

## DSP Tools

DSP stands for Digital Signal Processing. In this section, we describe several powerful signal processing functions built into the program. DSP techniques generally offer significantly more flexibility in dealing with audio signals. Some functions, routinely implemented using DSP, are not even available to their “analogue” counterparts.

DSP is the future of audio, so we need to embrace it and hopefully, use it to its full potential. If you look “under the hood” of a typical DSP process – you will find a lot of difficult mathematics. And this may prompt you to hesitate and label the concept as “too difficult to use”. However, this is where programs like SoundEasy come in handy. Here, the nuts-and-bolts of a particular DSP implementation are invisible to the user, and all that you see, are familiar filter names, filter parameters, and so on... So please think again. The DSP is really good for your speakers.

SoundEasy running DSP process on a PC is not a stand-alone solution for the actual loudspeaker. Therefore, once the process has been designed, tested and confirmed on a PC with a sound card, the functionality needs to be moved onto a free-standing DSP engine to form a heart of a powered loudspeaker system.

### System Phase Linearizer

The first device, the System Phase Linearizer, is designed to annihilate phase delays caused by the crossover, the enclosure, and the natural roll-off of the drivers’ frequency response. It would be used with fully completed loudspeaker system, and inserted in the audio chain just before the audio power amplifier. You may think of this function as pre-distorting the phase response to make it a mirror image of the system phase response, so that when these two are put together, the phase response comes out flat.

How do you pre-distort the phase response in a multi-way loudspeaker system?. The idea is rather simple. Using the provided controls, you mimic the **acoustic transfer function of each driver in the system**, and the summation of all channels is done automatically by the program. As a result, you should come very close to the actual acoustical transfer function of the system – that is, you obtain amplitude and phase response of the system – see Figure 19.1.

Thanks to the power of DSP, we can now disregard amplitude response and force it to be a flat line in frequency domain. And now, all we need to do, is to create a mirror image of the phase response.

A technique that allows us to reverse phase response of an LTI system is based on time-inverting its impulse response. In practical terms, you need to reverse time scale of the impulse response. Typical impulse response described in time domain starts with a high peak, followed by its decaying tail. Imagine reversing time scale – it would be like the start of the impulse was the tail end of it, and then we arrive at the peak. If you perform this operation as convolution, interestingly, the phase response becomes its own mirror image around zero degrees. And that’s exactly what we are looking for. In addition, if we are using FIR ( Finite Impulse Response) techniques to accomplish the convolution, the amplitude response of the process can be decoupled from the inverted phase response and then forced to become flat.

Please note, that the combined transfer function of the System Phase Linearizer plus the actual loudspeaker results in the same amplitude response as the loudspeaker, but flat phase response. If your drivers acoustical responses sums up to flat, you get a system with flat amplitude response and flat phase response.

The benefit of the above is perhaps most obvious when comparing time-domain system responses to a square wave excitation – see Figure 19.2

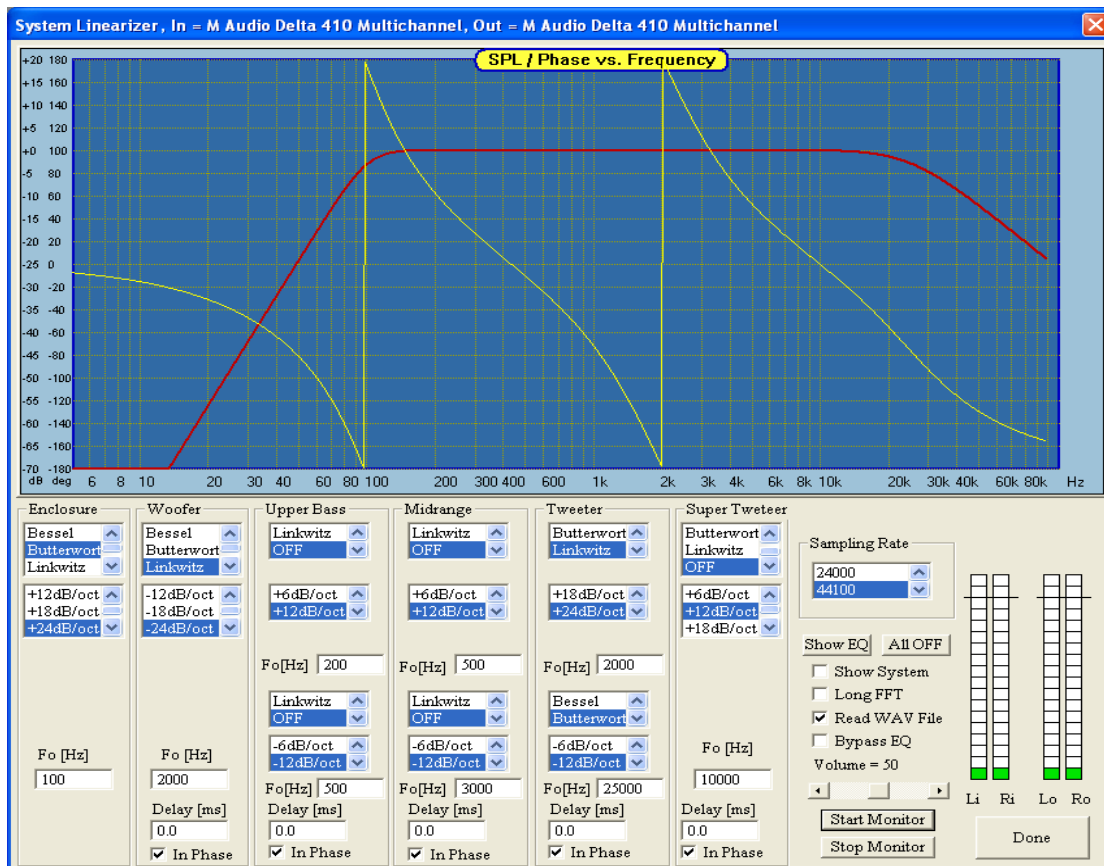


Figure 19.1. Simulated frequency/phase response of a loudspeaker system




Figure 19.2. Top trace – two-way crossover response with original phase delays.  
Bottom trace – same crossover response as top trace + inverted phase.

### Description of Controls - Midrange Section Example.

1. **Filter Type list box** – There are Butterworth, Linkwitz and Bessel-type high-pass filters to choose from.
2. **Filter Order list box** – You can select 6/12/18/24/48 dB/oct high-pass filters.
3. **Fo[Hz]** – 3dB cut-off frequency.
4. **Filter Type list box** – There are Butterworth, Linkwitz and Bessel-type low-pass filters to choose from.
5. **Filter Order list box** – You can select 6/12/18/24/48 dB/oct low-pass filters.
6. **Fo[Hz]** – 3dB cut-off frequency.
7. **Delay [ms]** – Data entry field for midrange delay. Must be less or equal to 1ms.
8. **“In Phase”** – Check box to select “in-phase” or “out-of-phase” midrange connection.

### Other Controls

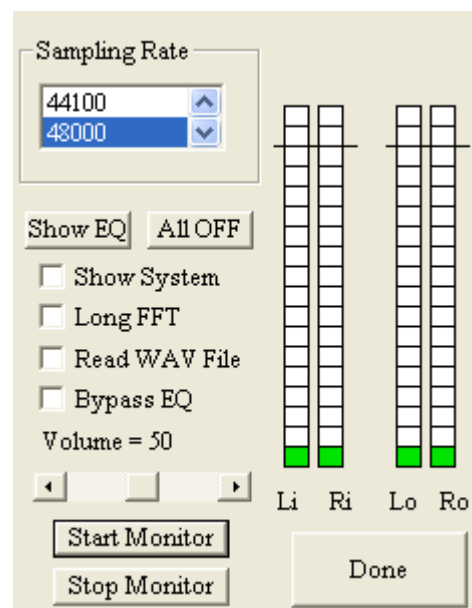


Figure 19.3. Phase Equalizer controls

1. **Sampling Rate** – This list box allows for selection of 44.1 kHz or 48kHz sampling rates.
2. **“All OFF”** – This button turns OFF all filters.
3. **“Show EQ”** – When “checked”, it will enable plotting phase response in addition to amplitude responses.
4. **“Show System”** – Shows the original system response.
5. **“Long FFT”** – Select longer FFT process.
6. **“Read WAV File”** – Check box to allow input sound data from .WAV file.
7. **“Bypass EQ”** – Turns OFF the EQ function.
8. **Volume control** – This slider controls output signal in 0 – 100% range.
9. **“Start Monitor”** – This button activates the Equalizer process.
10. **“Stop Monitor”** – This button terminates the Equalizer process.
11. **“Done”** – Closes the Equalizer screen.
12. There are also input and output level LED indicators.

### Synthesized Bass

So, when played on the new hi-fi system, your favourite music sounds overly bright ?. Perhaps it is due to poor recording practices, or perhaps the bass player himself did not explore his instrument fully?. This seems to be generally true, even for some great artists recorded in 60's and 70's. Contemporary recordings appear definitely better balanced tonally, but even so, the device described in this article may add unexpected dimension to your listening pleasure.

It would be desirable to be able to fully exploit the benefits of good quality loudspeaker system. A well-designed system of today, can deliver broad-spectrum sound, easily reaching down to 20Hz. But how often do you actually hear a really deep bass note coming out of your subwoofer?. In fact, the loudspeaker system can well expose all shortcomings of the recording studio practices, be it too much undesirable noise, or poor tonal balance. Graphic equalizer is perhaps a helpful tool in correcting some of the issues referred to above. Even so, I found myself re-adjusting the knobs from one song to another. Settings, that worked for one song, seemed to be doing the opposite for another. So, is there anything at all, that could help in this situation?. Perhaps the answer is provided by the music itself.

### **A little digression into music**

Musical **octave** is the interval between one musical pitch and another with half or double its frequency. For instance, if one note has a frequency of 100Hz, the note an octave above it will have 200Hz frequency, and the note an octave below will have 50Hz frequency. Interestingly, the human ear tends to hear both notes as being essentially the same, due to the same harmonic relationship. Not surprisingly, notes an octave apart are given the same note name. This is called octave equivalency, and assumes, that notes one or more octave apart are musically equivalent.

Given the above – what if we could make use of the low-frequency notes as a guiding reference for synthesizing the same notes, but one octave lower ?. Sub-harmonics addition - this is the basic idea, behind enhancing the low-frequency content of the existing recordings.

### **Back to technical issue**

Implementing FIR filtering techniques on a personal computer with a sound card is so well documented in the available literature, that I will only briefly mention it's existence later.

Standard pitch-shifting algorithm (also known as Phase Vocoder) is based on a two-stage processing. In Stage 1, or analysis stage, a sound sample is windowed and then a short-time Fourier transform (STFT) is applied to each segment, or “frame”. Typical values for the length of the frame are: 1024, 2048, 4096 and 8192 samples or “bins”. The resulting spectrum can be manipulated accordingly to the expected results and then re-synthesis, or Stage 2, is accomplished by applying the inverse STFT to each segment. This will return the sound sample back to time domain, and the segment is ready to be played out by the sound system. In order to avoid discontinuities at the frame boundaries, the segments/frames need to be windowed after re-synthesis and overlapped, which means, that much of the processed signal is fed back to the input and re-processed again and again, depending on “overlap factor”. The overlap factor tells you how many times the frame is re-processed. This process is known as WOLA – Weighted OverLap Add. Both: analysis and synthesis stage typically attempt to use amplitude and phase information in each of the FFT bins. It should be clear by now, that analysis of the frame and synthesis of the sample are done in frequency domain.

Unfortunately, low-frequency sound quality produced by such method is questionable. Sound notes parameters such as “attack”, “sustain” and “decay” all suffer degradation. The “attack” is slow and the sound builds itself over several frames. The “sustain” suffers from amplitude fluctuations and exhibits a fair amount of echo. The “decay” is too slow and smeared.

Overall, the sound suffers from “phasiness”, echo, and amplitude instabilities. Much of the amplitude instabilities are attributed to lack of phase synchronization between FFT bins. That is, if the frequency domain analysis stage uncovered a bass note that falls on the boundary of two bins, the phases of those bins are not synchronized, as the bins know nothing about each other. I have experimented with standard PV algorithm and tried various possible improvements, but finally discarded it as unsuitable for the bass synthesiser.

### **Somewhat different approach**

There are two processing streams within the bass synthesizer. Standard sound processor has 5 processing units: selection of high-pass filters, two shelving filter selections, a selection of low-pass filters and finally Q-notch/peak filter. Each filter has several characteristic parameters to choose from. This array of tools allows for much more professional sound processing than bass/treble set of controls

on a typical amplifier, or even graphic equalizer. One can cut excessive sub-audio, low-frequency noise without actually cutting off useful bass information. A 50Hz hum can be removed with Q-notch filter. Furthermore, one may prefer to reduce “harshness” or add “brightness” to the favourite recording.

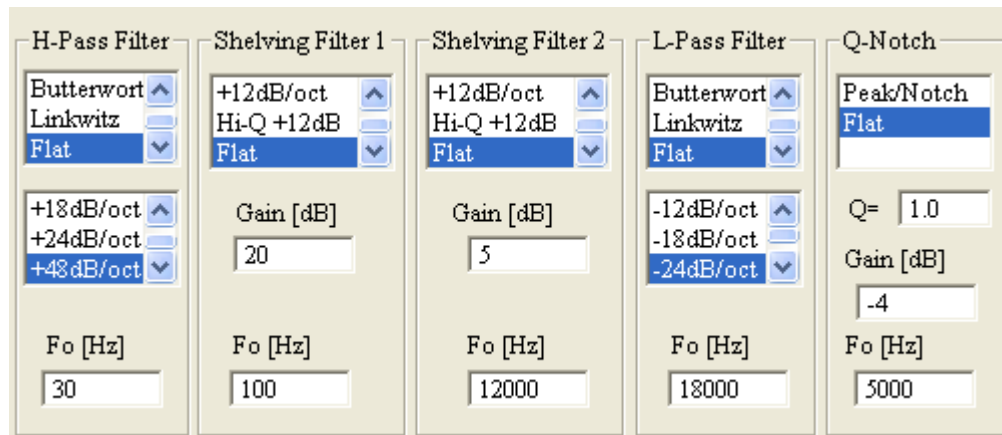


Figure 19.4. Bass Synthesizer DSP filtering section

## Bass Synthesizer

The second stream is the bass synthesizer. Algorithm implemented in the bass synthesizer is also a two-stage process. Collected sound sample is FFT transformed into frequency domain for analysis. The algorithm only scans frequency range between two nominated limits: F1 and F2.




Figure 19.5. Very simple bass controls of the bass synthesizer.

However, the analysis stage discards information about the phase of each FFT bin. The technique used here is known as STFTM – Short Time Fourier Transform Magnitude. The length of the frame is selectable as: 4096 or 8192 samples. The synthesis stage is however performed in the time domain. The synthesized waveform is then transformed into frequency domain, where is low-pass filtered and then mixed into both left- and right-channel sound streams. The synthesized waveform is always started with phase=0 degrees, so the “attack” time is practically zero. All the above happens within single “frame”. The phase of the synthesized waveform is maintained from one frame to another. The overall processing algorithm uses standard “Overlap-and-Add” DSP technique, but it does not employ the WOLA approach. Since there is no re-processing of synthesized frames, the resulting sound is free from “echo” and “phasiness”.

Operation of the bass synthesiser in frequency domain is shown on several spectrograms below. Figure 19.6 depicts a sample spectrogram of a well-known pop-music artist. The song comes from “A hard Day’s night” CD, and is titled: “I should have known better”, The Beatles, 1964 EMI records LTD. It is immediately observable, that low frequency content of this recording is very poor. The song itself is great, but to the listener’s ear, the music sounds too bright and ill-balanced.




Figure 19.6. Sample spectrogram of “I should have known better” – The Beatles.




Figure 19.7. Same song as on Figure 1, but with bass synthesizer turned ON.



The bass synthesized can work with already very low frequencies. A good example is shown below, where a 36Hz bass note is re-synthesized down to 18Hz. At this point it would be important to realize, that the bass synthesizer can inflict serious damage to subwoofer. For instance, by setting low frequency boundary to 20Hz, you allow the synthesizer to generate 10Hz sub-harmonics at high amplitude.




Figure 19.8. Artist: George Michael, song called “Spinning the wheel”, comes from CD “Older”.

### How does it actually sound

Good. The reason for it is, that the only thing that is different from the original recording is the addition of a sub-harmonic to the lowest bass notes in the recording. The vital amplitude-phase relationship of the original harmonics in the bass note has not only been preserved, but also enhanced with an extra sub-harmonic, one octave down.

In order to visualize the above, I have temporarily modified the algorithm to allow the synthesized sub-harmonics to be outputted from the sound card on a separate channel. Therefore, the music channel and the sub-harmonic channel can be simultaneously captured on a digital storage oscilloscope for comparison – as shown on Figure 19.9.

It can be observed on Figure 19.9, that the sub-harmonic, generated as a pure sine-wave, takes on characteristics of a live, instrument-generated sound, when combined together with the original set of harmonics. This is evident in the fact, that the sine-wave has changed it’s smooth, “mathematical” shape and after adding the original set of harmonics, it now fully contributes to the original instrument’s sound.

Interestingly, Figure 19.9 captures transition between two different frequencies and different amplitudes of bass notes. The uninterrupted transition within 10 ms is clearly evident, indicating high level of performance of the synthesizer.

Additional and desired characteristics relate to the envelope of the synthesized sub-harmonic, namely: zero attack time, fast amplitude updates every 185ms ( this is due to the STFTM frame length, which is 8192 bins and sampling frequency of 44100Hz so that frame length in time domain is:  $8192/44100 = 185\text{ms}$ ), and instantaneous decay. All these characteristics contribute to maintaining the “musical” character of the final sound.




Figure 19.9. Trace 1 – music channel + sub-harmonic channel, 1V/dev  
Trace 2 – sub-harmonic channel alone, 0.5V/dev.

## Conclusions

The device described in this article incorporates several audio DSP filters, a bass synthesizer and can record into a WAV file. Thus allowing the user to create his own, perfectly balanced, bass-enhanced “.wav” file recordings, that can be made into a standard CD using any off-the-shelf CD recording software, like “RecordNow”.

The bass synthesizer works best if the musical content has single, well defined bass notes within user selected frequency range. The resulting new bass notes actually sound quite musical, despite being completely synthetic. Perhaps mixing them back into the original sound stream makes them feel like the real artist playing. DSP processing artefacts are not audible during typical processing application.

By now, I have re-recorded all The Beatles song catalogue for my own use. I have actually stopped listening to the original CDs, as they sound too poor in comparison with the restored and enhanced material.

So, are you ready to keep your subwoofers really busy?.



## Advanced System Phase Linearizer

One of the most useful test signals in electronics is a humble square wave. The “ideal” square wave is a superposition of an infinite number of sine waves, each contributing its required amplitude and phase. It is due to this very feature, that when passed through an audio system, the square wave can reveal time domain performance issues of the system. This is because all of its sine wave components must be passed by the system without time distortion, or different delays, in order to recombine as a square wave at the output of the system under test.

Practically generated square waves have limited number of contributing sine waves, but the number is still sufficiently large, that for audio frequency range testing, we can fully utilize the “almost perfect” square waves.

It is important to realize, that the system time delay does not need to be zero. But it must be the same for all frequencies within the audio system pass band. Such condition will be easily met if the system under test has a zero-degree phase response.

### Real-life loudspeaker example

The system under test discussed here consists of a filter and a loudspeaker in an enclosure. The two components that will introduce time delay are the filter and the combination of driver and the enclosure itself.

To illustrate the above, a 12” guitar loudspeaker in a vented box was measured and its minimum-phase (MP) responses were obtained with a help of an MLS measurement technique – see Figure 1 below. It is immediately observable, that the loudspeaker has rather irregular frequency response. Since the loudspeaker is essentially a minimum-phase device, the corresponding phase response is also highly irregular, and definitely not flat.

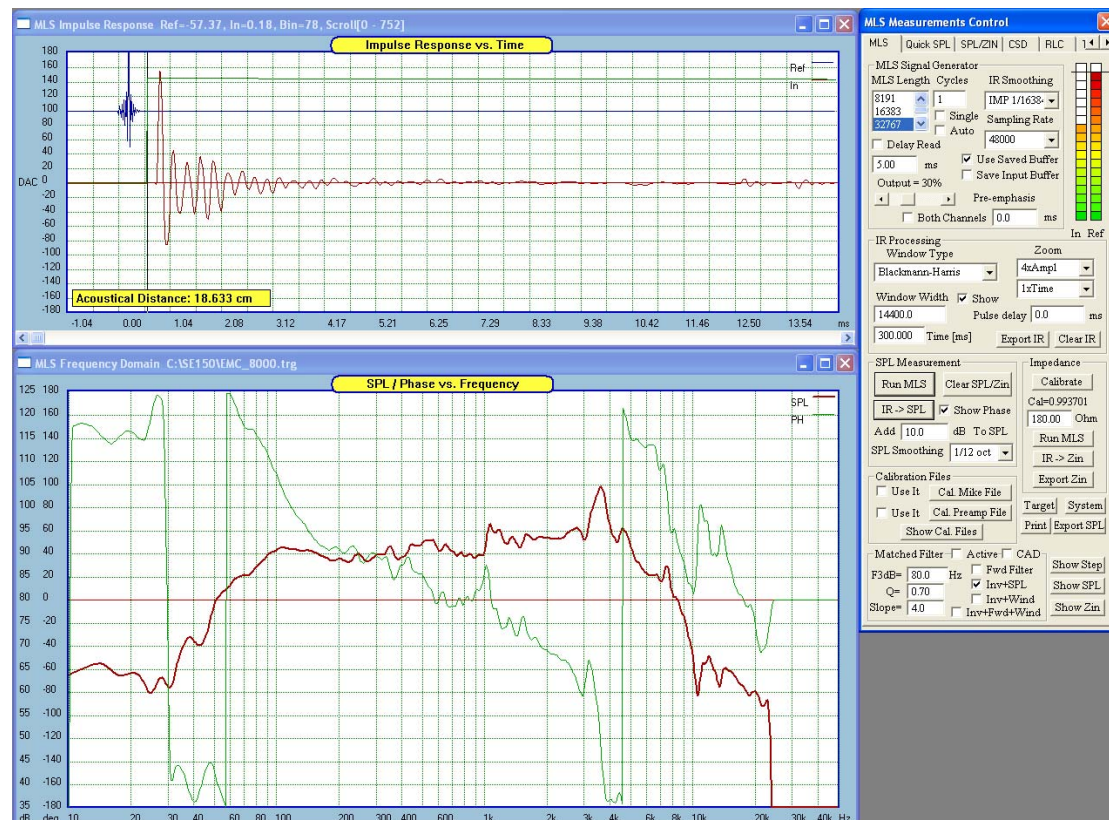


Figure 19.10. Measured system's magnitude (red) and phase (green).

Let's establish the **frequency response of interest**, which is the frequency range where the SPL will be equalized to flat response. In my example it will be: 91Hz – 5250Hz. A 300Hz square wave reproduced by this loudspeaker is highly distorted.




Figure 19.11. Distorted square wave passed through system on Figure 19.10.

The outgoing waveform resembles a triangular-wave with a high level of ringing imposed on it. The ringing is the result of highly irregular frequency/phase response from 1kHz to 6kHz, with an additional +10dB peak around 3.5kHz. The resulting square wave is shown on Figure 2.

Unfortunately, this is what most loudspeakers do – irregular frequency response, coupled with accumulated delays in the system prevent the loudspeaker from correctly recombining all sinusoidal components of the square wave. The result is quite poor reproduction of musical transients. Just like the vertical transitions of the square wave and their flat sections are reproduced as spikes with ringing, musical transitions are also highly distorted.

But musical transients are difficult to capture as they change constantly. More importantly, we do not really know what the “undistorted” musical transient should look like in time domain.

So, even though nobody listens to square waves, they provide a well-defined test signal, that will be very easy to compare with distorted square waves passing through the loudspeaker. The distortions, or deviations, are therefore immediately obvious, and provide clues as to the origins of the distortions. High level of ringing in the resulting square wave is related to irregular frequency response. This gives us the first clue in the quest to improve the shape of the outgoing square wave – flatten the frequency response.

A popular tool used for linearizing a transfer function of an LTI (Linear Time-Invariant) system is a Hilbert-Bode Transform (HBT). Just like Fourier Transform allows you to flip between time domain and frequency domains, the HBT allows you to move from magnitude response to phase response and vice-versa. I can therefore nominate a frequency range of interest within the loudspeaker's magnitude response, then attach flat “tails” on the low and high-side of this frequency range and apply this artificially created magnitude response to the HBT. As a result, I will get corresponding phase response, which in turn means, that I actually have full complex transfer function calculated via HBT.




Figure 19.12. Advanced System Linearizer controls.

### On-axis vs. off-axis equalization

Typical loudspeaker will exhibit somewhat different frequency response when measured off-axis. Figure 19.13 illustrates this situation very well. Since our HBT-based equalizer is being designed for on-axis performance, there is a need to consider its impact on off-axis performance of the loudspeaker.




Figure 19.13. Off-axis SPL. From "Loudspeaker Response Equalization Using Warped Digital Filters" – Matti Karjalainen, Esa Piirila, Antti Jarvinen – NORSIG 96.

It is observable, that the 0deg, or on-axis curve, is the flattest within the whole family of responses. This is a very common occurrence and is due to the fact, that loudspeakers are typically designed and optimized for the on-axis performance. Therefore the equalizer design for on-axis performance requires automatically the least amount of equalization to perform. This is actually quite often advantageous for off-axis performance as well. If you examine the off-axis curves on Figure 4, you will notice, that as the angle increases, the curves have gradually increased peaks and valleys, but the deviations hold their locations on frequency scale. If the same equalizer was applied to each of those curves, only the on-axis curve would be fully equalized. Other curves would be equalized to lesser degree, but they would not be over-equalized. For instance, the 5deg curve has very similar shape as the 0deg curve, but the notch at 18kHz is much deeper. If this curve was to be corrected with the same equalizer, then most of the curve would become almost perfectly flat, but the notch at 18kHz would still be evident to some degree. You can examine the curves for other angles and you would most likely conclude, that the equalizer would still help correcting many of the irregularities evident in off-axis curves. The above situation is typical, however, not all loudspeakers are as easy to handle as the one just discussed. If the peaks and valleys of the off-axis frequency response are unstable in frequency domain, than the effectiveness of equalizer will be reduced.

### Amplitude Equalizer design

The importance of having the error transfer function is, that it can be convolved with the loudspeaker's transfer function to linearize it's frequency response.




Figure 19.14. Amplitude Error Function – magnitude (blue), phase (orange)



The linearization will be accomplished, if the HBT-calculated transfer function is first inverted. The result of this operation is a mirror-imaged frequency and phase response calculated within the previously nominated frequency range, and flat response everywhere else.

We have now created an Amplitude Error Function. The thick blue line is the SPL of the Amplitude Error Function (notice, it's inverted already), and orange line is the phase of the Amplitude Error Function – see Figure 5. Please note mathematically correct phase response and it's transitions from irregular-to-flat sections. This is the HBT in-action.

Also, please note, that linearization process described so far, happens in frequency domain. Therefore convolution process, that is rather complex in time domain, is now reduced to simple multiplication of the loudspeaker's frequency response with the corresponding Amplitude Error Function. So, let's do it now.

The result of amplitude linearization is shown on Figure 19.15. The resulting SPL (pink curve) is now flat within the frequency range of interest, and the phase response (yellow curve) is almost smooth within the same frequency range.

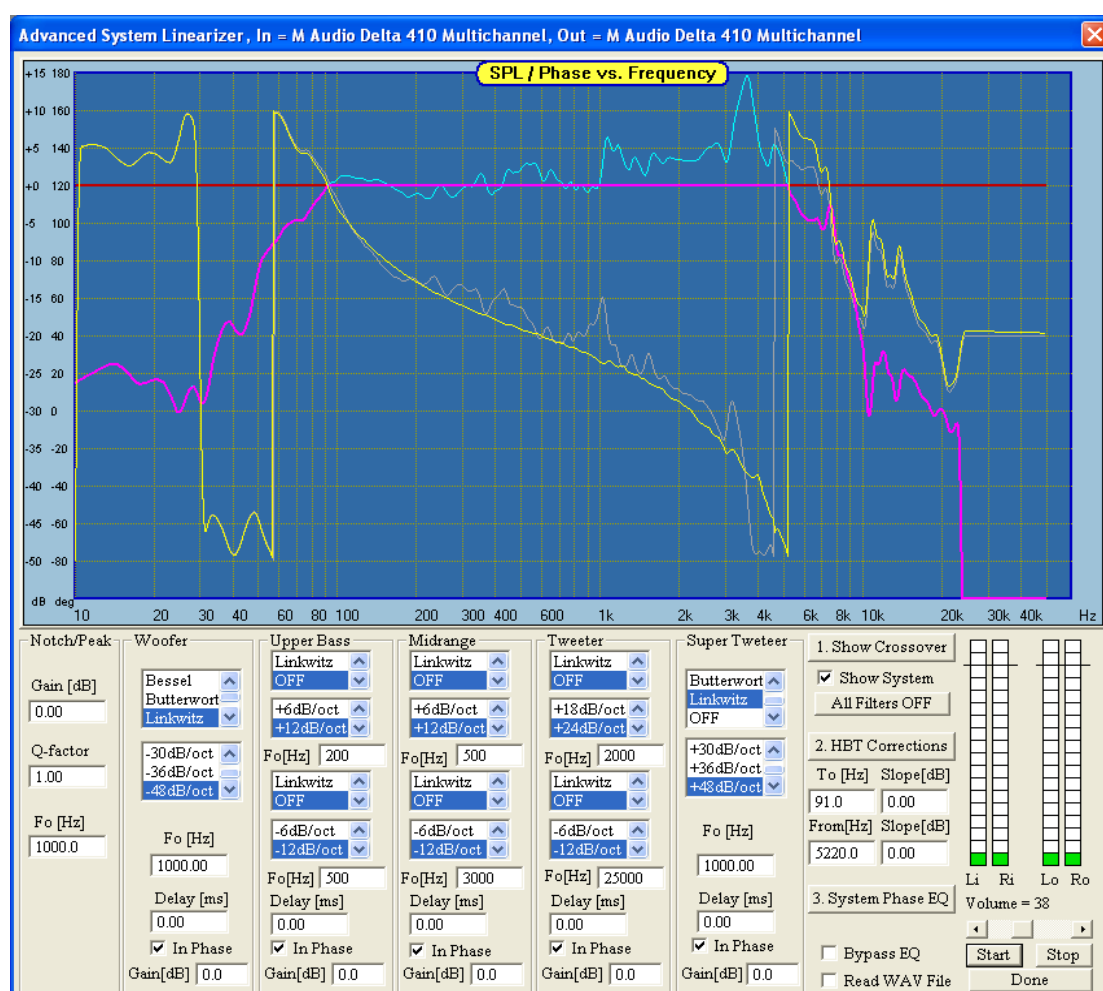


Figure 19.15. Loudspeaker linearized: magnitude (pink), phase (yellow)

At this point of time, we should examine, if we are making any real progress in time domain, as we are making it in the frequency domain. Let's have a look at the square wave being passed through our new, amplitude-linearized loudspeaker system – see Figure 19.16.

The news is good – we are making visible progress. The square wave after being passed through the amplitude-linearized loudspeaker system has now lost almost all ringing characteristics. It's still not a typical flat-top square wave, but we have already removed much of the imperfections.



Inspecting the non-flat shape of the square wave leaving the loudspeaker we can suspect, that still, not all sinusoidal components of the square wave recombine correctly. Otherwise, the square wave would be flat already. Since, we have taken care of maintaining the amplitudes of those sine waves by linearizing the frequency response of the loudspeaker, then the only remaining parameter, that is still causing recombination errors must be the phase of the system.

Indeed, the phase response (yellow curve on Figure 19.15) is smooth, but not flat. If we could somehow “reverse” the phase response, we stand a good chance, that combining the reversed phase with the loudspeaker’s own phase delay will yield a flat phase response of the system.



Figure 19.16. Square wave passed through amplitude-linearizing system.

### Inverting System Phase

A technique that allows us to reverse phase response of an LTI system is based on time-inverting it’s impulse response. In practical terms, you need to reverse time scale of the impulse response. Typical impulse response described in time domain starts with a high peak, followed by it’s decaying tail. Imagine reversing time scale – it would be like the start of the impulse was the tail end of it, and then we arrive at the peak. If you perform this operation as convolution, interestingly, the phase response becomes it’s own mirror image around zero degrees. And that’s exactly what we are looking for. In addition, if we are using FIR ( Finite Impulse Response) techniques to accomplish the convolution, the amplitude response of the process can be decoupled from the inverted phase response and then forced to become flat.

We have now created a perfect phase-reversal device with flat amplitude response - System Inverse Phase Function – see Figure 19.17. Remember, that flat amplitude response requirement is important here, because at this stage, we do not want any more amplitude corrections. We have done this already in the previous stage, using our Amplitude Error Function.

To summarize, we have now 2 transfer functions:

- **Amplitude Error Function** – created by inverted HBT – this is used to linearize the system's SPL and remove ripples from phase response within **SPL frequency range of interest**.
- **System Inverse Phase Function** – this is used to nullify the whole system phase response.

These two corrective functions, in that order, will now be applied to the loudspeaker frequency response. In other words, we multiply loudspeaker's transfer function first, by the **Amplitude Error Function**, then take the result and multiply it by the **System Inverse Phase Function**. The function that we obtain in the end, is the final transfer function of the fully equalized system. We can now try to pass our test square wave through such system and observe the result.




Figure 19.17. System Inverse Phase Function: magnitude (red), phase (yellow).

## Conclusions

It is clearly evident, that the proposed, two-stage equalization technique brings about very good results. The resulting outgoing square wave is almost perfectly recombined from individual sine waves constituting the input square wave. This would be a confirmation, that our equalized system has now flat amplitude within the frequency range of interest, and zero phase response. This loudspeaker can now play square waves.

Second equally important conclusion comes from analysing Figure 6. The Amplitude Error Function makes the magnitude and phase of our loudspeaker smooth, but it does not change macro-characteristics of the phase response, or should we say, it does not “unwrap” the phase. The importance of this is, that the HBT-based, Amplitude Error Function can be equally applied to smooth the magnitude and phase response of non-minimum phase systems, such as multi-way loudspeaker system, complete with crossover. Also, the System Inverse Phase Function inverts the phase of the complete system, as it was measured, and regardless of the trajectory of the phase response. Consequently, the whole two-stage equalization technique is fully applicable to multi-way loudspeaker systems.




Figure 19.18. Top trace – 300Hz square wave run through the loudspeaker alone.  
Bottom trace – the same square wave run through the loudspeaker +equalizer.

Third important conclusion is almost automatic. Linear phase response of the complete system (in our case, zero phase response) results in constant group delay also equal to zero, of the complete electro-acoustical system. As I explained at the beginning of this article, the loudspeaker used for this example was mounted in a vented enclosure. Therefore, the measured phase response already included characteristics ( or phase delay) of the vented box. Since the finally obtained phase response was totally flat, it shows, that the equalization process automatically eliminated group delay associated with the enclosure.

### Description of Controls

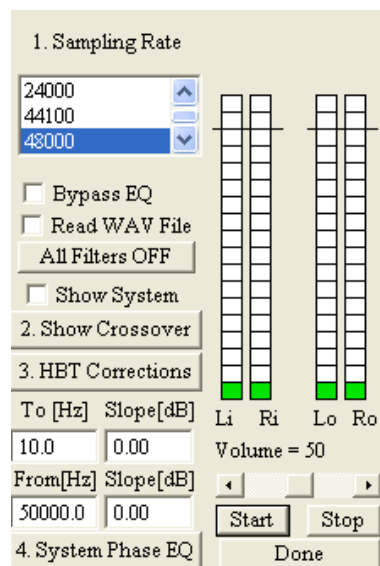


Figure 19.19 Advanced System Phase Linearizer controls.

13. **Sampling Rate** – This list box allows for selection of 44.1 kHz or 48kHz sampling rates.
14. **“All OFF”** – This button turns OFF all filters.
15. **“Show EQ”** – When “checked”, it will enable plotting phase response in addition to amplitude responses.

16. **“Show System”** – Shows the original system response.
17. **“Long FFT”** – Select longer FFT process.
18. **“Read WAV File”** – Check box to allow input sound data from .WAV file.
19. **“Bypass EQ”** – Turns OFF the EQ function.
20. **Volume control** – This slider controls output signal in 0 – 100% range.
21. **“Start Monitor”** – This button activates the Equalizer process.
22. **“Stop Monitor”** – This button terminates the Equalizer process.
23. **“Done”** – Closes the Equalizer screen.
24. There are also input and output level LED indicators.

## Digital Room Equalizer

Digital Equalizer is capable of “non-minimum phase” equalization of the loudspeaker system, and therefore, correcting amplitude and time domain delays of 5-way loudspeaker system. An extension of this concept is equalization of the “speaker-to-ear” room transfer function – the Digital Room Equalizer (DRE). The DRE is a “minimum-phase” equalizer, which in addition, includes a delay function as well. At low frequencies, a typical listening room transfer function is governed by the room modes (resonances). Room resonances, behave as minimum-phase phenomena, and can be dealt with using a minimum-phase equalizer devices. This new feature has been implemented to allow you to “deconvolve” or “equalize” areas of your listening room. The issue of room equalization has been described in many publications, and seven of them are listed below to start your creative thinking process flowing. Possibly start with item 1, as it is mostly descriptive, and then item 4.

1. **The Science of Audio** - a series of lectures by Floyd E. Toole, Ph.D. Vice President Acoustical Engineering Harman International Industries, Inc. (**easy, excellent reading** ).
2. **Digital equalization methods for audio systems** ( Preprint 2598 ), J. Mourjopoulos University of Patras, Patras, Greece.
3. **Loudspeaker and room correction using digital signal processing** ( Preprint 2756 ), Colin Bean, Peter Craven, B&W Loudspeakers, Worthing, England.
4. **Practical Adaptive Room and Loudspeaker Equaliser for Hi-Fi Use** (Preprint 3336) CRAVEN, Peter G. Consultant to B&W Loudspeakers, Steyning, United Kingdom, GERZON, Michael A. Technical Consultant, Oxford, United Kingdom.
5. **Digital room equalization.** Reproduced from Studio Sound, unknown issue 1991. From M. GERZON archive. <http://www.audiosignal.co.uk/Gerzon%20archive.html>
6. **Automated In-situ Frequency Response Optimisation of Active Loudspeakers.** Andrew Goldberg, and Aki Mäkivirta, Genelec Oy, Olvitie 5, 74100 Iisalmi, Finland.
7. **Listening Test Results from a new Loudspeaker/Room Correction System.** Lars G. Johansen (1), Per Rubak (2) (1) Aalborg University, Institute of Electronic Systems, (2) Engineering College of Aarhus, Denmark. Audio Engineering Society, Convention Paper 5323.

**Authors of reference (7) summarized at the end of their presentation:** “ *Our conclusions are quite optimistic. Given we are operating with an acoustically fine and well damped listening room, superb loudspeakers, and a nearly optimal positioning of loudspeakers and listener, it must be expected to be a hard task further to improve reproduction quality by electronic means. Yet, our results show that practically oriented correction techniques do pay off even under these conditions. Measurements and correction design using less optimal listening rooms and cheaper loudspeakers show that our design system works equally well as in the near ideal scenario – revealing even greater objective improvements.* ”.

Room equalization involves knowledge of psychoacoustic and general/traditional room acoustics expertise. Room equalization cannot perform miracles, and is not a panacea for badly reverberating rooms with long decay times. Possibly the best results are obtained by employing some common sense acoustical treatment in the room together with DRE techniques. The improvements obtained are generally quite noticeable.

The goal of room equalization is to improve the perceived quality of sound reproduction in a listening area. The room transfer function is position dependent, therefore perfect equalization within a reasonably large listening area is unlikely, and even an acceptable equalization is typically a far from a “flat line” SPL curve.



It is possible that equalization even degrades the sound in larger listening area if it is based only on a single point measurement and without trough's correction. The frequency response in nearby positions can actually become worse (more irregular) after the equalization designed using only a single point measurement is applied. A well known method to prevent this undesirable outcome is to use an average of SPL responses measured within the listening area.

The actual DRE process involves using MLS measurement system to capture averaged room IR and obviously the DRE screen itself. It is therefore a non-trivial process and you are well advised to take small steps on your learning curve.

### **Practical tips on Room Equalization**

1. We want the equalizer to flatten the “in-room” frequency response, but not at the expense of boosting dips or notches in the frequency response to the point where the boost causes amplifier and speaker overload or massive amounts of the boosted frequency at other listening positions. Fortunately, the ears are much less sensitive to the odd dip in a frequency response than to peaks of similar amplitude. We can take advantage of this phenomenon and be sensible about the boost associated with the troughs. Therefore, DON NOT attempt to fully equalize the troughs or dips in the original room transfer function. The program runs the RDE screen in “**Pixel Editor**” mode, so you can actually “draw” on the screen and correct the measured frequency response by smoothing the dips. By doing so, the “inverse” filter will not have peaks in it's frequency response.
2. Room equalization needs to be of minimum phase type to avoid problems with audible pre-responses when cancellation of the room response is imperfect. One of the remarkable things about minimum phase filters is that the overall frequency and phase response of the filter can be computed, by the HBT, just from a knowledge of the amplitude frequency response alone. In other words, if one knows what frequency response one wishes to correct, the phase response aspects of the filter follow automatically. As you can see, the whole equalization process is performed in the frequency domain rather than in the time domain. This allows you to correct the measured room frequency response to partially eliminate dips in the response.
3. Frequency response in higher frequency range (the statistical region) should be smoothed out by applying smoothing at capturing room IR stage. Typically, the SPL in the statistical region just resembles “noise”, and needs to be smoothed out before application of the DRE. Standard acoustical treatment of the room is also very helpful in this frequency range. Initially, you can simply “switch off” this region from DRE action. This is accomplished by applying a flat LP tail, starting from say 500Hz (or lower) to the measured room SPL. Since the measured SPL is now flat above 500Hz, the DRE filter will also be flat – no room correction above 500Hz. Gradual experiments into the statistical region would allow you to determine how much benefit you can gain in mid-high frequencies range using the DRE. The benefits are likely to be tied to individual room's acoustics.
4. Another advantage of being able to derive minimum phase equalizers just from a measurement of the amplitude of the frequency response is that you can measure the SPL response at several points across a listening area, and equalize for the averaged frequency response at all these points, rather than for the frequency response measured at just one point. This gives better results across the whole listening area, at the expense of less good results very near one single point. In typical domestic listening rooms, averaging for six points across, say a sofa, can give useful room equalization across the listening area.  
One advantage of averaging the measured response at several room points is that this tends to “fill in” measured dips in the frequency response, so you have less headache, trying to decide how much of the dip you need to leave untreated.
5. It is left entirely to your discretion to decide where to, on the frequency scale, apply tails for the HBT. Just a few points to remember:



- Application of tails should not result in frequency response with sharp “corners”. This type of frequency response is mathematically correct, but is not natural to realizable audio components.
- Do not attempt to create bass boost or treble boost in the overall DRE filter frequency response, especially at the low/high ends of the frequency range.
- Neutral gain (around 0dB) will be accomplished if the measured room SPL is around 90dB. Please monitor the “In” and “Out” LED level indicators during the DRE operation.

If you are a “seasoned acoustician” you should be able to run DRE with filtering function averaged over 5-10 locations and show the results of the DRE operation as a Cumulative Spectral Decay plots.

If you are new to this technology, you may want to initially concentrate on single measurement approach to the design of the inverse filter DRE and verifications measurement. This way, it may be easier for you to grasp all issues involved.

### Description of the DRE Process

The complete DRE process is broadly described below. Some of the steps are quite involving and will be detailed in the paragraphs that follow.

1. Select what DRE channel (Left / Right / Both) is being measured by “checking” the box in DRE control dialogue.
2. Make spatially averaged SPL measurements in your listening area, for Left / Right or Both channels. Use MLS to perform the measurements and averaging.
3. Transfer the averaged SPL into DRE using the Curve Arithmetic Buffer.
4. Remove excessive troughs / dips from the averaged SPL using DRE Pixel Editor screen.
5. Select HBT tail frequencies for low- and high-end of the frequency range. This may be different for each channel.
6. Select the desired slope of HBT tails, so that the final DRE filter has flat or rolling slopes at low- and high frequency extremes. Again, this may be different for each channel.
7. Plot all relevant DRE curves for visual inspection.
8. Run DRE.

Extensions to this process involve modifications to the target filter transfer function by means of tilting the whole filter or adding ripples to the transfer function. It needs to be emphasized here, that the DRE process is a “challenge”.

**We recommend, that you first practice without any amplifiers/speakers connected to the DRE, just to make sure, that you are comfortable with the process.**

### Making average SPL measurements.

You can use the MLS screen to make all room IR measurements. However, the process commences by selecting **Left / Right OR Both channels**, for which room transfer function is being currently measured. This selection is performed from the DRE control box, under Channel column. Please note, that you can move the results of Curve Arithmetic from the **Buffer** to the **Left / Right / Both channels of the Digital Room Equalizer**. First, please select where you want the results to go, by selecting appropriate check box on the **DRE** control dialogue box. Then select “**Save Buffer**” to **Curve 0**, on the MLS “Post-Processing” tab. This actions will save the Buffer content into driver data space (Curve 0) and also into the DRE data space, as requested from DRE dialogue box.

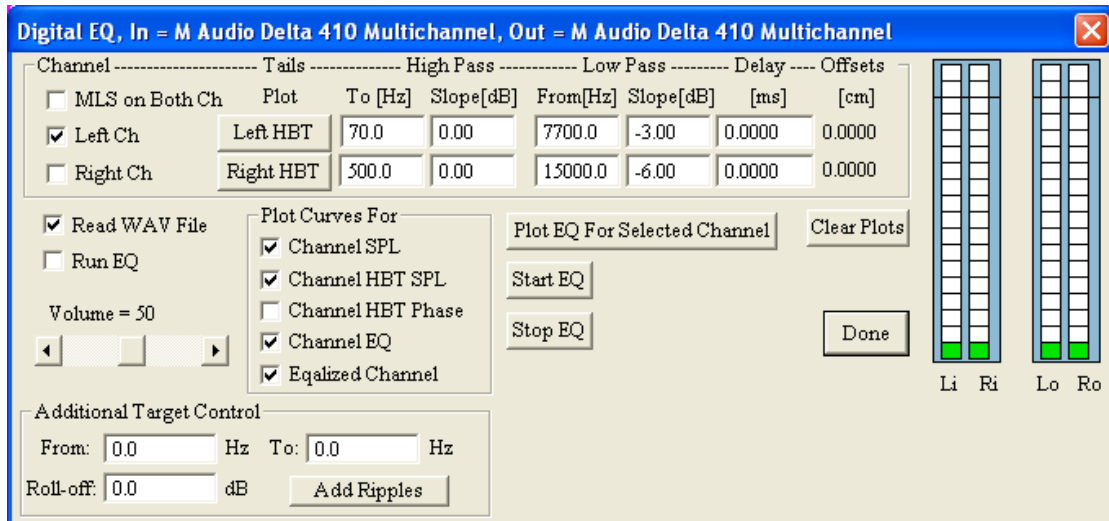


Figure 19.20 DRE control box.

1. Measure SPL at several (5-10) nominated locations. It would be recommended to save each individual measurement as IR-file (check “**Save Input Buffer**” box in the “**Acquire IR**” TAB) for future use. It is also useful, to label each measurement as a file name when saving the IR to a file. This will allow you to “mix-and-match” locations at will for different average locations at later stage.




Fig 19.21. MLS is used to measure room transfer function

2. Load each IR file into MLS screen and perform “**IR->SPL**” transformation. Store each resulting SPL curve into the five “**Post-processing**” locations.

- When all five curves have been processed, perform “Add” operation. If the resulting “averaged” curve ended up to high on the SPL scale, “Subtract” the required number of dB’s from it. Up to five files can be included in the averaging process at one time, so the averaged curve may end-up being 13-14dB higher. You can average more locations by performing the above process sequentially.

You should aim at having the spatially averaged SPL curve around 90dB level. This is a 0dB gain for the DRE.

- The results need to be moved into the “driver” data space (from Buffer to curve 0) in order to show up in the DRE screen. The rest of the operations can be performed from the DRE dialogue box.

### Removing troughs/dips.

When applied to the DRE, the concept of the “inverse filter” simply means “mirror image” of the measured SPL. The measured SPL dips become filter’s peaks and vice-versa. The obvious problem here is, that the 30-40dB cancellation troughs end-up as massive peaks and will damage the components of your audio chain. Clearly, the dips need to be reduced.

When you open the DRE window, the screen is automatically set to “**Pixel Editor**” mode. You can actually “draw” on the screen and correct the measured frequency response by smoothing the dips. By doing so, the “inverse” filter will not have peaks in it’s frequency response. As a first approximation, you may aim at having the deep peaks reduced to 5-6dB shallow, rounded dents in the measured SPL curve.

The process of “drawing” the SPL using the Left Mouse Button and cursor keys has been described in early chapters of this manual. An example of an undesired dip in the averaged room transfer function is shown below. Here, a large dip at 200Hz (pink curve) will result in corresponding peak in the DRE filter response (green curve).




Fig 19.22. Undesired dip in the averaged room transfer function






Fig 19.23. The dip is removed from the averaged room transfer function

However, if the dip is removed from the averaged room transfer function (pink curve), the peak on DRE filter response also disappears – see Figure 19.23 above.

### Selecting HBT parameters

Selecting the HBT parameters is quite an important issue. Let's assume, that the spatially averaged room transfer function looks like the pink curve on Figure 14.16 below.

**We know already, that anything above 90dB line will be attenuated and anything below 90dB line will be amplified by the DRE.**

Upon visual examination of the pink curve, we can determine, that we can gain a bit of bass EQ boost at the low-end (about 6dB), if the HBT high-pass filter (the "To" frequency) is selected at the short leveled section of the pink curve at around 70Hz. This is because this section is about 6dB below 90dB.

The slope for high-pass filter is selected to be flat (0dB). This way, all frequencies below 70Hz will be amplified by the constant 6dB factor. Please do not attempt to select positive slopes (positive value in this data field) here. This would result in DRE trying to unnecessarily boost the lowest end of the frequency spectrum. It would be advisable to enter negative value here to allow the DRE filter to roll-off towards the frequency extremes.

On the high-side of the frequency range, we would like the DRE to cancel the high peak at 3.5kHz. Therefore, the HBT low-pass (the "From" frequency) is selected as 7.7kHz. This is where a small peak is residing.

The slope for low-pass filter is selected to be rising -3dB (yes, negative value). This way, all frequencies above 7700Hz will be attenuated by the constant -3dB factor in the final DRE filter response. Please do not attempt to select positive slopes (positive value in the data field) here. This would result in DRE trying to unnecessarily boost the high end of the frequency spectrum. It would be advisable to enter negative value here to allow the DRE filter to roll-off towards the frequency extremes.




Fig 19.24. Example of HBT (blue plot).

Now, we are ready to confirm the above assumptions and plot the plot the HBT(blue curve) + DRE (green curve) filter responses – see Figure 19.25 below. It is observable, that the DRE filter is a mirror image of the HBT response, centered around 90dB level line.




Fig 19.25. DRE (green curve) centered around 90dB.

Finally, we can review the TOTAL frequency response of the DRE + room after equalization. This situation is depicted on Figure 14.18 below.



Please view Figure 14.18. The pink curve is the “un-equalized” room transfer function. The thin black curve is the corresponding, equalized room transfer function.

The example presented in this paragraph showed fairly smoother low-end room transfer function. There was no significant dips, therefore we did not have to use pixel editor to account for the excessive dips. Therefore, the final, simulated, equalized room transfer function looks very smooth and even within the pass-band. Please also note, that we have gained bass extension from 100Hz down to 65Hz. Real-life results will be compromised by the need to account for the dips, so the final equalized room transfer function will not be a flat line.




Fig 19.26 Equalized (black curve) room transfer function

## DRE Controls

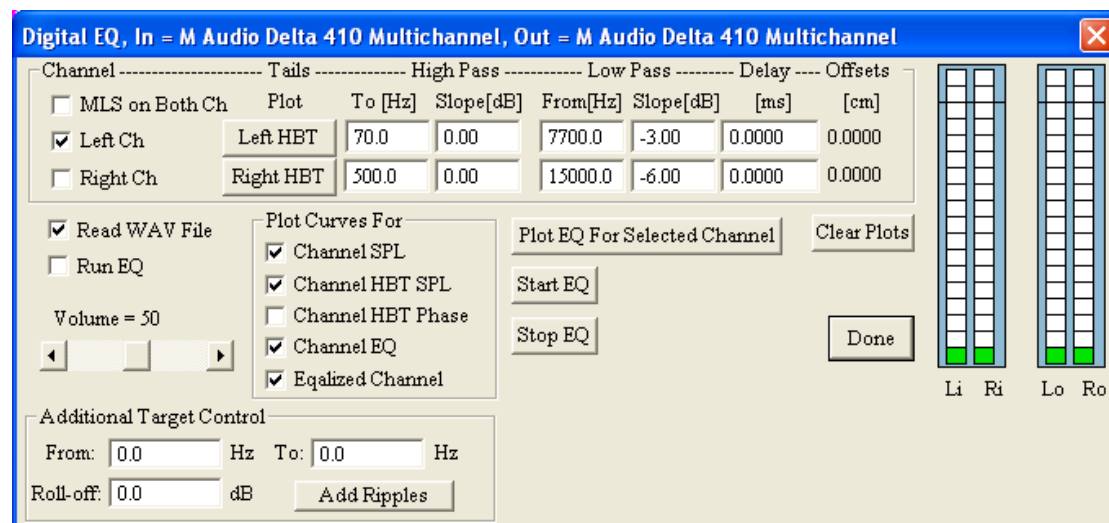


Fig 19.27 DRE controls.

### Channel group

1. **“MLS on Both Ch”** – check box to select the averaged SPL room transfer function measurement from MLS to be stored in both channels.
2. **“Left Ch”** - check box to select the averaged SPL room transfer function measurement from MLS to be stored in the Left Channel.
3. **“Right Ch”** - check box to select the averaged SPL room transfer function measurement from MLS to be stored in the Right Channel.

### Tails group

1. **“Plot Left HBT”** – button to plot left channel HBT.
2. **“Plot Right HBT”** – button to plot right channel HBT
3. **“To”** – data field to enter cut-off frequency for high-pass section of HBT for left / right channels.
4. **“Slope”** – data field to enter slope speed in dBs, for high-pass section of HBT for left / right channels.
5. **“Delay”** – data fields to enter delays for left / right channels.
6. **“Offsets”** – displays delay data converted to centimeters.

### Additional Target Control group

1. **“From”** – data field to enter starting frequency for the roll-off.
2. **“To”** – data field to enter ending frequency for the roll-off.
3. **“Roll-off”** – data field to enter steepness of the roll-off.
4. **“Add Ripples”** – button to activate control box for entering 3 ripples onto the target DRE curve.

### Plot Curves group

1. **“Channel SPL”** – check to display averaged SPL room transfer function
2. **“Channel HBT SPL”** – check to display HBT SPL of the room transfer function
3. **“Channel HBT Phase”** – check to display HBT Phase of the room transfer function
4. **“Channel EQ”** – check to display DRE transfer function
5. **“Equalized Channel”** – check to display DRE transfer function + averaged SPL room transfer function.

**“Read Wav File”** – check box to select signal source as WAV file directly from the HD.

**“Run EQ”**- check box to ENABLE / DISABLE the DRE.

**“Volume”** – slider to adjust output volume.

**“Plot EQ For Selected Channel”** – Button to plot any of the selected curves.

**“Start EQ”** – Button to start DRE play.

**“Stop EQ”** – Button to stop DRE play.

**“Clear Plots”** – clears the display.

**“Done”** – closes the DRE.

**“Li / Ri”** – Input level indicators for Left / Right channels.

**“Lo / Ro”** – Output level indicators for Left / Right channels.