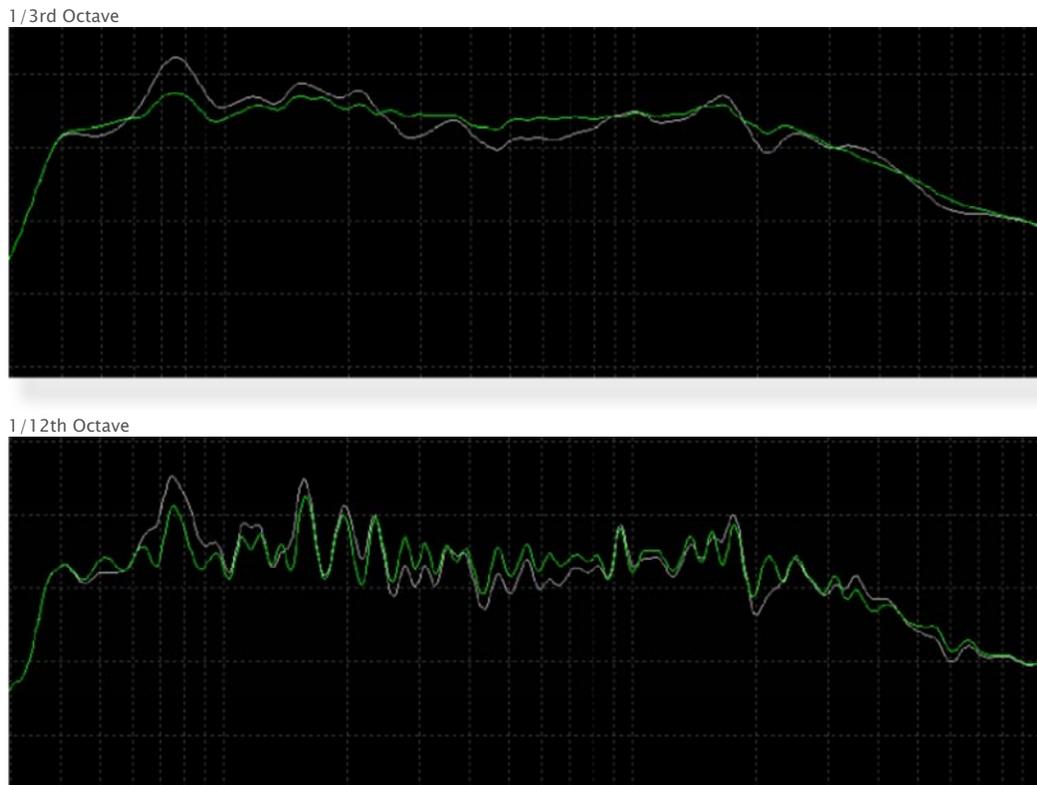


Figure 19: Close up of the seventh segment filtered average spectrum in 1/3rd and 1/12th octave resolution.

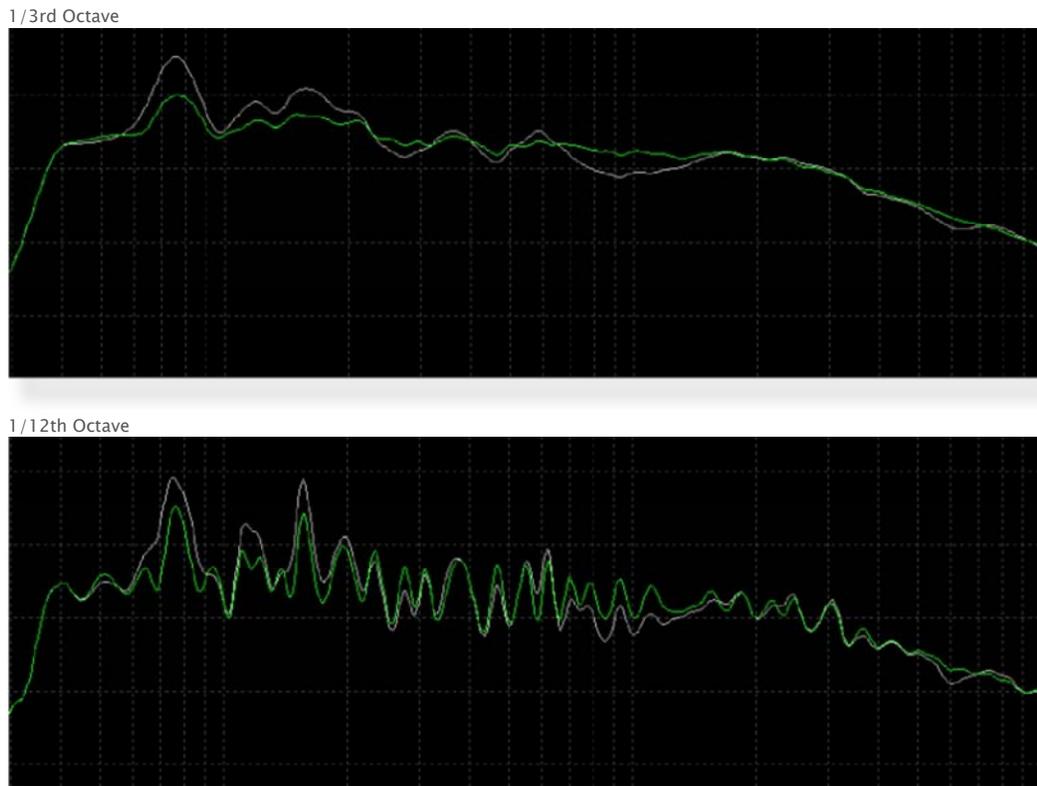


Figure 20: Close up of the eighth segment filtered average spectrum in 1/3rd and 1/12th octave resolution.

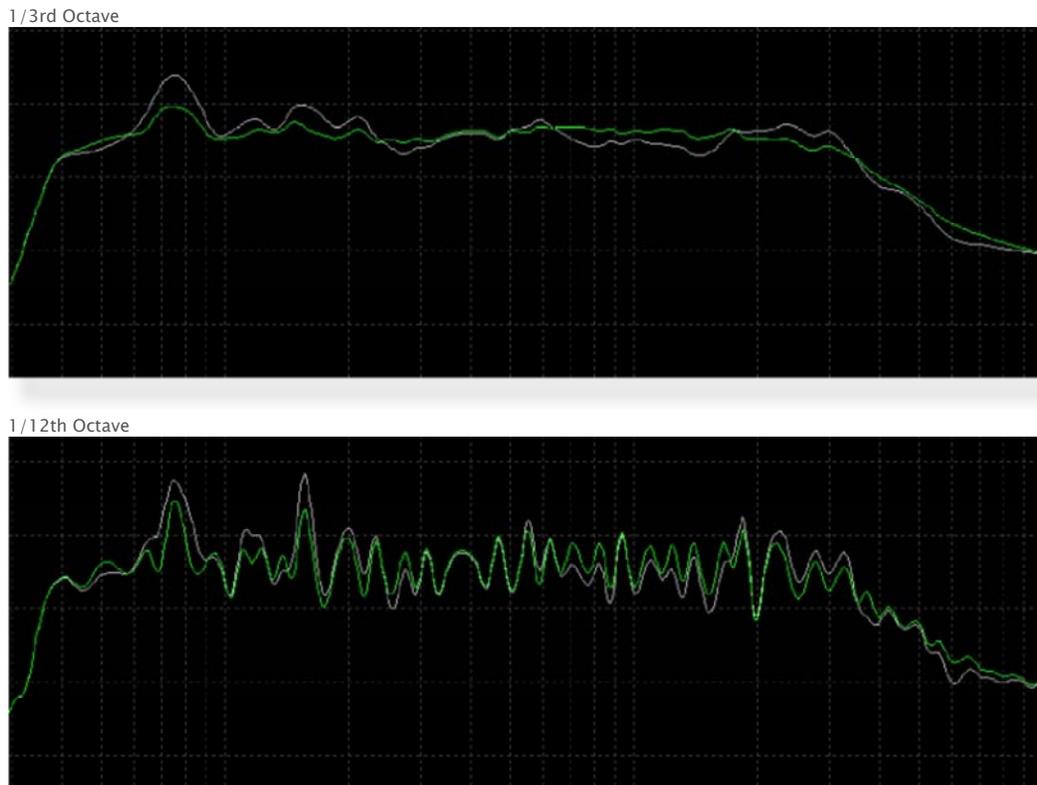


Figure 21: Close up of the ninth segment filtered average spectrum in 1/3rd and 1/12th octave resolution.

Of the ten segments in our track partitioning, the final one is yet to be addressed. Because the general sparseness of the final segment, the average spectrum is quite irregular and this irregularity can be difficult to interpret. A better way to deal with this segment is to base the segment equalisation upon the segment next to it, the ninth segment. But how to we do this?

Segment Referencing

It is typically common that an equalisation we design for one segment could be applied to another wholesale or with additional of minor customisations. Take for example, the filter we design for a chorus and a track with multiple choruses. It is likely that the ideal response for one chorus will be similar to another. Likewise, in our case of wanting a separate filter for the track exit, it is beneficial to be able to base the response on the adjacent segment to avoid having audible tonal discontinuity at the split point.

To be able to create filters for other segments based on an existing one, Har-Bal uses a concept called segment referencing. We select a segment as a reference by pressing down the left mouse button down on the "R" button in the time line and dragging the mouse to the segment selector we want as a segment reference and letting go of the mouse button. Then we select the next segment to work on. Figure 22 illustrates how we would do this using the timeline control.

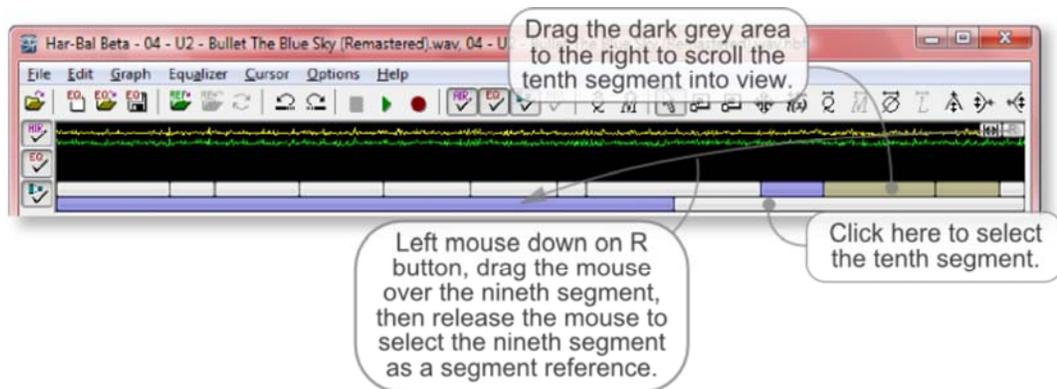


Figure 22: Selecting the ninth segment as segment reference (when the ninth segment is currently selected) and then selecting the tenth segment for editing.

Having selected the ninth segment as segment reference and the tenth segment for editing we can create a filter for the tenth segment based on the segment reference using the **New Filter** menu command. Doing so opens the **New Filter** dialog box. To copy the reference segment response into the selected segment, select the *create...* option "a filter for the selected segment based on the reference segment" as shown in Figure 23.

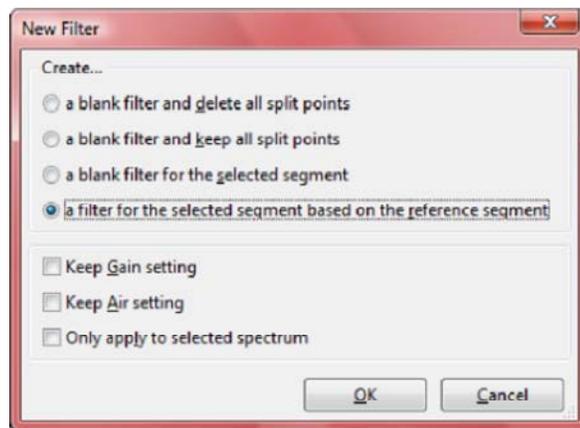
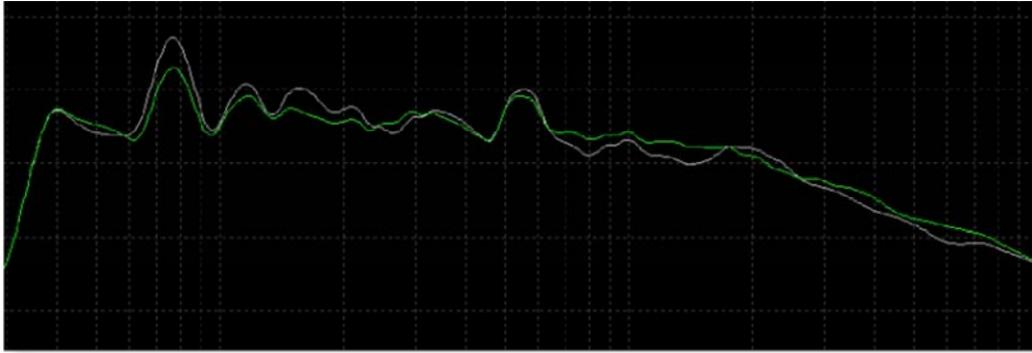


Figure 23: New Filter Dialog Box. Creating a New Filter Based on a Segment Reference.

The copied filter response is further modified to give the spectrums shown in Figure 24. Note that because of the sparsity of the instrumentation in the track exit and not wanting to destroy the character of the instruments, the editing is more conservative than in the other segments.

1/3rd Octave



1/12th Octave

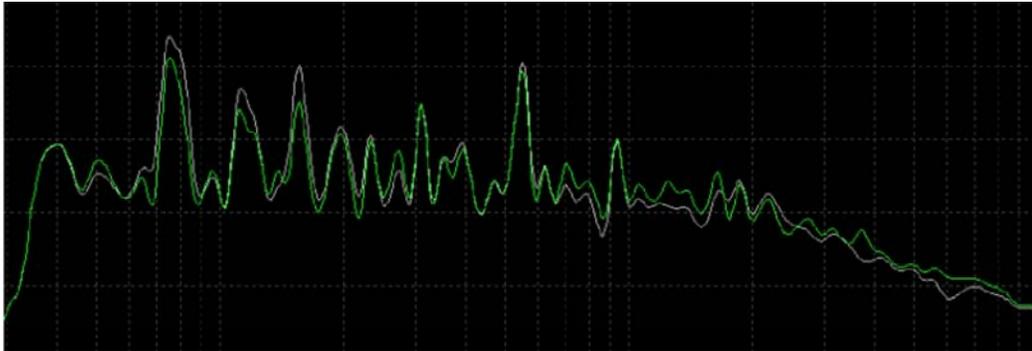


Figure 21: Close up of the tenth segment filtered average spectrum in 1/3rd and 1/12th octave resolution.

On completing the creation of filters using the empathetic equalisation approach applied to the average spectrum, we need to listen to the track for any disturbing transient material. In particular we should listen and watch out for high peak output at selected frequencies that cause distress. If our processing has created some problem areas then we need to address those.

Peak Correction

If on listening you hear some distressing output then it can generally easily be tracked down by looking at the peak spectrum in 1/12th octave view. Any stressful transient material will show as a large peak in the peak spectrum trace. As an example the large peak in segment eight of the filter causes audible stress and needs correction (See Figure 22). We do it using the same empathetic equalisation technique as before but now we apply it to the peak spectrum.

It is useful to have both the peak and average spectrum traces visible while making this type of adjustment because often the peak and average trends are at odds with one another : correcting the average leads to a large peak in the average and correcting the peaks leads to a large dip in the average. In such cases we have to come to a compromise between the peak and average and we base our decision of judicious listening. In future Har-Bal will have a means of decoupling this problem and tackling peaks on their own but at this point in time "*peak taming*" is not implemented.

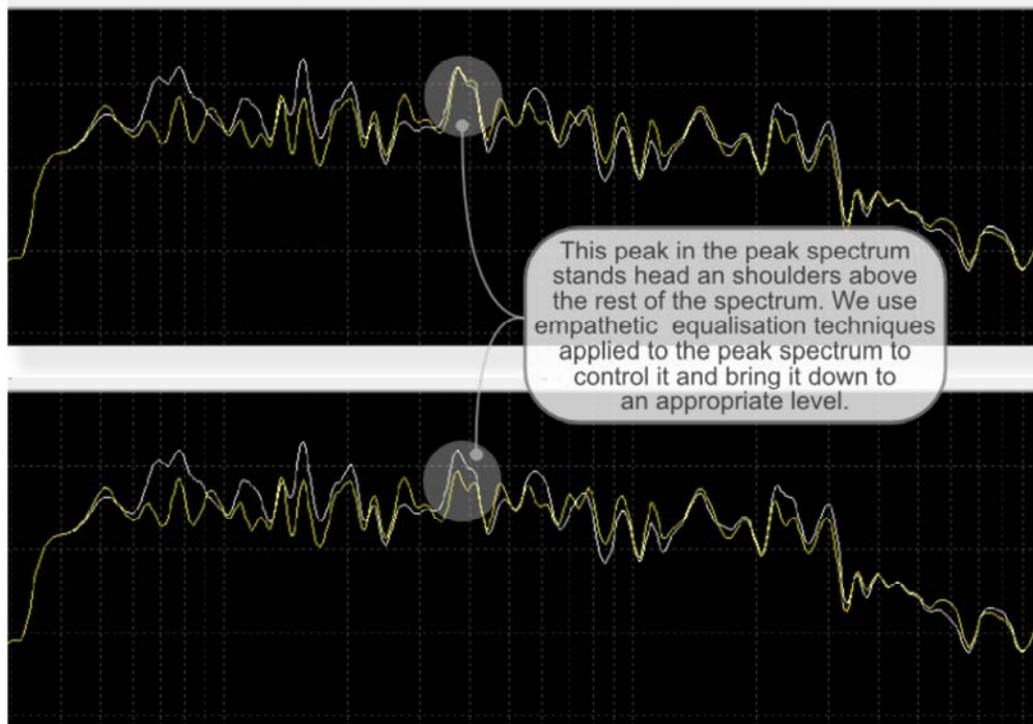


Figure 22: Identify and correct peaks that cause audible distress in the peak spectrum of all segments. Only address a peak if it causes audible distress.

The Final Steps

Having corrected the average and peak spectrums we need to re-visit our overall equalisation and see if it is true to the recording. We do so by careful listening and regularly [toggling equalisation](#) in and out for the entire track. Doing so reveals that whilst the clarity has been improved we have made the bass subdued and mid-range rather dominant in comparison to the original.

Rather than attempt to address this issue by adding extra filter edits to the segment filters there is an easier way. In addition to the segment filters, Har-Bal includes an over-arching filter that combines with the segment filters to create the final outcome. When we need to change the overall balance of the track we use this filter. If you have selected a segment filter but need to now select the over-arching filter all you need do is press the double arrow button in the timeline control as shown in Figure 23.



Figure 23: Press the double arrow button to return to the over-arching filter of a segmented filter.

Having selected the over-arching filter, it can be seen from the 1/3rd octave resolution average spectrum that the bass region does not nominally overlay the original spectrum but largely lies beneath it, which is why it sounds bass deficient compared with the original. However, we do not want to just boost the bass region as it will mask the mid-range so we would like to lift the entire range from about 1.7kHz down to zero so that the weight of the bass end more accurately trends with the original. We can easily do this using the low shelf cursor, as shown in Figure 24. I also introduce a low shelf cut with cutoff frequency of 40Hz to stop the amplification of subsonic content (not shown).

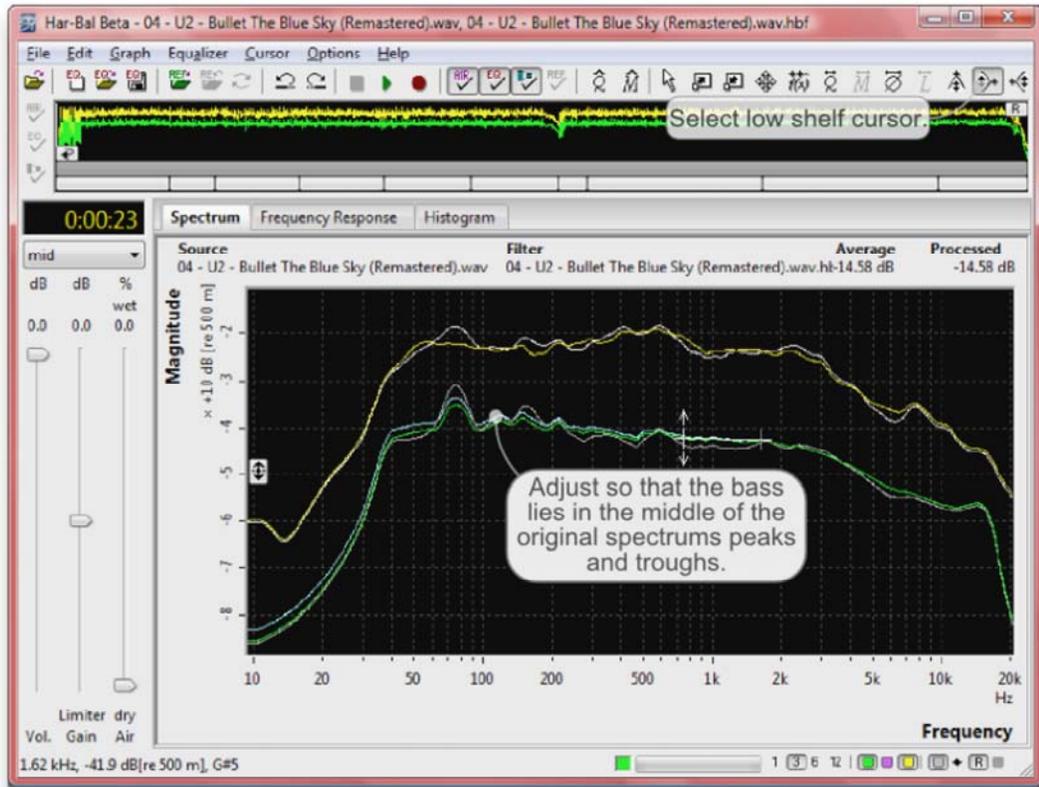


Figure 24: Adjusting the low end balance using the low shelf cursor.

Because of the added weight in the mid range that our re-balancing brings, the top end takes on a subdued sound. To address that issue we use the high shelf cursor with a knee at around 3kHz to bring that area more presence (see Figure 25).

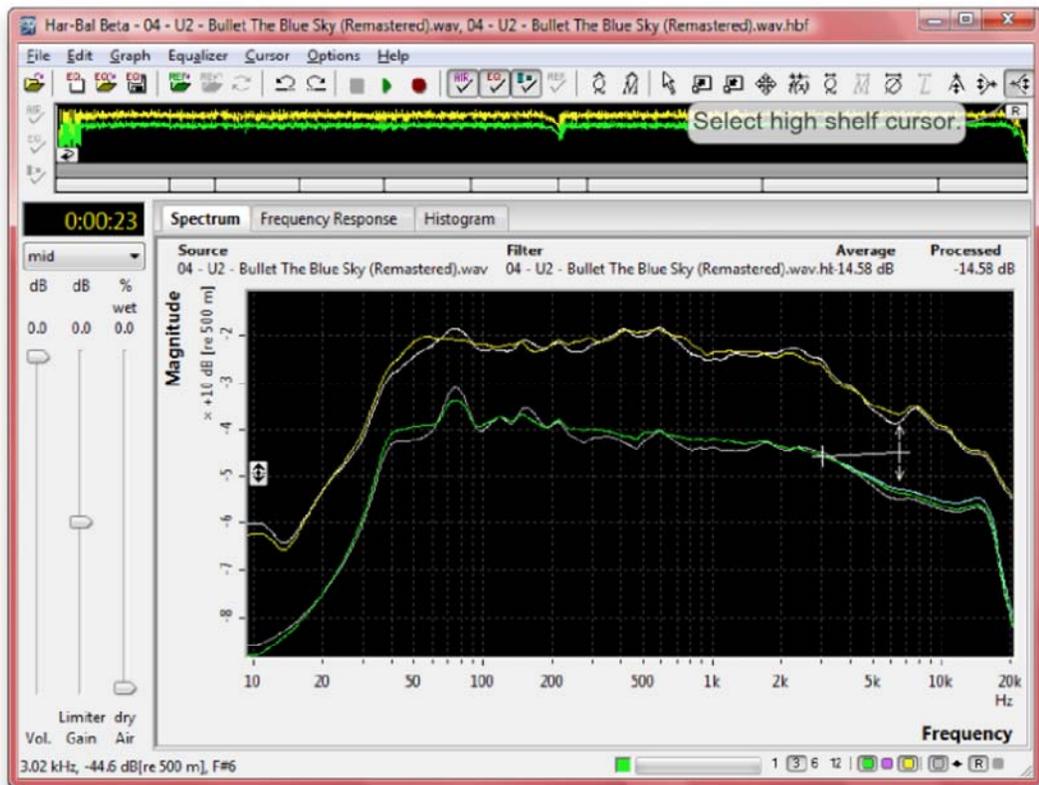


Figure 25: Adjusting the high end balance using the high shelf cursor.

Throughout this tutorial so far we have made all our adjustments to the mid spectrum. This is generally the region you will adjust mostly when building filters in Har-Bal. However, occasionally you will also want to introduce modification in the side spectrum to alter the stereo imaging of the track. This track is a good case in point as the bass in this track is quite widely panned. This does not present a big problem when listening through speakers but with the ubiquity of headphone listening in this day and age, wide panned bass sounds quite disturbing. To improve the sound for monitoring in headphones we can modify the side spectrum to bring this change about. By using the low shelf cursor to introduce a cut from 300Hz down we can reduce the panning of the bass parts, producing near mono bass (see Figure 26). To understand how this works consult the [Har-Bal introduction](#).

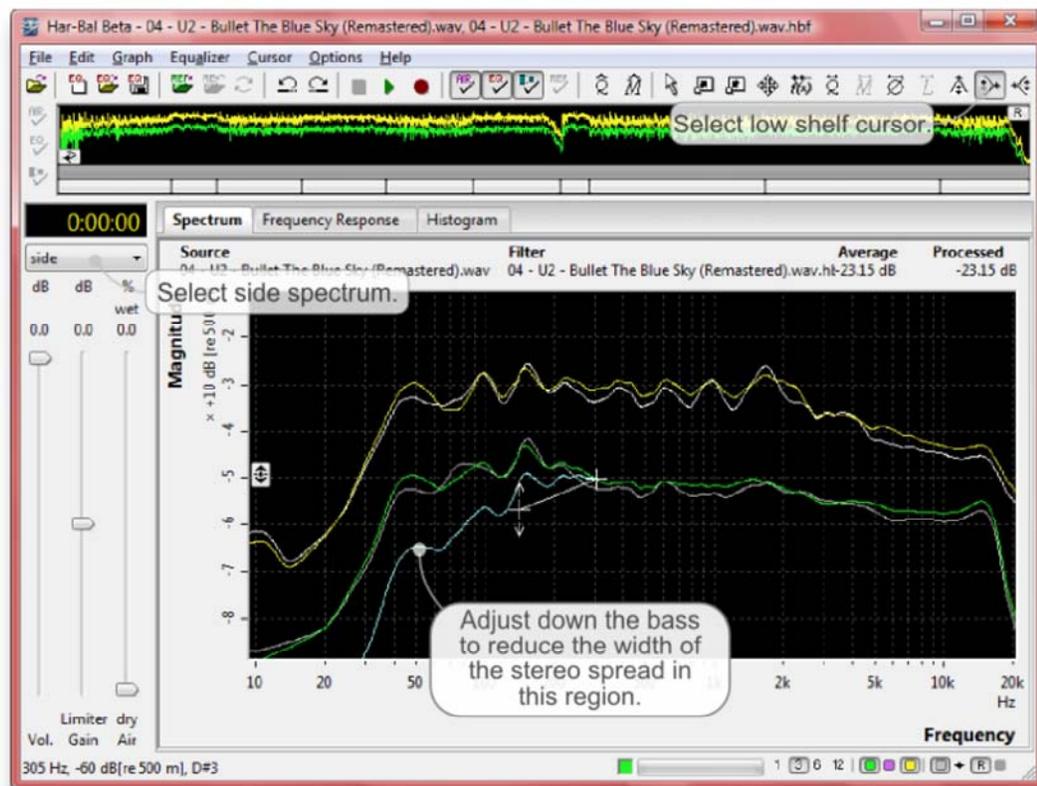


Figure 26: Reducing the bass stereo spread in the side spectrum using the low shelf cursor.

With final listening of the track (particularly when listening through headphones), it is apparent that the track was recorded in a fairly dead recording studio and little ambiance was added during the recording process. This can lead to a sense of detachment from the performance so to add a sense of coherent space to the recording we add 15% air to the recording.

Bullet the Blue Sky : Concluding Remarks

Having completed a comprehensive Har-Bal filter for the U2 track *Bullet the Blue Sky* and hearing the results, it should be patently clear to anyone that harmonic balancing, as applied through Har-Bal, is capable of sonic changes simply not possible via other approaches. The process is complicated but the results are rewarding. The detail required to mimic the process by any other means is simply not possible, owing to the sheer number of parametric equalisation stages required.

Now you might think that I am merely eulogising Har-Bal and without a demonstration of this filter that would certainly be the case. So to bring a taste of reality back into this discussion the filter discussed and developed here is freely available from our website (along with filters to a number of other demonstration tracks). Just download the filter and the track and hear for yourself what Har-Bal can do. I apologise whole heartedly to you for having to pay for the track (we don't have a distribution contract with U2 or the other artists) but for less than the price of a good cup of coffee you can be hearing the revelation for yourself.

The only caveat I'd place on such a test is to ensure that you download the track from the link provided because a lot of material these days is re-mastered multiple times so the only way you can be sure of having the same master as the one the filter was developed for, and hence being able to hear the same results, is if you download from the same location. If you happen to have the same track in your CD collection it is tempting to try that version but it may well be a different master, which will undoubtedly give a biased impression of Har-Bal.

Artist	Track	Filter
Coldplay	Fix You	04 - Coldplay - Fix You.wav.hbf
Paul Dempsey	Fast Friends	02 - Paul Dempsey - Fast Friends.wav.hbf

Foo Fighters	In your Honour	01 – Foo Fighters – In Your Honor.wav.hbf
Foo Fighters	Best Of You	03 – Foo Fighters – Best Of You.wav.hbf
Soundgarden	Room A Thousand Years Wide	08 – Soundgarden – Room A Thousand Years Wide.wav.hbf
U2	With Or Without You (Remastered)	03 – U2 – With Or Without You (Remastered).wav.hbf
U2	Bullet the Blue Sky (Remastered)	04 – U2 – Bullet the Blue Sky (Remastered).wav.hbf

More on the Har-Bal Timeline

Playback Navigation and Looping

Navigating playback to a particular point in the track is as simple as clicking on the desired location in the time line. You can also do this during playback.

The starting location of playback is the selected location in the time line. This location is returned to when playback is stopped unless the *"return to start when play stopped"* option in the Har-Bal [Preferences](#) dialog is not checked. When not checked, playback continues from the last position reached. If the option is checked and the playback position is changed while playing the start position does not change. It only changes if the change is made when not in playback.

Looping is achieved by setting up a looping interval before pressing play. The start of the loop is the starting location for playback and the end of a loop is set by dragging the looping button to the desired location (see Figure 27). A loop can be cleared by dragging the looping position to a time in the time line before the start time. During playback the loop end time can be altered.

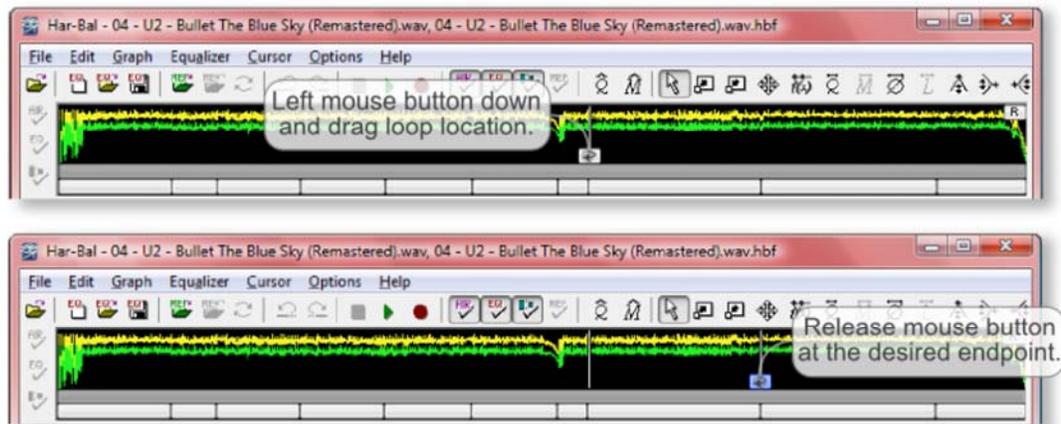


Figure 27: Setting playback loops.

Playable References

With the [Open Reference](#) command you can open a sound file to use as an audible reference which you can [toggle](#) in and out of playback as an aid to building compilations with a consistent tone. Note that for a sound file to be a *"playable"* reference it must have the same sample rate as the session sound file. If they differ the [toggle reference](#) option will be greyed out.

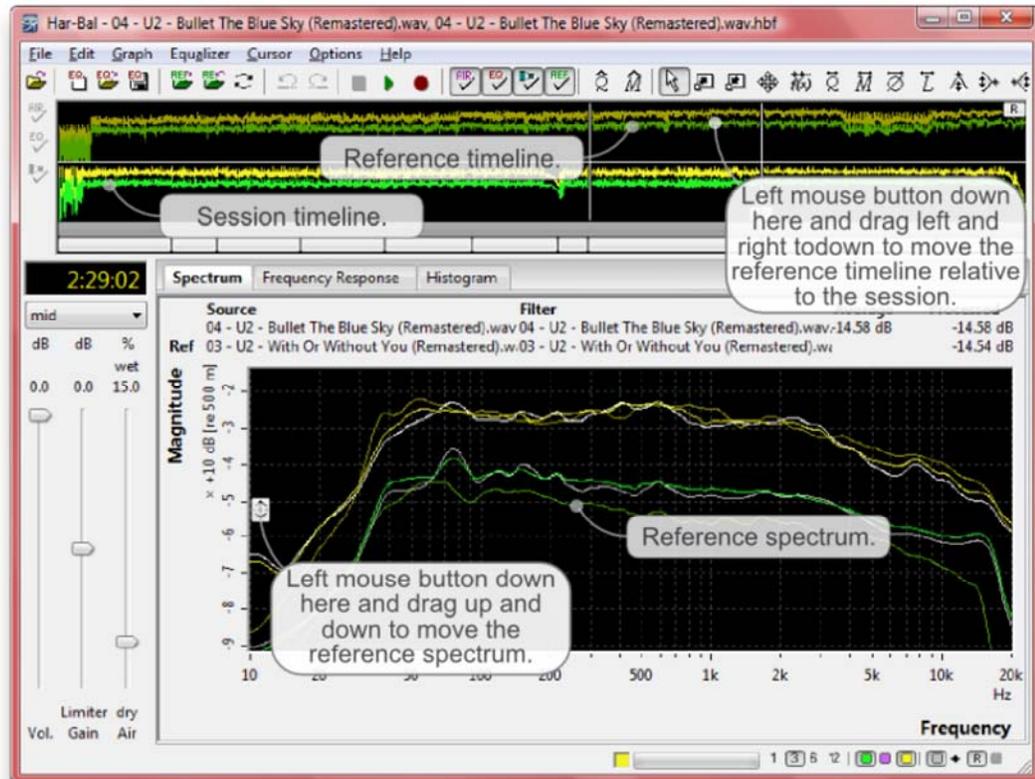


Figure 28: Playable references.

When a playable reference is opened the timeline for the reference is displayed above the timeline for the session and the spectrum for the reference displayed in the spectrum view (see Figure 28). The reference timeline can be dragged horizontally relative to the session timeline so that any part of the reference can be played in sync with any part of the session. Similarly, the reference spectrum can be dragged vertically to align with the session spectrum as an aid to understanding and correcting for differences in tonality.

Creating, Deleting and Moving Split Points

Split points are created by pressing and holding down the left mouse button at the location of the split on the timeline and then pressing Ctrl+S on the keyboard. The best way to determine the split location is by playing the track and clicking on the start location. When you release the left mouse button playback will resume from that location and through listening we can easily hear if the location captures the content we want in a segment. If it doesn't then we reposition and try again. When we have found the correct location we press and hold in the left mouse button and press Ctrl+S on the keyboard to create the split.

If the split location needs adjusting we can do so by pressing the left mouse button down while over the split marker and waiting for the cursor to change to horizontal a resize cursor. When it has we can then drag the split point left or right of its current position and release the left mouse button to establish the new position.

To delete a split point we click on the split marker to select it. Then once selected simply press the delete key on the keyboard to delete the split point. These different scenarios are illustrated in Figure 29.

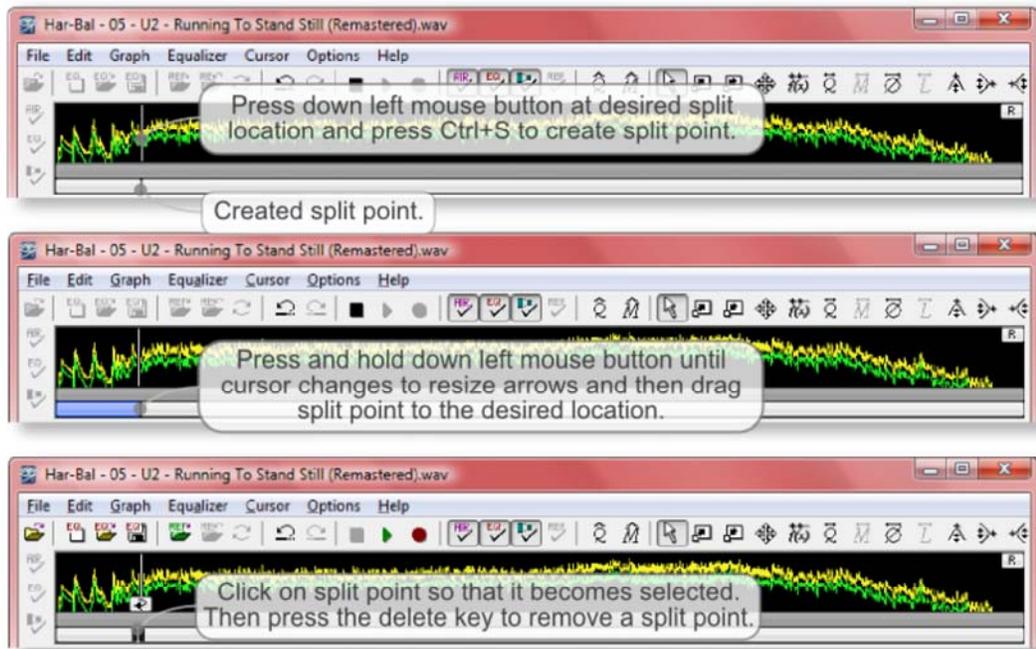


Figure 29: Creating, Deleting and Moving Split Points.

Selecting and Navigating Segments

When in full timeline view you select a segment by mousing over the desired segment in the top bar and left mouse button clicking (see Figure 30). When a segment is already selected you do not need to return to full timeline view to select a new segment. An easier way is to drag the selected range in the top bar so that it overlays the segment you would like to select and then click on the segment in the bottom bar to change the selection to it.

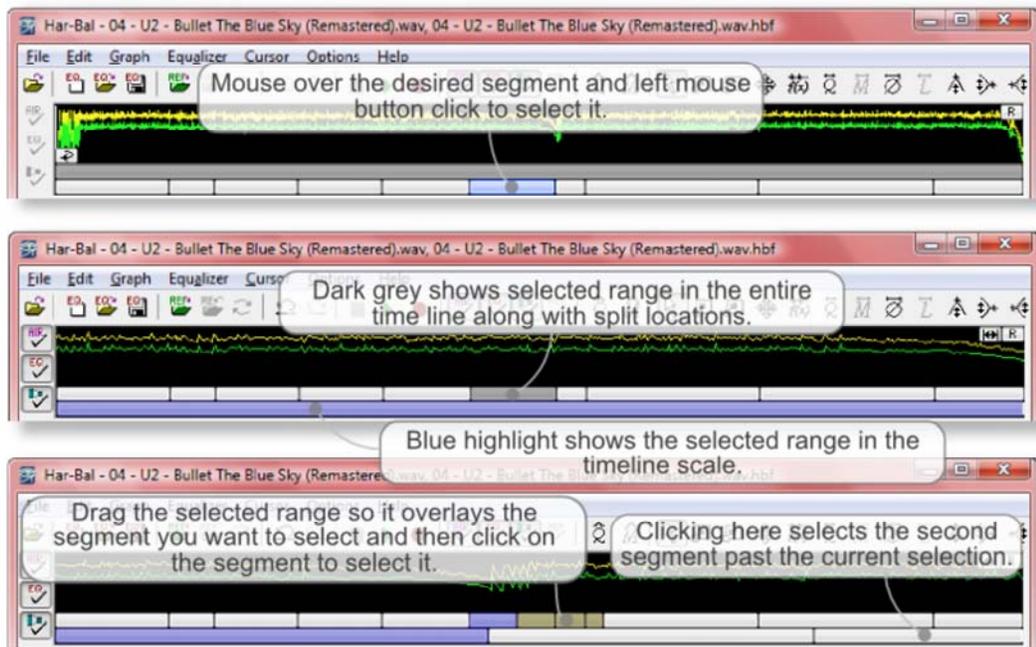


Figure 30: Navigating Segments.

Concluding Remarks

This tutorial is by no means extensive. For the sake of getting you up and running we have deliberately left it as a shortish introduction of what we think you should know. There is much more that we could discuss but for the sake of brevity we have left that untouched.

If you haven't already done so you can find extra details on the workings of Har-Bal and some of the new functionality in the [introduction](#), and a discussion of the concept of loudness compensation [here](#). For the features untouched in this tutorial, you can easily find help using the help system in Har-Bal. To use the help system consult the instructions found [here](#). Finally, the [Har-Bal forum](#) has a wealth of information on it regarding much of what has been left out of this tutorial. I encourage you to make use of it.



Har-Bal Harmonic Balancer

Har-Bal™ new thinking, new directions...

Menu Commands

The top level menus provided by Har-Bal Harmonic Balancer include:

- [File Menu](#)
- [Edit Menu](#)
- [Graph Menu](#)
- [Equalizer Menu](#)
- [Cursor Menu](#)
- [Options Menu](#)
- [Help Menu](#)

File Menu

- [Open Command](#)
- [Close Command](#)
- [Record As Command](#)
- [Export Reference Command](#)
- [Batch Analysis Command](#)
- [Manage Files Command](#)
- [Filter Submenu](#)
 - [New Command](#)
 - [Open Command](#)
 - [Save Command](#)
 - [Save As Command](#)
- [Reference Submenu](#)
 - [Open Command](#)
 - [Close Command](#)
 - [Swap Command](#)
- [Recent Sessions Command](#)
- [Recent References Command](#)
- [Print Command](#)
- [Print Setup Command](#)
- [Exit Command](#)

File | Open Command

 shortcut: Ctrl+O

Use the Open command to select and open a file to process in a Har-Bal session. The source file formats supported by Har-Bal include:

- Wave files (.wav)
- AIFF files (.aif)
- flac files (.flac,.fla)
- ogg files (.ogg)
- NeXT/Sun files (.au)
- Amiga IFF/8SVX files (.iff,.svx)
- Sphere NIST files (.nist)
- Raw PCM Data files (.pcm)
- Paris files (.paf)
- mu-law files (.ulaw)
- A-law files (.alaw)
- A-law files (.alaw)
- mp3 files (.mp3)

Note that support for mp3 format is provided by using the built in CODEC in your operating system. On opening an mp3 format Har-Bal will first convert the file to a .wav or .aif format file and then open the decoded file. Also note that VBR mp3 file encoding is not supported in Har-Bal on the Windows platform.

On opening a source file for the first time Har-Bal will analyze the file contents and save the analysis in an analysis file stored in the same folder as the source file. The analysis file has the same name as the source file but with a .hba file extension.

You should only open files in a folder that resides on the computer hard drive and the folder must be writable by Har-Bal. Har-Bal requires high speed disk access for the real-time spectrum display and playback to function correctly.

File | Close Command

The close command closes the currently open session.

File | Record As Command

The File | Record As command renders the processed track to disk to an output file of your choosing. By default the output file has the same format as the source file and the file name as the name of the source file with **_eq** appended. For example, if the source file is **my_track.wav** then the default output file is **my_track_eq.wav**. You can change the file name to one of your own choosing and you can change the output file format by selecting a different **Save As type** in the Save as dialog box.

Recording files occurs in the background to allow you to continue working on another track in another Har-Bal session. The status of outstanding recording jobs is displayed in the [Recording](#) dialog. This dialog remains visible while recording jobs are queued and closes automatically when all jobs are completed. You can minimise this dialog and continue doing other work with Har-Bal at any time.

File | Export Reference Command

The File | Export Reference command allows you to create reference analyses from landmark recordings. The reference analysis contains a snapshot of the spectrums and histograms but no time dependent information (real time spectrum). Reference files are given the extension **.hbref** and can be opened as a session reference with the [Open reference](#) command.

File | Batch Analysis Command

The File | Batch Analysis command allows you to analyze any number of source files in the background, thus freeing up Har-Bal to do other work at the same time. Selecting this command opens an open file common dialog that allows multiple file selection. In the file list view select the sound files that you wish to analyze. After pressing the OK button the selected files are then queued to be analyzed and the status of the analysis jobs displayed in the [Analyzing](#) dialog. This dialog remains visible while analysis jobs are queued and closes automatically when all jobs are completed. You can minimise this dialog and continue doing other work with Har-Bal at any time.

File | Manage Files Command

The File | Manage Files command opens the [Manage Files](#) dialog box. Use this dialog box to manage and delete the working files that Har-Bal creates and uses to process your tracks. Har-Bal will delete **.hba**, **.hbfb**, **.hbfbfs**, **.hbfr**, **.hbfs**, **.hbir** files but not **.hbf** files. This is by design as the **.hbf** file contains all the information required to re-create the processing designed for a given session and is something you are likely to want to archive for future reference. If you want to delete superfluous **.hbf** files then use Explorer (on Windows) or Finder (on Mac OS X) to do so.

File | Filter | New Command



Use the Filter | New command to create a new filter. This command opens the [New Filter](#) dialog box. On opening a session file a new filter is created with a flat frequency response on all channels.

File | Filter | Open Command



shortcut: Ctrl+F

Use the Filter | Open command to open and use a previously saved session filter.

File | Filter | Save Command

Use the Filter | Save command to save the session filter to file. The filter may then be re-loaded at a later time to reproduce the processing applied to a given source file.

File | Filter | Save As Command

Like the Filter | Save command, the Filter | Save As command allows you to save the session filter to a specific file.

File | Reference | Open Command



shortcut: Ctrl+R

Use the Reference | Open command to open a given source file or reference file for use as a reference. The analysis of the reference source will be overlayed on the analysis plots of session. However, unlike the session spectrum, the reference spectrum is not altered by the filter. If the reference is playable then the [toggle reference](#) feature will be enabled. Note that for a reference to be playable it must be a sound file **AND** must have the same sampling rate as the session file.

File | Reference | Close Command



Copyright © 2011
Har-Bal™. All rights
reserved.



Use the Reference | Close command to discard the reference file. The reference spectrum will be removed from the spectrum display.

File | Reference | Swap Command



The Reference | Swap command allows you to swap the reference file and filter with the session file and filter. This command is only enabled if the reference file is playable. This feature is particularly useful when performing tonality and loudness balancing between two tracks when you need to make filter adjustments for both the source and the reference. Note that for a reference to be playable it must be a sound file **AND** must have the same sampling rate as the session file.

File | Recent Sessions Command

The File | Recent Sessions command lists the 10 most recently used source (and corresponding filter) files as a popup menu. The name of the files appear in the popup menu items.

File | Recent References Command

The File | Recent References command lists the 10 most recently used reference (and corresponding filter) files as a popup menu. The name of the files appear in the popup menu items.

File | Print Command

Use the File | Print command to print any of the analysis views to printer. To select the printer use the Print Setup command.

File | Print Setup Command

Use the File | Print Setup command to select the printer to use when printing analysis views and to configure the printer.

File | Exit Command

shortcut: Alt+F4

Use the File | Exit command to close Har-Bal Harmonic Balancer.

Edit Menu

- [Undo Command](#)
- [Redo Command](#)

Edit | Undo Command



shortcut: Ctrl+Z

Use the Edit | Undo command to undo changes made to the filter. Each change made is recorded in an unlimited Undo stack. View changes are recorded in the stack to keep track of editing context so that you don't inadvertently undo something unintended and so that you can see what is being undone.

Edit | Redo Command



shortcut: Ctrl+Y

Use the Edit | Redo command to redo undone changes made to the filter. Each change made is recorded in an unlimited Undo stack. View changes are recorded in the stack to keep track of editing context so that you don't inadvertently undo something unintended and so that you can see what is being undone.

Graph Menu

- View Submenu
 - [Spectrum Command](#)
 - [Frequency Response Command](#)
 - [Histogram Command](#)
- Show Spectrum Submenu
 - [Peak Command](#)
 - [Real time Command](#)
 - [Average Command](#)
 - [Original Command](#)
 - [Playback Focus Command](#)
 - [Reference Command](#)
 - [Segment Reference Command](#)
 - Spectrum Resolution Submenu
 - [1 Octave Command](#)

- 1/3 Octave Command
- 1/6 Octave Command
- 1/12 Octave Command
- Zoom all Command
- Clipboard copy Command
- Properties Command

Graph | View | Spectrum Command

shortcut: Alt+1

Use the Graph | View | Spectrum command to display the session and reference spectrums. This view shows the session spectrum, the effect of the filter on the spectrum and the reference spectrum. You can edit the spectrum using the Gain, Low Shelving and High Shelving cursors.

Graph | View | Frequency Response Command

shortcut: Alt+2

Use the Graph | View | Frequency Response command to display the frequency response of the selected filter. You can edit the frequency response using the Gain, Low Shelving and High Shelving cursors.

Graph | View | Histogram Command

shortcut: Alt+3

Use the Graph | View | Histogram command to display the histogram of session levels. You can edit the histogram using the Dynamics Node cursor.

Graph | Show Spectrum | Peak Command

 shortcut: Ctrl+1

Use the Graph | Show Spectrum | Peak command to toggle the visibility of the peak power plot on the spectrum and histogram views. If checked then the peak power is visible. You can use this command to temporarily hide the peak power plot.

Graph | Show Spectrum | Real time Command

 shortcut: Ctrl+2

Use the Graph | Show Spectrum | Real time command to toggle the visibility of the real time power plot on the spectrum view. If checked then the real time power plot is visible. You can use this command to temporarily hide the real time power plot.

Graph | Show Spectrum | Average Command

 shortcut: Ctrl+3

Use the Graph | Show Spectrum | Average command to toggle the visibility of the average power plot on the spectrum and histogram views. If checked then the average power plot is visible. You can use this command to temporarily hide the average power plot.

Graph | Show Spectrum | Original Command

 shortcut: Ctrl+4

Use the Graph | Show Spectrum | Original command to toggle the visibility of the power plots corresponding to the original unprocessed track on the spectrum and histogram views. If checked then the original power plots are visible. You can use this command to temporarily hide the original power plots. Alternatively, showing these plots immediately shows you how the spectrum and histogram has been changed by overlaying the processed plots on top of the originals.

Graph | Show Spectrum | Playback Focus Command

 shortcut: Ctrl+5

Use the Graph | Show Spectrum | Playback Focus command to toggle the visibility of the focus nodes on the spectrum and histogram views. If checked then the focus nodes are visible and are displayed at octave intervals on the plots that have focus. That is the plots corresponding to the spectrum and histogram corresponding to what you hear during playback.

Graph | Show Spectrum | Reference Command

 shortcut: Ctrl+6

Use the Graph | Show Spectrum | Reference command to toggle the visibility of the power plots of the reference track on the spectrum and histogram views. If checked then the reference power plots are visible. You can use this command to temporarily hide the reference power plots.

Graph | Show Spectrum | Segment Reference Command

 shortcut: Ctrl+7

Use the Graph | Show Spectrum | Segment Reference command to toggle the visibility of the power plots of the segment reference on the spectrum and histogram views. If checked then the segment reference power plots are visible. You can use this command to temporarily hide the segment reference power plots. This command is only active if a segment reference is selected.

Graph | Show Spectrum | Spectrum Resolution | 1 Octave Command

 shortcut: Ctrl+Shift+1

Use the Graph | Show Spectrum | Spectrum Resolution | 1 Octave command to show the spectrum analysis with a 1 Octave band resolution. This command is used to visualise how the spectrum trend changes with resolution. When displayed with high resolution it becomes harder to visualise the trend but specific issues become more apparent.

Graph | Show Spectrum | Spectrum Resolution | 1/3 Octave Command

 shortcut: Ctrl+Shift+2

Use the Graph | Show Spectrum | Spectrum Resolution | 1/3 Octave command to show the spectrum analysis with a 1/3 Octave band resolution. This command is used to visualise how the spectrum trend changes with resolution. When displayed with high resolution it becomes harder to visualise the trend but specific issues become more apparent.

Graph | Show Spectrum | Spectrum Resolution | 1/6 Octave Command

 shortcut: Ctrl+Shift+3

Use the Graph | Show Spectrum | Spectrum Resolution | 1/6 Octave command to show the spectrum analysis with a 1/6 Octave band resolution. This command is used to visualise how the spectrum trend changes with resolution. When displayed with high resolution it becomes harder to visualise the trend but specific issues become more apparent.

Graph | Show Spectrum | Spectrum Resolution | 1/12 Octave Command

 shortcut: Ctrl+Shift+4

Use the Graph | Show Spectrum | Spectrum Resolution | 1/12 Octave command to show the spectrum analysis with a 1/12 Octave band resolution. This command is used to visualise how the spectrum trend changes with resolution. When displayed with high resolution it becomes harder to visualise the trend but specific issues become more apparent. The bin centres in 12th octave view is aligned with the western equal tempered scale.

Graph | Zoom all Command

shortcut: Alt+Z

Use the Graph | Zoom All command to zoom the spectrum or histogram graph into full view so that all the graph features are visible.

Graph | Clipboard copy Command

Use the Graph | Clipboard Copy command allows to place a captured image of the spectrum into the clipboard, which can then paste into another application that accepts images.

Graph | Properties Command

Use the Graph | Properties command to adjust the range (zoom factor) for the spectrum and histogram views using the keyboard rather than the mouse. This command opens the [Graph Properties](#) dialog box.

Equalizer Menu

- [Stop Command](#)
- [Play Command](#)
- [Record As Command](#)
- [HB Air In Command](#)
- [EQ In Command](#)
- [Dynamics In Command](#)
- [Reference In Command](#)
- [Apply intuitQ Command](#)
- [Apply intuitMatch Command](#)
- [Match Loudness Command](#)

Equalizer | Stop Command

 toggle shortcut: spacebar

Use the Equalizer | Stop command to stop playback of the session.

Equalizer | Play Command

 toggle shortcut: spacebar

Use the Equalizer | Play command to play the session.

If the [Air In](#) menu is checked or the AIR toolbar button is depressed then playback will include the air ambiance processing. You can toggle the Air In menu during playback to hear the effect of Air on the source material.

If the [EQ In](#) menu is checked or the EQ toolbar button is depressed then playback will include the equalization filter. You can toggle the EQ In menu during playback to hear the effect of EQ on the source material.

If the [Dynamics In](#) menu is checked or the dynamics toolbar button is depressed then playback will include the dynamics processing. You can toggle the Dynamics In menu during playback to hear the effect of dynamics procession on the source material.

If a reference file is loaded and it is playable you can toggle playback between the session file and the reference file using the [Reference In](#) menu command or REF In toolbar button. Playback of the reference will include equalization filtering if a filter file was specified when the reference was opened. Pressing any of the processing related buttons mentioned above has no effect on the reference playback. This feature is particularly useful for normalising songs to have consistent tonality and loudness.

Equalizer | Record As Command

 shortcut: Ctrl+D

The Equalizer | Record As command renders the processed track to disk to an output file of your choosing and includes any processing that is enabled at the time of invoking this command. By default the output file has the same format as the source file and the file name as the name of the source file with **_eq** appended. For example, if the source file is **my_track.wav** then the default output file is **my_track_eq.wav**. You can change the file name to one of your own choosing and you can change the output file format by selecting a different **Save As type** in the Save as dialog box.

Recording files occurs in the background to allow you to continue working on another track in another Har-Bal session. The status of outstanding recording jobs is displayed in the [Recording](#) dialog. This dialog remains visible while recording jobs are queued and closes automatically when all jobs are completed. You can minimise this dialog and continue doing other work with Har-Bal at any time.

Equalizer | HB Air In Command

 shortcut: Ctrl+A

The Equalizer | HB Air In command switches the air ambiance processing in and out of playback and record operations in real time so that you can evaluate its effect on the source material. Use air to add a sense of space to stereo recordings that lack a coherent space.

Air uses cross coupled FIR filters to realise a Haas zone ambiance effect. Moderate use of air will not colour the track tonality. Heavy use of air will have some effect on colour though the degree of colouration is small owing to the ambiance being in the Haas zone.

You can use air processing to stereo'ise monophonic material. In this application heavy use of air can be useful.

Note that the Air effect is most noticeable in stereo reproduction through loudspeakers in a reasonably dry listening environment. The more diffuse ambience present in the listening environment the less audible the air effect is. Headphones are an idea way of judging the air processing.

Equalizer | EQ In Command

 shortcut: Ctrl+Q

The Equalizer | EQ In command switches the equalization filter in and out of playback and record operations in real time so that you can evaluate the effect of the filter on the session.

Equalizer | Dynamics In Command

 shortcut: Ctrl+C

The Equalizer | Dynamics In command switches the dynamics processing in and out of playback and record operations in real time so that you can evaluate the effect of dynamics processing on the session.

Equalizer | Reference In Command

 shortcut: Ctrl+W

The Equalizer | Reference In command switches playback between the session file and the reference file in real time so that you can compare the tonality and loudness of the session with the reference.

Equalizer | Apply intuitQ Command

 shortcut: Ctrl+I

The Equalizer | Apply intuitQ command uses a smoothing algorithm to loosely smooth the average spectrum shape of the session file.

Equalizer | Apply intuitMatch Command

 shortcut: Ctrl+Alt+I

The Equalizer | Apply intuitMatch command uses a fitting algorithm to loosely match the session spectrum shape to that of the loaded reference.

Equalizer | Match Loudness Command

shortcut: Ctrl+M

The Equalizer | Match Loudness command matches the average loudness figure of merit of the session track with the current reference. Loudness matching using this approach is generally only accurate when the style of the music in the reference is comparable to the session track. The most accurate means of matching loudness is through A/B comparison and manual gain adjustment.

Cursor Menu

- [Disabled Command](#)
- [Zoom In Command](#)
- [Zoom Out Command](#)
- [Pan Command](#)
- [Harmonic Command](#)
- [IntuitQ Command](#)
- [IntuitMatch Command](#)
- [IntuitNull Command](#)
- [Match Loudness Command](#)
- [Gain Command](#)
- [Low Shelving Command](#)
- [High Shelving Command](#)
- [Dynamics Node Command](#)

Cursor | Disabled Command



Use the Cursor | Disabled command to disable editing of the session filter. If you have completed designing a filter for the session track, select this cursor to avoid accidentally altering the design through un-intended mouse clicks.

Cursor | Zoom In Command



Use the Cursor | Zoom In command to zoom into a specific feature in the spectrum, frequency response or histogram views. Once selected the cursor icon changes to the zoom in cursor. Zooming can be done in two different ways : by clicking or by clicking and dragging a zoom frame. If you click on a specific region then the view is zoomed in by 50% centred upon the point clicked. If clicking and dragging a zoom box, the view is zoomed into the region bounded by the box. During this cursor mode the axis values at the cursor position are displayed in the status bar.

Cursor | Zoom Out Command



Use the Cursor | Zoom Out command to zoom out of a specific feature in the spectrum, frequency response or histogram views. Once selected the cursor icon changes to the zoom out cursor. Zooming can be done in two different ways : by clicking or by clicking and dragging a zoom frame. If you click on a specific region then the view is zoomed out by 50% centred upon the point clicked. If clicking and dragging a zoom box, the view is zoomed out of the region bounded by the box such that the corners of the view are mapped to the corners of the zoom frame. During this cursor mode the axis values at the cursor position are displayed in the status bar.

Cursor | Pan Command



Use the Cursor | Pan command to pan the spectrum, frequency response or histogram view in any direction. Once selected the cursor icon changes to the pan cursor. You can then pan the view by clicking and dragging the pan arrow in the desired direction. When you release the left hand mouse button the point where the mouse button was pushed down is moved to the

point where the mouse button was released. During this cursor mode the axis values at the current cursor position are displayed in the status bar.

Cursor | Harmonic Command



Use the Cursor | Harmonic command to uncover the harmonic relationships of the peaks in the track spectrum and their values. Pressing and holding in the left mouse button will highlight the first ten harmonics with cross-hairs. Moving the mouse left or right while the left mouse button is down will move the cross-hairs along the spectrum line. The status bar shows the frequency and magnitude of the fundamental (the left most cross-hair). The frequency value is displayed in both Hz and the closest equal temperament scale note. You can display the values of the other harmonics in the status bar by pressing and holding in one of the number keys: the harmonic displayed corresponding to the number pressed. That is, to display the value of the second harmonic press and hold in the '2' key. The tenth harmonic is displayed by pressing the '0' key.

Cursor | IntuitQ Command



Use the Cursor | IntuitQ command to apply the intuitQ spectrum smoothing algorithm to a selected frequency range of the session spectrum. To do so, move the cursor to the starting point of the frequency range, click and hold down the left mouse button and then drag the cursor to the end point of the desired frequency range. When you release the left mouse button intuitQ will be applied to the selected frequency range.

Cursor | IntuitMatch Command



Use the Cursor | IntuitMatch to apply the intuitMatch spectrum matching algorithm to a selected frequency range of the session spectrum. To do so, move the cursor to the starting point of the frequency range, click and hold down the left mouse button and then drag the cursor to the end point of the desired frequency range. When you release the left mouse button intuitMatch will be applied to the selected frequency range. IntuitMatch is only available if a reference is loaded.

Cursor | IntuitNull Command



Use the Cursor | IntuitNull command to loosely undo the effect of intuitQ or intuitMatch over a given frequency range. To do so, move the cursor to the starting point of the frequency range, click and hold down the left mouse button and then drag the cursor to the end point of the desired frequency range. When you release the left mouse button intuitNull will be applied to the selected frequency range.

IntuitNull can be thought of as a 'loose' undo operation for IntuitQ or IntuitMatch. For many tracks you will want to apply intuitQ or IntuitMatch to the full spectrum range of the track but some areas in that range may become over-emphasised (typically when the track has few instrument parts). By selecting the overemphasised parts of the spectrum you can undo the emphasis that intuitQ created whilst keeping the changes in the rest of the spectrum that intuitQ made.

Note that because of the way in which intuitNull operates this operation should typically only be used before any use of the shelving tools. If you happen to do it after using those tools the accuracy can be severely affected. This is an algorithmic limitation.

Cursor | Match Loudness Command



Use the Cursor | Match Loudness command to do loudness matching based on a selected frequency range of the track spectrum. Typically when performing loudness matching of a quiet track to a loud reference (say a piano vocal duo with a rock song) full spectrum loudness matching gives poor results, with the quiet track being over amplified.

This behaviour arises because the instrumentation is very different between the two tracks. In the rock song the entire spectrum is filled by sound from all the instruments whereas the piano vocal duo is largely confined to the spectrum below 2kHz. In effect, if you do a full spectrum match in such a case the match loudness function demands that the piano and vocal play much louder than they would normally to make up for the shortfall of loud instruments present in the rock track.

What we really want to do is make the track loudness match the reference in the spectrum range that the instruments in that track occupy. Performing the matching over this limited range where the track spectrum energy is high results in better level match between loud and soft tracks.

To perform match loudness based on a limited frequency range move the cursor to the starting point of the frequency range, click and hold down the left mouse button and then drag the cursor to the end point of the desired frequency range. When you release the left mouse button match loudness will be applied based on the selected frequency range. Match loudness is only available if a reference is loaded.

Cursor | Gain Command



Use the Cursor | Gain command to design a matching equalization filter for the session file. Once selected the cursor icon changes to the gain cursor. You can then use the gain cursor to alter the shape of the spectrum plot.

The gain cursor is, in essence, equivalent to a parametric equalizer. The centre frequency is chosen by moving the cursor to the desired frequency on the spectrum plot and then pressing and holding down the left mouse button. The gain is controlled by moving the mouse up or down and the Q is controlled by moving the mouse left or right. During the dragging operation a preview of the filtered spectrum follows the mouse movement. To keep the Q fixed while making boost/cut adjustments hold the SHIFT key down (when the left mouse button is down). To keep the Q at the maximum value press and hold the M key down. When you release the left mouse button the filter change is made and the spectrum plots updated. If you wish to cancel an edit press the ESCape key prior to releasing the left mouse button.

Note that when using the gain cursor you can switch focus from average to peak spectrum plots by pressing the tab key whilst the Gain cursor is active.

Cursor | Low Shelving Command



Use the Cursor | Low Shelving command to design a matching equalization filter for the session file. Once selected the cursor icon changes to the low shelving cursor. You can then use the low shelving cursor to introduce a low shelf boost or cut at a given corner frequency with a given slope.

The low shelf cursor is similar to the shelving response of a bass tone control without the plateau. The corner frequency is chosen by moving the cursor to the desired frequency on the spectrum plot and then pressing and holding down the left mouse button. The slope is controlled by dragging the mouse up or down. During the dragging operation a preview of the filtered spectrum follows the mouse movement. When you release the left mouse button the filter change is made and the spectrum plots updated. If you wish to cancel an edit press the ESCape key prior to releasing the left mouse button.

Note that when using the low shelf cursor you can switch focus from average to peak spectrum plots by pressing the tab key whilst the low shelving cursor is active.

Cursor | High Shelving Command



Use the Cursor | High Shelving command to design a matching equalization filter for the session file. Once selected the cursor icon changes to the high shelving cursor. You can then use the high shelving cursor to introduce a high shelf boost or cut at a given corner frequency with a given slope.

The high shelf cursor is similar to the shelving response of a treble tone control without the plateau. The corner frequency is chosen by moving the cursor to the desired frequency on the spectrum plot and then pressing and holding down the left mouse button. The slope is controlled by dragging the mouse up or down. During the dragging operation a preview of the filtered spectrum follows the mouse movement. When you release the left mouse button the filter change is made and the spectrum plots updated. If you wish to cancel an edit press the ESCape key prior to releasing the left mouse button.

Note that when using the high shelf cursor you can switch focus from average to peak spectrum plots by pressing the tab key whilst the high shelving cursor is active.

Cursor | Dynamics Node Command



Use the Cursor | Dynamics Node command to design a dynamics processing scheme for the session file. Once selected the cursor icon changes to the dynamics node cursor. You can then use the dynamics node cursor to create, edit and delete dynamics control nodes.

Dynamics control nodes are created by clicking and/or dragging the dynamics node cursor on the histogram view. Doing so creates a control node comprising two circles and vertical lines. This schema represents the mapping of input level to output level at that point. The solid circle represents the input level and the open circle the output level that the input level is mapped to. The vertical positioning of the nodes has no bearing on the transfer function characteristic that the nodes define, only the horizontal position has.

Once placed a control node can be edited by mousing over the appropriate circle and clicking and dragging to the desired location. The input level circle if dragged, moves the output level in tandem and maintains the difference in level between them. The output level circle if dragged, does not move the input level circle. A control node can be deleted by dragging it off the histogram view.

For more information on how dynamics processing is achieved in Har-Bal, consult the discussion of dynamics processing in the [introduction](#).

Options Menu

- [Preferences Command](#)
- [Open Tips Command](#)
- [Show Tips Command](#)

Options | Preferences Command

Use the Options | Preferences command you to customise the behaviour and setup of Har-Bal. This command opens the [Preferences](#) dialog box.

Options | Open Tips Command

Use the Options | Open Tips command to open a tip file that can be displayed on the spectrum view as a guide to filtering tracks. Har-Bal is supplied with a number of tip files that can be found the the HarBal\tip_files folder and include:

- General.etp
- Brass_french_horn.etp
- Brass_trombone.etp
- Brass_trumpet.etp
- Brass_tuba.etp
- Percussion_bassdrum.etp
- Percussion_cymbals.etp
- Strings_bassviol.etp
- Strings_cello.etp
- Strings_guitar.etp
- Strings_viola.etp
- Strings_violin.etp
- Voice_alto.etp
- Voice_bass.etp
- Voice_soprano.etp
- Voice_tenor.etp
- Woodwind_bassoon.etp
- Woodwind_clarinet.etp
- Woodwind_flute.etp
- Woodwind_oboe.etp

When filtering post-mixdown tracks use the General.etp tip file. When filtering mono tracks from a multi-track source (ie. prior to mixdown) use the tip file most closely corresponding to the type of instrument recorded in that channel. Note that you can create your own tip files with a text editor. Open any one of the above files and use it as a guide in creating your own tips files.

Options | Show Tips Command

Use the Options | Show Tips command to toggle tips on and off. If a tip file is loaded and the Show Tips command is checked then tips will be displayed on the spectrum view when using the Gain, Low Shelf and High Shelf cursors.

Help Menu

- [Contents Command](#)
- [Example Tutorial Command](#)
- [Using help Command](#)
- [About Command](#)

Help | Contents Command

shortcut: F1

Use the Help | Contents command to open the help file for Har-Bal at the contents page.

Help | Example Tutorial Command

Use the Help | Example Tutorial command to open the example tutorial document. New users of Har-Bal are encouraged to study the tutorial in detail as it provides a good introduction to Har-Bal and effective ways of using it.

Help | Using help Command

Use the Help | Using Help command to get help using Har-Bal help.

Help | About Command

Use the Help | About command to open the About dialog which displays version and copyright information for Har-Bal and contact details for the purposes of product queries and/or complaints.



Har-Bal Harmonic Balancer

Har-Bal™ new thinking, new directions...

Dialog Boxes

- [Open Reference](#)
- [Graph Properties](#)
- [Manage Files](#)
- [New Filter](#)
- [Tasks](#)
- [Progress](#)
- [Preferences](#)
 - [Equalizer Options](#)
 - [Driver Options](#)
 - [Graph Options](#)

Open Reference

Use the Open Reference dialog box to select a sound file or a reference file (.hbref) a reference. The analysis of the reference is displayed along with the analysis of the session file. If the reference file is a sound file and has the same sampling rate as the session file then the [reference playback](#) feature will be enabled.

Source File Edit Box

Use the source file edit box to enter the filename and path of a sound file or a reference file to use as a reference. Press the ... button to select the file using an open dialog box.

Filter File Edit Box

Use the filter file edit box to enter the filename and path of a filter file to combine with the reference sound file. This file is optional. If specified the reference spectrum is altered according to the characteristics of the filter. Press the ... button to select the filter file using an open dialog box.

Graph Properties

Use the Graph Properties dialog box to alter the displayed range for graphical results. Note that you can also change the displayed ranges graphically using the zoom cursors.

X and Y Axis Maximum and Minimum Edit Boxes

The X and Y axis maximum and minimum edit boxes are used to specify the displayed range on either axis. To change the displayed range, type in new minimum and maximum values for either axis.

Manage Files

Use the Manage Files dialog to manage and delete the working files that Har-Bal creates and uses to process your tracks. Har-Bal will delete .hba, .hbfb, .hbfbfs, .hbfr, .hbfs, .hbir files but not .hbf files. This is by design as the .hbf file contains all the information required to re-create the processing designed for a given session and is something you are likely to want to archive for future reference. If you want to delete superfluous .hbf files then use Explorer (on Windows) or Finder (on Mac OS X) to do so.

Active Files List Box

Select the active files whose associated working files you would like to delete. You can select all by pressing the **Select All** button or select no files by pressing the **Clear All** button. Pressing the **Delete** button will delete the working files associated with the selected active files.

New Filter

Tasks

Analysing

When performing background analysis the tasks dialog shows the queue status of pending and current analysis jobs. Job status is displayed in a list view with column headings for the Input File, Analysis File and % Complete. The first item in the list view is the current job and the % Complete column updates as the job progresses. You can delete any job from the list by selecting the job and pressing the **Del** key.

Recording

When Recording the tasks dialog shows the queue status of pending and current recording jobs. Job status is displayed in a list view with column headings for the Input File, Filter File, Output File and % Complete. The first item in the list view is the current job and the % Complete column updates as the job progresses. You can delete any job from the list by selecting the job and pressing the **Del** key.

Progress

When performing analysis the progress dialog shows the progress of analysis. Analysis has two phases. The first phase does spectrum analysis of the entire track. The second phase builds indexing information by sorting the spectrums. When opening mp3 files a progress dialog shows the progress of converting the mp3 to a .wav or .aif file.

Preferences

Equalizer Options

Use the Equalizer Options tab to customise the behaviour of Har-Bal.

Filter Topology Radio Buttons

Har-Bal can realise FIR filters with either linear phase or minimum phase characteristics. By default Har-Bal uses minimum phase filter realisations.

Return to Start Position when Play Stopped Check Box

When checked Har-Bal returns the playback position to the position the track was cued to before to pressing play. When not checked the cue position remains at the position where playback was stopped.

Show Splash Screen on Startup Check Box

When checked Har-Bal displays a splash screen when the application is starting. When not checked no startup splash screen is displayed.

Restore Files on Startup Check Box

When checked Har-Bal loads the files that were opened when Har-Bal was last shut down. When not checked Har-Bal starts with no files open.

Driver Options

Use the Driver Options tab to select the device to use for playback.

Output Interface Radio Buttons

In Windows Har-Bal supports three different interfaces to communicate with sound card hardware. You can choose between ASIO, Direct X or Multimedia interfaces. In Mac OS X sound output is handled through AU drivers.

Driver List Box

This list box shows all the drivers installed for the particular interface that you have selected. The driver that is selected from this list is the driver that Har-Bal will attempt to use when previewing your harbalized track.

Patching Dropdown List Boxes

Select the desired channel associations from the left and right channel drop down list boxes. Note that you cannot select the same physical channel for left and right channels.

Graph Options

Use the Graph Options tab to change the visual appearance of the graphs in Har-Bal and the graphs exported to the clipboard or sent to the printer. You can change the colour of selected graphical features and the fonts used to display text in the views.

Normal Colours Tab

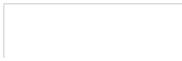
Use the normal colours tab to customise the colours used in drawing the graphs on screen in the spectrum, frequency response and histogram views.

Colour Catagory List Box

Use the colour catagory list box to select the graphical feature whose colour you wish to alter. After selecting the desired catagory the current colour of that feature will be indicated by the selected colour in the colour palette. You can change the colour of that feature by selecting a different colour in the colour palette.

Colour

Use the (H)ue, (S)aturation and (L)ightness slider controls to alter the colour of a particular graphical feature. After selecting the desired colour catagory from the colour category list box realise the desired feature colour by using the sliders or edit fields.



Copyright © 2011
Har-Bal™. All rights
reserved.

Printer Colours Tab

Use the printer colours tab to customise the colours used in printing the graphs on screen in the spectrum, frequency response and histogram views. These colours are also used when the [Graph|Clipboard copy](#) menu command is used.

Colour Catagory List Box

Use the colour catagory list box to select the graphical feature whose colour you wish to alter. After selecting the desired catagory the current colour of that feature will be indicated by the selected colour in the colour palette. You can change the colour of that feature by selecting a different colour in the colour palette.

Colour

Use the (H)ue, (S)aturation and (L)ightness slider controls to alter the colour of a particular graphical feature. After selecting the desired colour catagory from the colour category list box realise the desired feature colour by using the sliders or edit fields.

Title Font Tab

Use the title font tab to select the font to use when drawing the axis titles in the spectrum, frequency response and histogram views.

Choose Font Button

Click the **choose font** button to open a font chooser dialog box and select the font to use in drawing the axes title text.

Font Preview Group Box

The font preview group box shows a preview of the appearance of the selected font as would be seen on screen. Note that the actual preview text is arbitrary and has no bearing on the views.

Unit Font Tab

Use the unit font tab to select the font to use when drawing axis units (numbers) in the spectrum, frequency response and histogram views.

Choose Font Button

Click the **choose font** button to open a font chooser dialog box and select the font to use in drawing the axes units text.

Font Preview Group Box

The font preview group box shows a preview of the appearance of the selected font as would be seen on screen. Note that the actual preview text is arbitrary and has no bearing on the views.

Tips Font Tab

Use the tips font tab to select the font to use when drawing tips.

Choose Font Button

Click the **choose font** button to open a font chooser dialog box and select the font to use when drawing tips.

Font Preview Group Box

The font preview group box shows a preview of the appearance of the selected font as would be seen on screen. Note that the actual preview text is arbitrary and has no bearing on the tips.



Har-Bal Harmonic Balancer

Har-Bal™ new thinking, new directions...

Keyboard Shortcuts

Windows Shortcut	Mac OS X Shortcut	Function
V + 0	V + 0	Zero volume slider
V + Up Arrow	V + Up Arrow	Increment volume slider by 1 dB
V + Down Arrow	V + Down Arrow	Decrement volume slider by 1 dB
V + SHIFT + Up Arrow	V + SHIFT + Up Arrow	Increment volume slider by 0.1 dB
V + SHIFT + Down Arrow	V + SHIFT + Down Arrow	Decrement volume slider by 0.1 dB
G + 0	G + 0	Zero gain slider
G + Up Arrow	G + Up Arrow	Increment gain slider by 1 dB
G + Down Arrow	G + Down Arrow	Decrement gain slider by 1 dB
G + SHIFT + Up Arrow	G + SHIFT + Up Arrow	Increment gain slider by 0.1 dB
G + SHIFT + Down Arrow	G + SHIFT + Down Arrow	Decrement gain slider by 0.1 dB
A + 0	A + 0	Zero air slider
A + Up Arrow	A + Up Arrow	Increment air slider by 1%
A + Down Arrow	A + Down Arrow	Decrement air slider by 1%
A + SHIFT + Up Arrow	A + SHIFT + Up Arrow	Increment air slider by 0.1%
A + SHIFT + Down Arrow	A + SHIFT + Down Arrow	Decrement air slider by 0.1%
F1	CTRL + ?	Open application help
F4	CMD + Q	Quit application
CTRL + SHIFT + 1	CTRL + SHIFT + 1	Select 1 octave resolution
CTRL + SHIFT + 2	CTRL + SHIFT + 2	Select 1/3rd octave resolution
CTRL + SHIFT + 3	CTRL + SHIFT + 3	Select 1/6th octave resolution
CTRL + SHIFT + 4	CTRL + SHIFT + 4	Select 1/12th octave resolution
ALT + 1	ALT + 1	Select spectrum view
ALT + 2	ALT + 2	Select frequency response view
ALT + 3	ALT + 3	Select histogram view
CTRL + 1	CTRL + 1	Toggle peak power plot
CTRL + 2	CTRL + 2	Toggle realtime power plot
CTRL + 3	CTRL + 3	Toggle average power plot
CTRL + 4	CTRL + 4	Toggle original (unfiltered) plots
CTRL + 5	CTRL + 5	Toggle playback focus markers
CTRL + 6	CTRL + 6	Toggle reference plots
CTRL + 7	CTRL + 7	Toggle target reference plots
ALT + Z	ALT + Z	Zoom all applied to view
CTRL + I	CTRL + I	Apply intuitQ
CTRL + ALT + I	CTRL + ALT + I	Apply intuitMatch
CTRL + O	CMD + O	Open a session sound file
CTRL + F	CMD + F	Open a session filter file
CTRL + R	CMD + R	Open reference sound and filter files
CTRL + Z	CMD + Z	Undo edits
CTRL + Y	CMD + SHITF + Z	Redo edits
CTRL + D	CTRL + D	Render / Record filtered session to file
CTRL + A	CTRL + A	Toggle air in / out of playback and rendering
CTRL + C	CTRL + C	Toggle dynamics processing in / out of playback and rendering
CTRL + Q	CTRL + Q	Toggle equalisation in / out of playback and rendering
CTRL + W	CTRL + W	Toggle playable reference in / out of playback
CTRL + M	CTRL + M	Apply match loudness
SPACEBAR	SPACEBAR	Toggle playback on / off

Copyright © 2011 Har-Bal™. All rights reserved.