REW V5.20 Help

REW Help Contents

[REW homepage](#)

Welcome to REW

REW (Room EQ Wizard) is a Java application for measuring room responses and countering room modal resonances. It includes tools for generating test signals; measuring SPL; measuring frequency and impulse responses; generating phase, group delay, spectral decay plots, waterfalls, spectrograms and energy-time curves; generating real time analyser (RTA) plots; calculating reverberation times; displaying equaliser responses and automatically adjusting the settings of parametric equalisers to counter the effects of room modes and adjust responses to match a target.

REW uses a logarithmically swept sine signal for its measurements. This is much faster than manual measurements, more accurate, less likely to suffer from clipping at resonances, less sensitive to system non-linearity than MLS and allows the impulse response of the room to be determined, which in turn is the basis for many additional features. When using the Real Time Analyser displays REW can generate Pink Periodic Noise sequences for much better visibility of low frequency behaviour than obtained using random Pink Noise without the need for lengthy averaging.

The REW home page is at www.roomeqwizard.com. For hints, tips and help visit the REW forum at www.avnirvana.com.

Requirements

- REW can be used on Windows XP/7/8/8.1/10, macOS 10.11 or later and Linux
- Minimum screen resolution: 1024 x 768
- Minimum RAM: 2 GB, 4 GB or more recommended
- The Windows and Linux installers include a private Java runtime for use by REW if a Java 8 runtime is not already installed. Do not install a headless version of the Java runtime, they lack audio I/O
- The macOS download includes a private Java runtime for use by REW. Java does not need to be installed. Microphone access permissions are part of the code signature and should be automatic. To force a mic access permission prompt for REW run

  ```
  tccutil reset Microphone
  ```

  from a terminal before starting REW
- RS232 serial communications (only used to communicate with TAG McLaren Audio AV32R DP and AV192R AV Processors) only function on Windows 32-bit systems.
- Midi communication (used to set filters on Behringer BFD Pro DSP1124P and FBQ2496 equalisers) is supported on Windows. macOS should also support Midi.

Windows XP, Vista, 7, 8 or 8.1

When installing REW on Windows XP, Vista, 7, 8 or 8.1 you may encounter a somewhat cryptic error message:
"The program can't start because api-ms-win-core-timezone-l1-1-0.dll is missing from your computer. Try reinstalling the program to fix this problem.".
This is caused by the Java runtime using components from the Windows Universal C Runtime which are not present on older Windows installations. To fix it see this Microsoft knowledge base article, KB2999226 at https://support.microsoft.com/en-us/help/2999226/update-for-universal-c-runtime-in-windows.

Linux installation

- Open a terminal window in the folder where the installation file was downloaded

- After downloading the installation file set permissions to allow it to run and install it as follows (example is for 5.19):

```
chmod 777 REW_linux_5_19.sh
```

```
sudo ./REW_linux_5_19.sh
```

- If your terminal working directory is not where the installer was downloaded include the path to the downloaded file. For example, if the terminal is open in your home directory and the installer was downloaded to the Downloads folder put Downloads/ in front of the installer filename.
- Start REW by running

```
roomeqwizard
```

Diagnostics

REW saves diagnostic logs in the user's home directory, the location is displayed in the **Help → About REW** window. The logs contain information from the last 10 startups, including any error messages or warnings that may have been generated. On macOS the logs are in the user library folder, which no longer appears in Finder by default. To display it either browse to ~/Library/Logs or open your Home folder in Finder (Shift Command + H), select View > Show View Options then select the checkbox to show the library folder.

REW saves a temporary copy of each measurement made in the logs folder. A normal shutdown deletes those temporary copies. If REW did not shut down normally it will warn of temporary file existence on startup and offer to load them.

The startup preferences for REW on Windows systems are stored in this registry key:

```
HKEY_CURRENT_USER\Software\JavaSoft\Prefs\room eq wizard
```

On macOS the startup preferences are stored in:

```
~/Library/Preferences/com.apple.java.util.prefs.plist
```

under a key called room eq wizard.

Preferences can be deleted from **Preferences → Delete Preferences and Shut Down** or by running the REW uninstaller.

No Warranty

THIS SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE

Acknowledgements

REW uses the JTransforms pure Java FFT library available at JTransforms

ASIO interfaces are supported using Martin Roth's JAsioHost, available at JAsioHost

REW's installers are built with the Install4J multi-platform installer builder, available at ej-technologies

Thanks to Gerrit Grunwald for ideas and elements from his Steel Series components, available at Harmonic Code

REW uses the Hola Multicast DNS Service Discovery library available at Hola

REW uses the Java WebSocket client available at Java WebSocket

Copyright © 2004-2020 John Mulcahy All Rights Reserved

Get started with REW

Send feedback to feedback@roomeqwizard.com

Help Index

Getting Started with REW

REW is a software package that measures the transfer functions of acoustic systems and displays the corresponding frequency, phase and impulse responses and various quantities derived from them. If that sentence made you ask yourself *what on earth is he talking about?* It would be worth taking a few minutes to read through the introduction to Signals and Measurements, which explains the basic concepts. Even if you are already familiar with the terminology a quick browse through that introduction may be helpful.

How REW makes Measurements

REW uses a logarithmic sine sweep method and a basic explanation follows. If you'd rather not know, you can skip the explanation.

You can find out a great deal about the log sweep method, and the various alternatives, by reading the paper "Transfer-Function Measurement with Sweeps" by SWEN MÜLLER and PAULO MASSARANI.

To make a measurement we need a sound source (a loudspeaker or subwoofer) and a microphone (SPL meters contain a microphone and many can be used instead of a mic). A logarithmic sweep signal is sent to the source, starting low and steadily increasing to a higher frequency. A logarithmic sweep is one that takes the same time to double in frequency. It takes the same time to go from 40 to 80Hz or 4kHz to 8kHz as from 20 to 40Hz. The mic picks up the sweep, sensing the sound that travels directly from the source and all the sound that first bounces off the room surfaces. Some sound paths may bounce off several room surfaces before reaching the microphone.
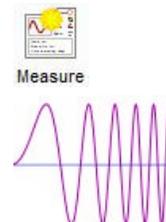
Once the sweep completes, analysis starts. A process called "Fast Fourier Transform" (FFT) is used to calculate the phase and amplitude of the frequency spectrum that made up the sweep. By comparing the amplitudes and phases of the spectrum of the received microphone signal with those sent in the sweep we can work out how each frequency has been affected by the loudspeaker *and room* we are measuring. This is called the "Transfer Function" of the room from the location of the source to the location of the mic. Note: A different source position or different mic position will have a different transfer function. In other words, *our measurement is only valid for one specific source and mic position.* Having worked out the transfer function we can use an "inverse FFT" to get from the frequency amplitude and phase information to a time signal that describes the way *any* signal is changed when travelling from the source to the mic. That time signal is called the "impulse response" - like the transfer function it is derived from, it is only valid for one specific source and mic position.

The impulse response is actually exactly the same signal we would see if we could emit a very short but loud click at the source position and record what the mic picks up afterwards ("very short" meaning lasting just the time of 1 sample at the sample rate we are using for our analysis, so for a 48 kHz sample rate that would be just 1/48,000 of a second which is 21 millionths of a second). You might ask why we don't just use a click then. One difficulty is that the click, because it is so short, needs to be extremely loud for us to be able to pick up what happens after the initial click over the background noise of the room. We could no longer use a speaker to generate that, we would need something like a starting pistol or to pop a balloon. We would also need a mic that could cope with both the extremely loud click itself and the much quieter echoes of the click produced by the room. You are likely to find that your family and neighbours are not that keen on you repeatedly firing a pistol to figure out what your room is doing, and even if they put up with it your results would not be as good as using a sweep. To be technical about it, you can achieve a much higher signal-to-noise (S/N) ratio by the sweep method. The S/N is determined by the background noise level and by how much energy is in the test signal, which in turn depends on how loud the signal is and how long it lasts. An impulse is extremely short, a few millionths of a second, so to get any significant energy it needs to be extremely loud. A sweep can last for many seconds, so even at a modest volume its total energy can approach as much as a million times more than an impulse.

Once the impulse response has been obtained, it can be analysed to calculate information about how the room behaves. The simplest analysis is the FFT, to show the frequency response between the source and mic positions. However, we have some control over it. Altering which part of the impulse response is analysed by the FFT changes what aspect of the room's response we see. The early part of the impulse response corresponds to the direct sound from the source to the mic, the shortest path between them. Sound that has bounced off the room's surfaces has to travel further to reach the mic, which takes longer, so the later parts of the impulse response contain the contributions of the room. "Windowing" the impulse response to look at only the initial part shows us the frequency response of the direct sound with little or no contribution from the room. A window that includes later parts of the response lets us see how the room's contribution alters the frequency response. The ability to separate the contributions of the direct and later (reflected) sound is a key difference between the frequency response derived from an impulse response and one from a Real Time Analysis (RTA), which can only show the total combined response of source and room.

Other information we can get from the impulse response include a "waterfall" plot, which is generated by moving a window in steps along the response and plotting the various frequency responses in order to produce a 3D picture of the way the response changes over time, and the room's "RT60" data, which is the time it takes sound in various frequency bands to decay by 60dB (or 1,000 times smaller than the original level).

Equipment Needed

The first requirement is a way to capture the test signal. There are a few options:

- A USB microphone that comes with a calibration file. Such a mic can be used for low frequency or full range measurements. If the cal file has sensitivity data in a format REW recognises, it can also act as a calibrated SPL meter.
  The miniDSP UMIK-1 is recommended and has calibration data in an REW-friendly format, see www.minidsp.com or Amazon.

- An alternative to a USB mic is an SPL meter with a line level analogue output. The Radio Shack (RS) meter is perfectly adequate for low frequency room acoustics work, either the analogue or digital display version. The Galaxy CM-140 meter has better tracking of the C-weighting curve and better behaviour above subwoofer frequencies than the RS meter, but is more expensive.

- A final option is an analog microphone, but most mics will require a preamplifier to produce line level and to supply the mic with phantom power. For full range measurements the mic response must be calibrated for accurate results. An SPL meter is still ideally required to provide a reference SPL figure against which to calibrate REW's SPL display.

- A tripod to support the mic or meter (hereafter just called "mic"). Small movements of the mic can result in large variations in the measurements, for repeatable results a means of supporting the mic for a prolonged period is essential. For low frequency measurements (below a few hundred Hz) the mic can be set pointing straight up. This avoids having to move it to measure different speakers and makes it easy to read the display on a meter. Use a "90 degree" mic calibration file if pointing the mic upwards. To make measurements at higher frequencies it is best to point the mic directly at the speaker being measured. Use a "0 degree" mic calibration file in that case. In both cases the mic should be placed at ear height in your usual listening position.

- If you are using a USB mic your computer's headphone output can generate the test signals REW uses without an external interface. If you are using an SPL meter or a mic with a preamp then an audio interface (internal or external) with line inputs and line or headphone outputs is required. **Note: Older PC and laptop mic-only inputs are NOT suitable and should not be used** as they had too much gain and most suffered from high noise levels and limited bandwidth. Combination mic/line inputs can be used successfully. Inexpensive or built-in interfaces are typically adequate, a reference measurement of a loopback connection can be used to remove the interface's frequency response from the measurement.

- Cables to connect from your SPL meter or mic preamp's output to your interface (if you are not using a USB microphone) and from the interface's line or headphone output to an input on your AV processor or equaliser. The leads need to be long enough to reach from your computer to your listening position (where your mic will be placed) and to your AV processor or equaliser. If your interface has phono (RCA) connectors phono-phono leads will be needed, if the interface has 3.5mm (1/8") sockets you will need a pair of **stereo** jack plug to stereo phono plug leads (also called Y adaptor cables) or stereo audio adaptors , see images below.

If you use male to male Y adaptor cables you will also need two phono socket-to-phono socket adaptors (also called RCA phono plug couplers, see image below) to connect to the leads that run to your SPL meter and AV processor. (RCA Y-cables come in all formats: M->2M, M->2F, F->2F, F->2M and a selection often comes in handy when making measurements.)

If connecting to a BFD Pro DSP1124P or FBQ2496 you will need phono to mono 1/4" jack plug adaptors like that pictured below.

- An interface to your equaliser if using REW's RS232 or Midi communication capabilities to set up the equaliser's filters. Requirements for communicating with TAG McLaren AV processors are detailed here, for BFD Pro DSP1124P or FBQ2496 (Midi interface) look here.

Connections

The overall setup for measuring when using an SPL meter is shown below. If you are using a USB microphone you do not need to make any connections to the interface inputs, just connect the mic to a USB port on your computer.
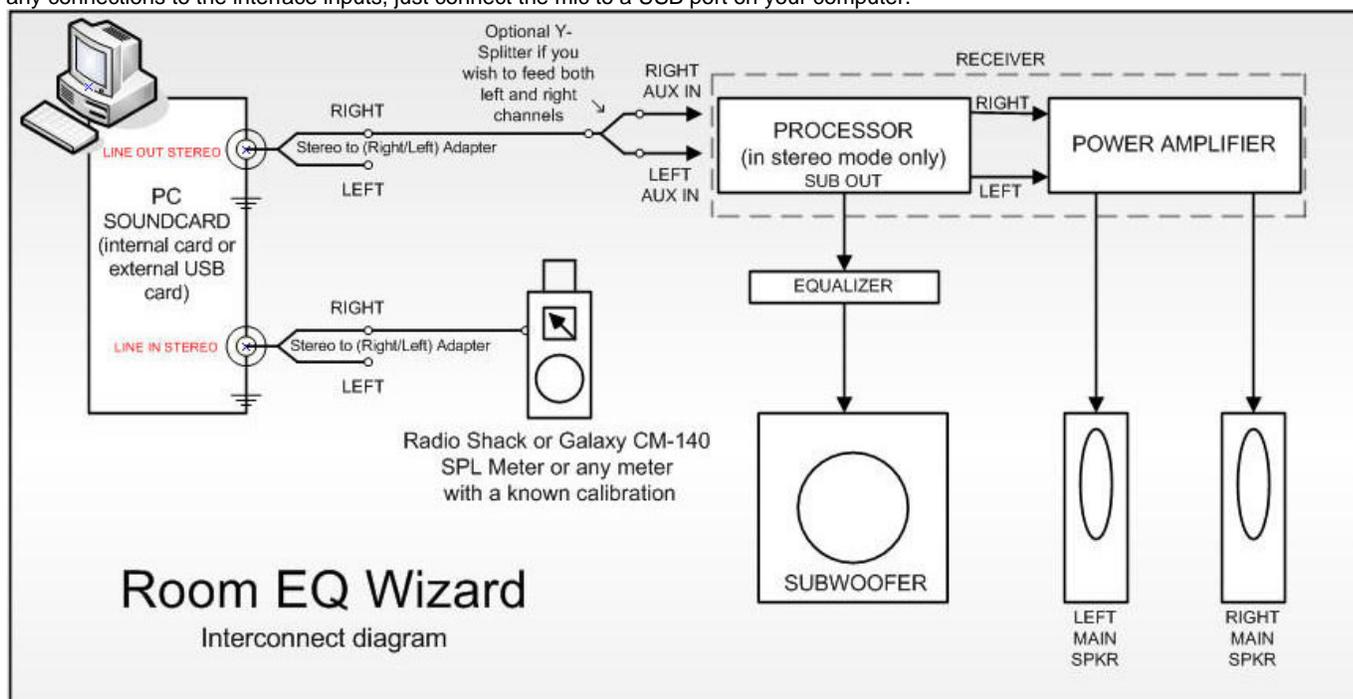

- When not using a USB mic, one of the interface's input channels is used to measure the sound pressure signal from your mic or meter. It must be connected to the meter's or mic preamp's analogue output. The default is to use the Right input, but either input can be used. A control in the Soundcard Preferences tells REW which input to listen to.
- One of the interface's output channels (typically the Right) must be connected to an input channel on your AV processor or to the input of your equaliser. Connecting to your AV processor allows you to make measurements that will show the response of main speakers as well as the subwoofer, and to view the integration between subwoofer and main speakers. The effects of your AV processor's bass management can be included in the measurements. Connecting to the left or right channel of an analog input will allow the corresponding main speaker and the subwoofer responses to be measured. Turn off or disconnect the main speaker or the sub to exclude them from a measurement.
- If your computer and AV processor both support audio over HDMI you can also use an HDMI connection to carry the audio. The HDMI audio driver needs to be selected as the output device under Preferences -> Output Device. If you are on Windows and want to access multichannel outputs over HDMI you will need to use an ASIO wrapper like FlexASIO or ASIO4All to access the outputs and inputs.
- The other input and output channels do not need to be used for basic measurement. The response of the interface itself can be compensated for by taking a reference measurement with the output connected directly to the input and configuring REW to subtract that measured response from subsequent room measurements. However, it is also possible to use a loopback connection from the interface's other output to its other input as a timing reference for REW to automatically compensate for the time delay in the interface and operating system when it makes a measurement. A timing reference is required to compare phase or time delays between measurements or for getting speaker delay settings correct in multi-channel systems. If you wish to do this you may need an additional RCA phono plug coupler to make the loopback connection. Whether REW uses

one channel as a timing reference is controlled by a selection on the Measure dialog. When the input is via a USB microphone instead of an interface an acoustic timing reference can be used. If an acoustic timing reference is used REW will generate a timing signal on the output that has been selected to act as the reference before it generates measurement sweeps on the channels being measured. The timing signal is a high frequency sweep to allow accurate timing so a subwoofer cannot be used as the reference channel. Measurements will have a time delay that corresponds to the difference in their distance from the microphone compared to the distance of the reference speaker - if the reference speaker is further away the delay would be negative. **Note** that multiple sweeps cannot be used with an acoustic timing reference.

Equaliser Connections

If using an equaliser(such as BFD Pro DSP1124P or FBQ2496) to optimise your subwoofer's response it should be connected between your AV processor's LFE/Sub output and the subwoofer's low level input. *For a BFD Pro the operating level switches on the rear panel should be pressed in to select the -10 dBV range.*
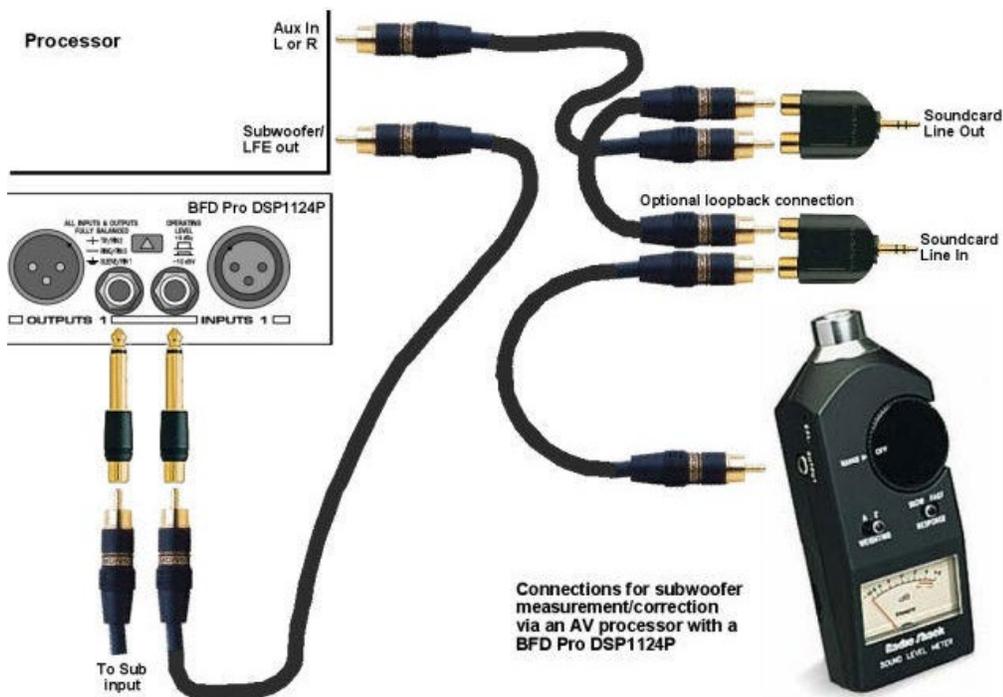
If the AV processor input being used has an anti-clipping feature (automatically reducing its sensitivity if it detects large signals) this should be turned off, as it could shift the measurement levels. The sensitivity of the input should ideally be set to 0.5V, though this is not critical.

The TAG McLaren AV32R DP and AV192R allow the test signal input to be routed to any speaker output via the Test Signal entry within the TMREQ filter menus for each speaker, which is handy for measuring other speakers (see this note for details). They seem to be the only AV processors with such a facility, other processors may have 5.1 or 7.1 analog inputs that can be used to similar effect, but in some cases bass management will not be applied to such inputs, limiting the ability to check sub/main speaker integration.

SPL Meter Range

If you are using an SPL meter as the input the meter's range should be set to the value normally used for speaker level calibration and must not be altered while using REW. If you are using the Radio Shack meter, select the 80dB range if you calibrate your system at 75dB (this is the standard level recommended by Dolby[TM]). Set your meter to C weighting and "slow". Select the **Input device is a C weighted SPL meter** option in Preferences -> Cal Files.

REW Overview

Having assembled the equipment required we can look it how REW itself is organised for making and analysing measurements in the REW Overview.

Help Index

Signals and Measurements

"What does all this stuff *mean* anyway?"

To make sense of the measurements you can take with REW it is helpful to have an understanding of what the measurements are. This topic gives an overview of the basics of signals and measurements and explains how the various graphs in REW are generated and how they relate to what we have measured.

Signals, Sample Rate and Resolution

The first thing to understand is what a signal is, at least in the context of making acoustic measurements. The signals we are interested in are sounds recorded through a microphone or SPL (Sound Pressure Level) meter. The sound pressure generates electrical signals in the mic/meter which are captured by our audio interface. The interface takes measurements of the electrical level at its input. Each measurement is referred to as a **sample**. How often it takes its samples is controlled by the **sample rate**. REW supports a range of sample rates depending on the capabilities of the interface, common rates are 44.1kHz or 48kHz - which means the interface is capturing the level at its input either 44,100 or 48,000 times every second. Three seconds of a signal sampled at 48kHz means a sequence of 3*48,000 = 144,000 measurement values. The highest frequency that can be captured at any given sample rate is half the rate - we need at least two samples for each cycle of the frequency to reproduce it. At 48kHz sampling that means the highest frequency we can capture is 24kHz. Frequencies higher than half the sample rate would cause **aliasing**, they would appear to be lower than they actually were. For example, a 25kHz signal sampled at 48kHz would actually look like a 23kHz signal. To prevent this, the inputs of the interfaces have anti-aliasing filters that try to block signals higher than can be captured, but they are not completely effective so we always need to consider the frequency content of the signals we are trying to capture.

The **resolution** of the interface measurements is typically either 16 bits or 24 bits. 16 bit resolution is the same as used on CDs, and is the resolution REW supports. Having 16 bit resolution means the individual measurement values can range from -32768 to +32767 (numbers that can be represented with 15 binary digits, plus a 16th binary digit to store the sign of the number). Rather than use the measurement numbers directly, it is convenient to refer to them in terms of how close they are to the largest number, which is referred to as **Full Scale** and abbreviated as **FS**. The full scale values are -32768 and +32767. The smallest non-zero measurement value is 1, which as a percentage of full scale is 100*(1/32768) or approximately 0.003% FS. Anything smaller than that is seen by the interface as zero. The full scale value will correspond to a certain voltage at the interface input - that is usually around 1 Volt. Soundcards that have higher resolution, such as 24 bit, usually have the same maximum input voltage (around 1 Volt) but can use a wider range of numbers to measure the voltage. For a 24-bit interface the full scale measurement values are -8388608 and +8388607. That still is only 1 Volt (typically), the largest input voltage has not changed, but the 24-bit interface has higher resolution - the smallest value it can detect is 100*(1/8388608) percent of full scale, 0.000012% FS. It is with the very smallest signals that higher resolution has benefits. The full scale value is often treated as corresponding to a value of one, and everything below full scale as being the corresponding proportion of one, so half full scale would be 0.5 and so on.

Clipping

If the signal gets larger than the full scale value the interface is unable to follow it - the measurement value cannot get higher than full scale no matter what is actually happening at the input. When the signal has gone beyond the range the input can measure it is said to have been **clipped**. Clipping shows up in input signals as flat parts of the response. If the clipping happens at the interface input it will be at +100% FS or -100% FS and REW will warn you, but sometimes clipping can happen before the signal gets to the interface (in a mic preamp whose gain is set too high, for example). In that case the measurement values may never reach the interface's FS levels but the signal is clipped nonetheless. Clipping must be avoided when measuring,

because the captured signal no longer represents what was actually happening at the input and that corrupts the measurement.

Viewing Signals

One way to look at signals is to plot the measurement values against time. When captured signals are plotted in REW on the **Captured** graph they are shown as **% FS**, a signal that reaches 100% FS is the largest the interface can capture. An example of an REW Scope plot is shown below, displaying a sweep signal REW has generated and (in red) the resulting signal captured from a microphone.
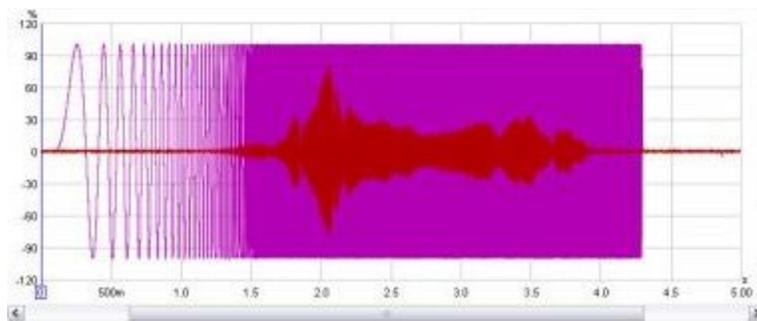

We are usually interested in more than just the sample values. The frequencies that make up the signal may also be of interest. The range of frequencies that make up a signal is called its **Spectrum** and we can calculate them using a Fast Fourier Transform or **FFT**. The FFT works out the amplitudes and phases of a set of cosine waves that, when added together, would give the same set of measurement values as the time signal. The amplitudes and phases of those cosine waves are a different way of representing the time signal, in terms of the frequencies that make it up rather than its individual measurement values. The amplitudes are easy to understand, a larger amplitude means a bigger cosine wave. The phases indicate the starting value for the cosine waves at the time of the first sample in the sequence that was measured. A phase of zero degrees would mean the starting value was amplitude*cos(0) = amplitude. A phase of 90 degrees would mean a starting value of amplitude*cos(90) = 0. We are more often interested in the amplitudes than the phases, but we shouldn't forget about the phases entirely - they contain half the information about the shape of the original time signal.

When an FFT is used to calculate the spectrum it uses a set of frequencies that are evenly spaced from DC (zero frequency) up to half the sample rate (the maximum that can be properly represented). The spacing depends on the length of signal we analyse in the FFT. FFT calculations are most efficient when the signal lengths are powers of two, such as 16k (16,384), 32k (32768) or 64k (65536). To calculate a 64k FFT from a signal that is sampled at 48kHz we need 65536/48000 seconds of the signal, or 1.365s. The frequencies would be spaced at 24000/65536 = 0.366Hz. If the FFT were generated from 16k samples the frequencies would be 1.465Hz apart. The fewer samples used to generate the FFT, the further apart the frequencies are so the lower the frequency resolution. For high frequency resolution we need to analyse long time periods of signals.

RTA

A common way of viewing the spectrum of a time signal is to use a Real Time Analyser or RTA. The RTA shows a plot of the amplitudes of the frequencies that make up the signals it is analysing. However, whereas the FFT produces signals that are at uniformly spaced frequencies, an RTA groups them together in fractions of an octave. An octave is a doubling of frequency, so the span from 100Hz to 200Hz is one octave. So is the span from 1kHz to 2kHz - the actual frequency span of an octave fraction is more the higher the frequency gets. For a 1/3 octave RTA the span is about 4.6Hz at 20Hz, but is 4.6kHz at 20kHz. For a 1/24 octave RTA the spans are 1/8th as wide. Within the span of an octave fraction many individual FFT values may be used to produce the single value the RTA assigns to that band of frequencies. Below is an image of the REW RTA displaying the spectrum of a 1kHz tone and its distortion harmonics.
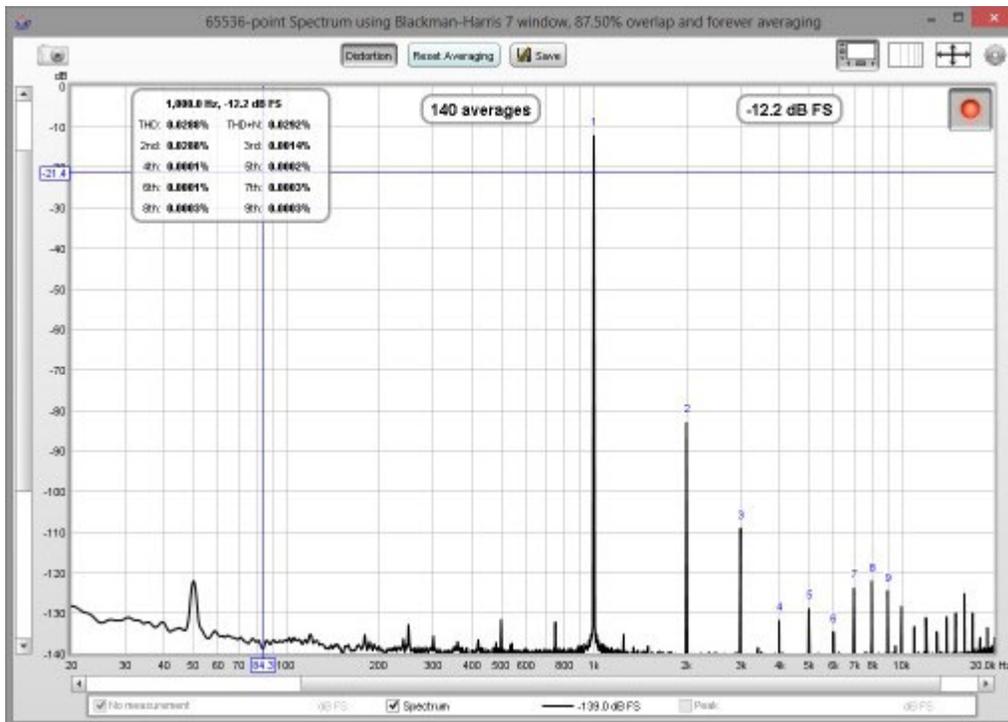
Systems and Transfer Functions

Viewing the spectrum of a signal has its uses, but we are also interested in how the equipment we use *alters* the spectrum of signals. The way a system changes the spectrum of signals that pass through it is called the system's **Transfer Function**. The transfer function has two components, the **Frequency Response** and the **Phase Response**. The frequency response shows how the amplitudes of frequencies are changed by the system, the phase response shows how the phases of frequencies are changed. A complete description of the system needs both responses, very different systems can have the same frequency response but their different phase response lets us distinguish them.

Note that it is important not to confuse a system's frequency response with the spectrum of the system's output. The spectrum of a signal shows us what that signal is made up of in terms of the frequencies it contains. The transfer function's frequency response tells us how the system *changes* the spectrum of signals. The purpose of measurement software like REW is to measure transfer functions, and REW's **SPL & Phase** graph shows the transfer function's frequency and phase responses. The frequency response amplitude is shown as an SPL trace. Below is a plot of the frequency response (upper trace, left hand axis) and phase response (lower trace, right hand axis) from a room measurement, showing the span up to 200Hz.
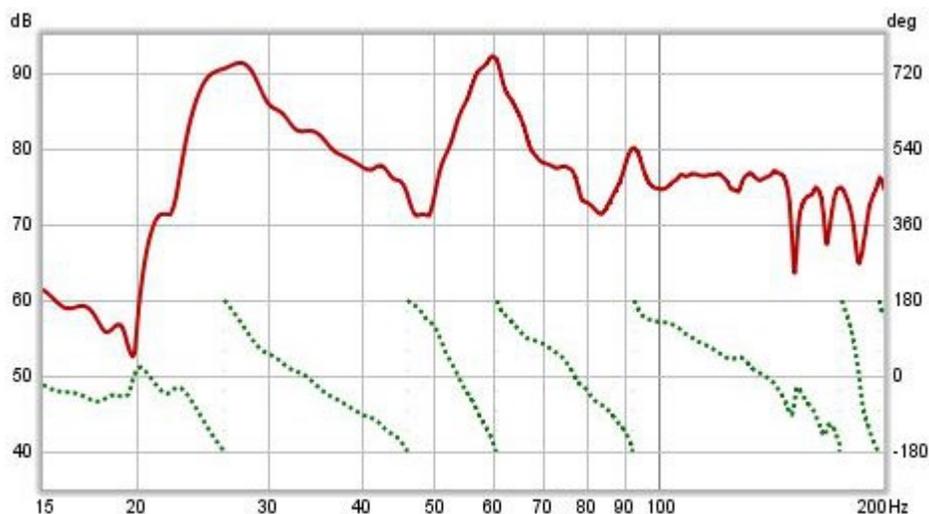
The Impulse Response

The transfer function shows us, through the frequency and phase responses, how the system affects the spectrum of signals that pass through it. It characterises the system in what is called the **frequency domain**. But what about the signal itself? How do we describe how the individual samples of the signal are changed by the system, its **time domain** behaviour? The way a system changes the samples of a signal is called its **impulse response**. The reason for the name will become clear. The impulse response (IR) is itself a signal, consisting of a series of samples. Signals that are input to the system overlap the IR as they pass through, sliding along it sample by sample. When the signal first appears, its first sample lines up with the first sample of the impulse response. The system output for that first input sample is the first IR sample value multiplied by the first signal sample value:

output[1] = input[1]*IR[1]

One sample interval later, the input has a 2 sample overlap with the IR. The output for this time period is the 2nd input sample times the first IR sample, plus the first input sample times the second IR sample:

output[2] = input[2]*IR[1] + input[1]*IR[2]

Another sample period later the input overlaps the IR by 3 samples, the output is

output[3] = input[3]*IR[1] + input[2]*IR[2] + input[1]*IR[3]

And so it goes on, as each successive input sample appears. That process of multiplying input signal samples by IR samples is called **convolution**. Typically the impulse response has a fairly short duration, much less than a second for a measurement of a piece of equipment and a second or two for a measurement of a domestic-sized room, so eventually the output at each time period consists of the length of the IR multiplied by the same length of the input signal, with all the individual products added up to give the output for that time period.

So why call it "impulse response"?

What output would we get if the input signal consisted of a single sample at full scale, to which we will assign a value of one, followed by zeroes for all other samples? The initial output sample would be

output[1] = input[1]*IR[1] = **IR[1]**

The next output sample would be

output[2] = input[2]*IR[1] + input[1]*IR[2] = 0*IR[1] + 1*IR[2] = **IR[2]**

The third sample would be

output[3] = input[3]*IR[1] + input[2]*IR[2] + input[1]*IR[3] = 0*IR[1] + 0*IR[2] + 1*IR[3] = **IR[3]**

and so on. The output would consist of each sample of the IR in turn. An input that has just a single full scale sample followed by zeroes is called an **impulse**, so the output of the system when fed that input is called the *impulse response*.

Relationship between Transfer Function and Impulse Response

As the transfer function and the impulse response are both descriptions of the same system we might reasonably expect that they are related, and they are. The transfer function is the FFT of the impulse response, and the impulse response is the inverse FFT of the transfer function. They are both views of the same system, one in the frequency domain and the other in the time domain. The transfer function is simply the spectrum of the impulse response.

Viewing the impulse response

The REW Impulse graph displays the impulse response. It shows the values as either % FS or **dBFS**. The dB scale is useful to see a wider dynamic range of the signal, rather than plot the values directly it plots the base 10 log of the values multiplied by 20. The top of the dB plot is 0 dBFS, which corresponds to 100% FS. A level of 50% FS would be 20*log(0.5) = -6 dBFS. 10% FS is 20*log(0.1) = -20 dBFS. The dBFS scale is useful to see how the lowest levels of the impulse are behaving and where it gets lost below the noise level of the measurement. The images below show an impulse response with % FS as the Y axis then the same response using dBFS. In the second image we can see the impulse takes longer to decay into the noise floor of the measurement than it might seem from the % FS plot.
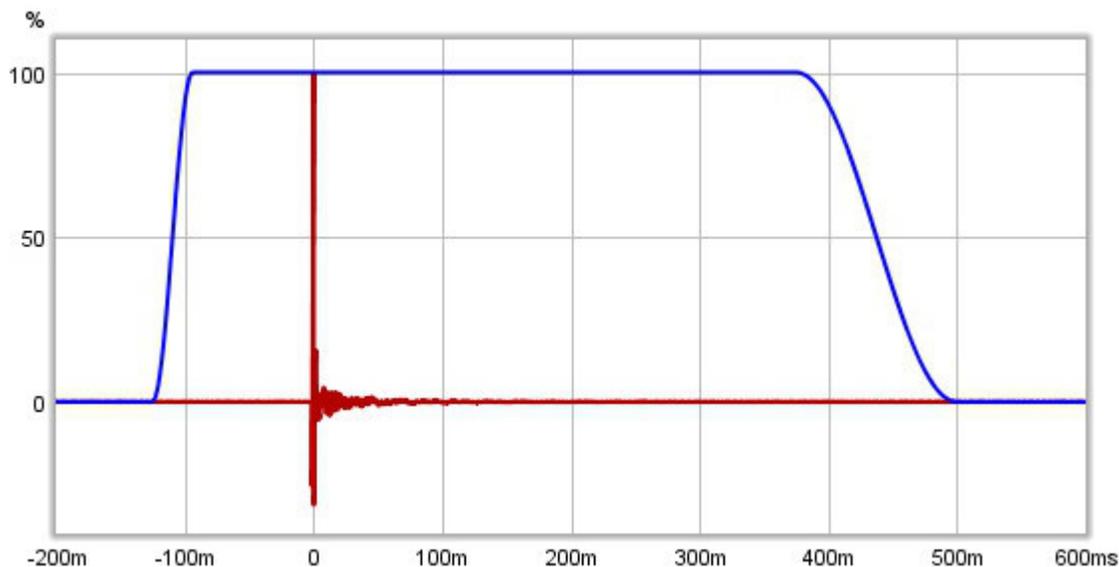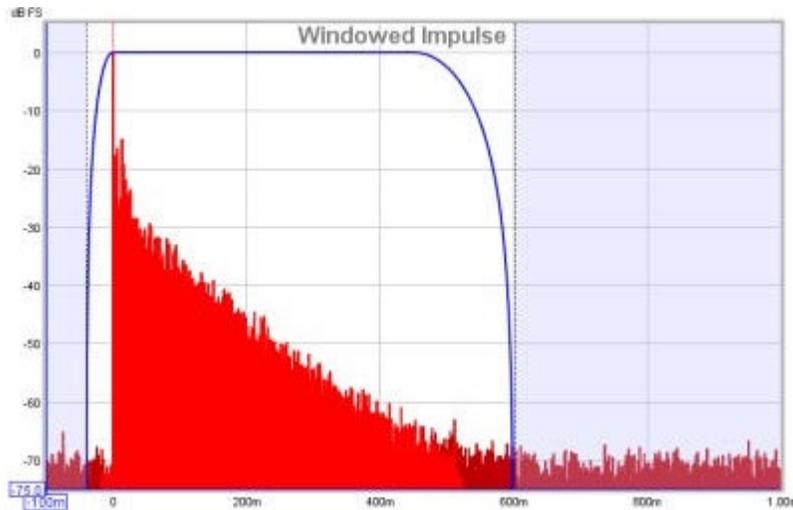
Windowing the Impulse Response

The system we want to measure might be a piece of equipment, like a loudspeaker, but in acoustics the system we are *actually* measuring includes other equipment and environments in the path between the signal generated for the measurement and the signal picked up for analysis. These include amplifiers, the microphone, the interface and most importantly the room itself. The system we are actually measuring includes all those elements, so to focus on one part of it we will need ways of removing the influence of the parts we are not interested in.

The response of the interface can be calibrated out by measuring it separately, as can the response of the microphone. Removing the effect of the room is more difficult. It may be the effect of the room is what interests us, especially if we are studying what we are hearing at our listening position, but if we are trying to isolate the performance of a loudspeaker the room's contribution can obscure details of the loudspeaker's performance.

The signal that reaches the microphone travels along a direct path, which is the shortest route from the loudspeaker and so takes the shortest time. The sound from the loudspeaker also radiates outwards and bounces off the room's surfaces. The reflections from those surfaces travel further before they reach the microphone, so they take longer to arrive. If the signal was an impulse, we would expect to see the direct arrival first, then the arrivals from the reflections. Those later arrivals are delayed by the extra time taken to travel the additional distance. The shortest that extra time can be is the time it takes sound to travel to the nearest surface - if that nearest surface was 3 feet away, for example, it would take at least 3 milliseconds longer for a reflection from that surface to reach the mic than the direct sound from the speaker (in practice it would take a little longer than that as the path distance would be a little more than 3 feet).

If we were to examine just the first few ms of the impulse response we would see the part that corresponds to the initial arrival, which came directly from the loudspeaker without a contribution from the room. Looking at a small portion of the impulse response in that way is called **windowing** the response (in the impulse response images a few paragraphs above the blue trace shows the window). If we calculate an FFT for that windowed portion of the IR we can see the transfer function for that direct arrival, which would be the transfer function of the loudspeaker alone. There is a drawback, however. If we take the FFT of a short signal, we can only see the response down to a limit that depends on how long the signal was. If we had a whole second of signal we can get a frequency response that goes down to 1Hz. If we only had 1/10th of a second, we only get a frequency response that goes down to 10Hz. In general, if the length of signal we analyse is T seconds, the lowest frequency is 1/T - so if our window was only 3ms long, the frequency response would only go down to 1/0.003 = 333Hz. To see low frequency responses free of room influences the nearest surface needs to be as far away as possible. To adjust the window settings in REW click the **IR Windows** button. By default REW uses window settings that include more than 0.5s of the impulse response, so that the effect of the room can be seen.

Waterfalls

The SPL & Phase and Impulse graphs are the most useful for studying the transfer function we have captured, but there is another graph that gives us useful information about what the room is doing to the sounds we play in it. That graph is the **Waterfall**. The waterfall is a plot of how the spectrum of a section of the impulse response changes as time progresses. It is produced by windowing an initial part of the response, typically a few hundred ms when looking at room responses, and calculating an FFT of that windowed section. The FFT produces the first slice of the waterfall. We then move the window along the impulse response a little and calculate another FFT to produce the second slice of the waterfall. Moving the window along a little further gives us the third slice, then the fourth and so on. As we move further along the waterfall we start to lose the initial contribution from the loudspeaker and increasingly see just the contribution of the room. The room's response is strongest at frequencies where there are **modal resonances**, which are frequencies at which the sound bouncing back and forth between the room's surfaces reinforces itself to produce stable, slowly decaying tones. Those frequencies stand out as ridges in the waterfall plot, with the worst modal resonances having the highest ridges that take the longest to decay.
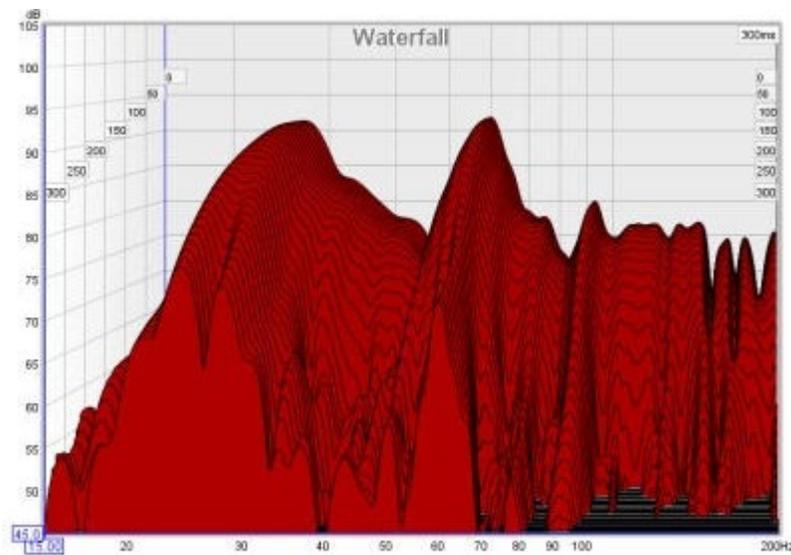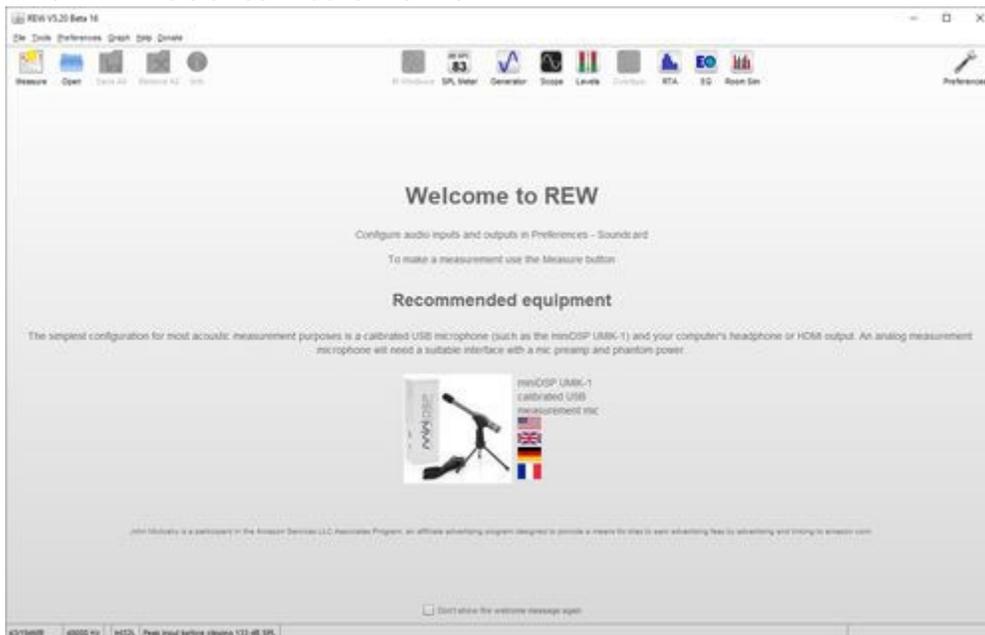

That was a very quick introduction to the basic signal and measurement concepts. If you have stuck with it all the way to the end, well done. Now you have the information needed to better understand how REW makes measurements.

Help Index

REW Overview

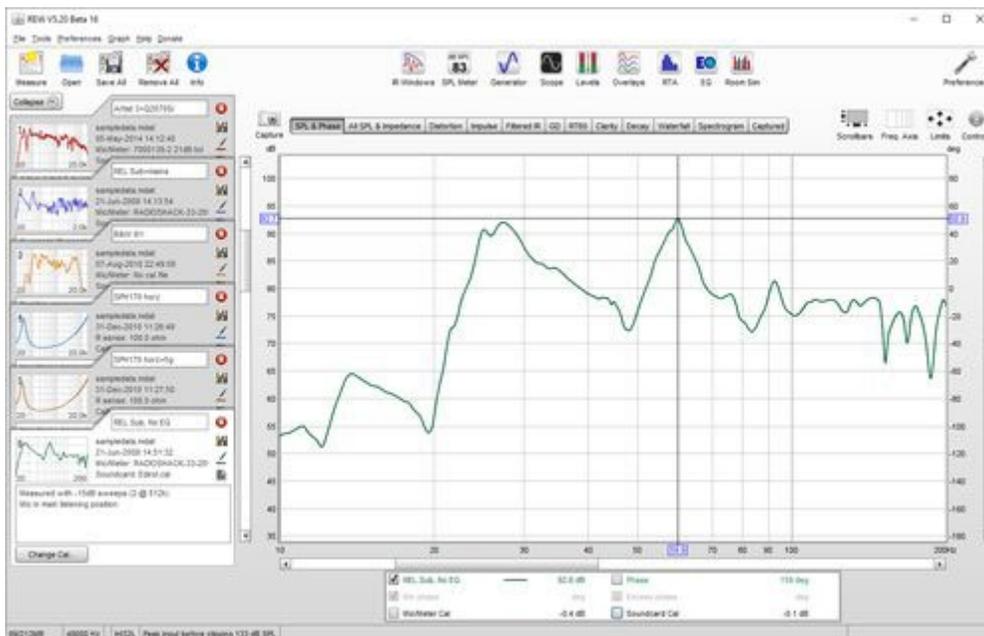When REW is started it looks like this:

The main window is blank until we either make a measurement or load some existing measurements. The SPL Meter, Signal Generator, Scope and Level Meters can be used without loading any measurements,

as can the RTA window, EQ window and the Room Simulator,

After making or loading some measurements the main window looks like this:

The measurements appear on the left, in overlapping panels. The currently selected measurement has a white background, the others are grey. On the right is the graph area for the current measurement.

All the toolbar buttons are now enabled. They are in 3 groups, firstly the buttons related to measurements:
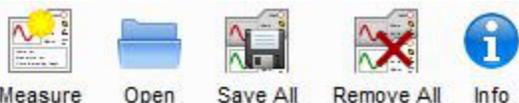
These buttons allow a new measurement to be made, existing measurement files to be opened, all the current measurements to be saved in a single measurement (.mdat) file, all the current measurements to be removed, and an Info panel to be opened that shows additional information about the current measurement.

The next group has the various tools

The IR Windows button opens a window that allows the type and extent of the Impulse Response windows for the current measurement to be changed. Next to that are the SPL Meter, Signal Generator, Scope and Level Meters buttons. Then there is a button to open the Overlays graph window, which allows any or all of the loaded measurements to be plotted on the same graph. The last three buttons are for the RTA window, the EQ window (which is used to study the effects of EQ on the current measurement) and the Room Simulator.

The final toolbar button brings up the Preferences panel
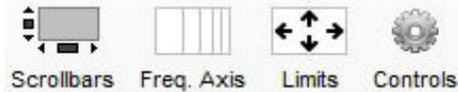

In the graph area there is a button to capture the current graph as an image

a selector strip to pick the graph type

and buttons to turn scrollbars for the graph area on/off, toggle the frequency axis between log and linear, set the graph limits and show the graph controls menu.

Below the graph is a legend area that shows the trace values at the cursor position

If smoothing has been applied the octave fraction (1/48 in the image above) appears between the trace name and its value.

The first step in getting REW running is to set up the audio input and output and calibrate the soundcard.

Help Index

Getting set up for measuring

When using REW to make measurements it is best to exit any other applications, disconnect from the internet and disable any wireless networking. Interference on audio inputs from the wireless interface or high processor demand from other applications, anti-virus updates and the like can cause gaps in the generated or captured audio signals leading to incorrect measurement results.

The initial steps required to make room measurements are:

1. Choose the audio input and output
2. Calibrate the interface (**not applicable when using a USB microphone**)
3. Check levels
4. Calibrate the SPL reading (not when using a USB microphone that has REW-compatible sensitivity cal data)

The various calibrations and level checks usually only need to be done once. If running REW for the first time it is best to read through these initial help chapters in sequence rather than jumping directly to the individual setup steps, however if your computer has already been set up using other acoustic measurement software you may be able to skip directly to Making Measurements.

Choosing the Audio Input and Output

On Windows platforms REW can use either the Java audio drivers or ASIO drivers, on other platforms only Java drivers are supported. The Windows Java drivers support sample rates up to 192 kHz but **only 16-bit data** (based on the Java 8 runtime at the time of writing, Dec 2020). Use ASIO on Windows for the full sample resolution the interface supports. On macOS and Linux 32-bit or 24-bit data is used if the interface offers it. Sample rates up to 192 kHz are supported on macOS, rates up to 1536 kHz on Linux. The sample rate and sample format/bit depth are shown at the bottom of the REW main window. Java drivers permit the input and output to be on different devices and allow volume control from REW.

ASIO drivers on Windows support up to 1536 kHz if the interface offers that and a variety of formats depending on the driver. ASIO drivers support one ASIO device which must be used for both input and output and REW has no control over levels. Pseudo-ASIO drivers such as FlexASIO and ASIO4All create an ASIO wrapper around the WDM drivers for devices, allowing input and output through different devices.

If using the Java drivers REW defaults to using the audio input and output which have been set as the defaults in your OS. Those defaults may change as you connect or disconnect audio devices, and when using the default setting REW cannot determine the specific capabilities of the device such as the bit depths or sample rates it supports, so it is strongly recommended to select the specific devices, inputs and outputs you wish to use for measurement in the Soundcard Preferences panel and do not leave the selection as default. Click the Preferences button in the toolbar to display the panel. The device lists show all audio interfaces that REW has detected, when an interface has been selected the input and output lists show the available inputs/outputs on that interface. ASIO devices appear in the list if they have an installed driver even if the interface is not connected, connect the interface to populate the list of supported sample rates.

Note 1: *if a USB interface is plugged in after REW has been started it may take up to 1 minute for it to appear in the list of devices - this is a feature of the Java Runtime Environment.*

Note 2: *if ASIO4All is being used and a change is made in the ASIO4All control panel to select other devices or inputs/outputs REW will need to reload the ASIO4All driver before it can display the newly available channels. To do that use the* **Reload** *button to the right of the* **ASIO Control Panel** *button.*

When using the Java drivers the lists may include both internal and external devices and default drivers offered by the operating system.
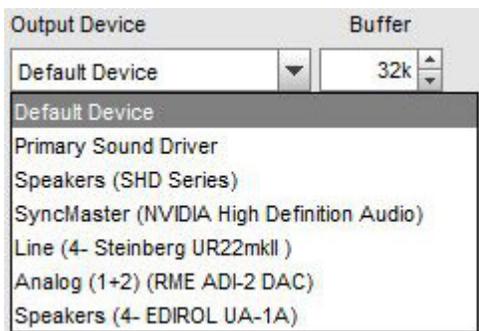
Where possible, select the interface itself rather than the OS drivers "Primary Sound Capture Driver", "Primary Sound Driver", "Java Sound Audio Engine" or similar. REW needs direct access to the controls on the interface if it is to automatically detect and adjust levels, this may not be possible if the OS drivers are selected. Java Sound Audio Engine is also prone to pops and clicks during playback which degrade measurements.

Once the devices have been chosen, the input and output can be selected. When using Java drivers the input will typically be called "LINE_IN" or "MICROPHONE" and the output will be "SPEAKER" or "LINE_OUT", however these names may be different for USB interfaces - for example, the input may be labelled "Digital Audio Interface". ASIO devices have more specific names for the available inputs and outputs and each mono channel will be listed separately.

When using a USB mic with a cal file that contains a sensitivity figure REW needs to read the input volume setting to correctly show SPL, to allow that the input device and input for the mic must be selected (they must not be left as "Default Device").

There seems to be a growing trend for Windows laptops and PCs to apply noise cancellation processing to stereo microphone inputs, suppressing the content that is common to both channels. With USB microphones like the UMIK-1 the signal is the same on both channels, so it is all common content, and may be heavily suppressed. If there are user controls for noise cancelling processing on inputs make sure they are off for the measurement input. Processing for an input can be turned off by going to the Windows **Sound Settings**, then selecting **Sound Control Panel** under **Related Settings**. Select the microphone to be used from the **Recording** tab then click the **Properties** button and select the **Advanced** tab. If there is an **Enable audio enhancements** box make sure it is **not** checked.

Trouble-shooting tip: To prevent REW from accessing interface volume controls, uncheck the **Control input volume** *and* **Control output volume** *boxes. Any level control settings will then need to be made using the OS volume controls and/or the interface's mixer.*

Multichannel Output

When using Java drivers a channel may be selected after choosing an input or output. The Java drivers typically provide stereo inputs and outputs, so the measurement input can be Left or Right and the output can be Left, Right or, to send the measurement signal to both channels, Left+Right. On Windows the Java drivers currently (Java 8 update 202) only offer stereo outputs even if the interface connected is configured for or supports multi-channel, for example an HDMI output offering multi-channel PCM audio. Multichannel access on windows requires ASIO drivers. To access multichannel outputs over HDMI you will need to use an ASIO wrapper like FlexASIO or ASIO4All.

On macOS and Linux a device configured for multi-channel can make all its channels available. An example of selecting from any of the 8 outputs of a multichannel card on macOS is shown below, the card is configured for 8-channel, 16-bit data at 48 kHz as shown in the Audio Midi Setup image. Note that macOS has an unusual channel order for this card, with the back channels appearing in the list before the side channels. REW labels the channels in order as Left, Right, C, LFE, SL, SR, SBL, SBR but on this card selecting SL (for example) in REW produces output on the left channel of the connector marked Rear on the

interface and SBL produces output on the left channel of the connector marked Side. That oddity aside the channels are individually selectable, and any unusual channel assignments can be dealt with using the Output Channel Mapping control.
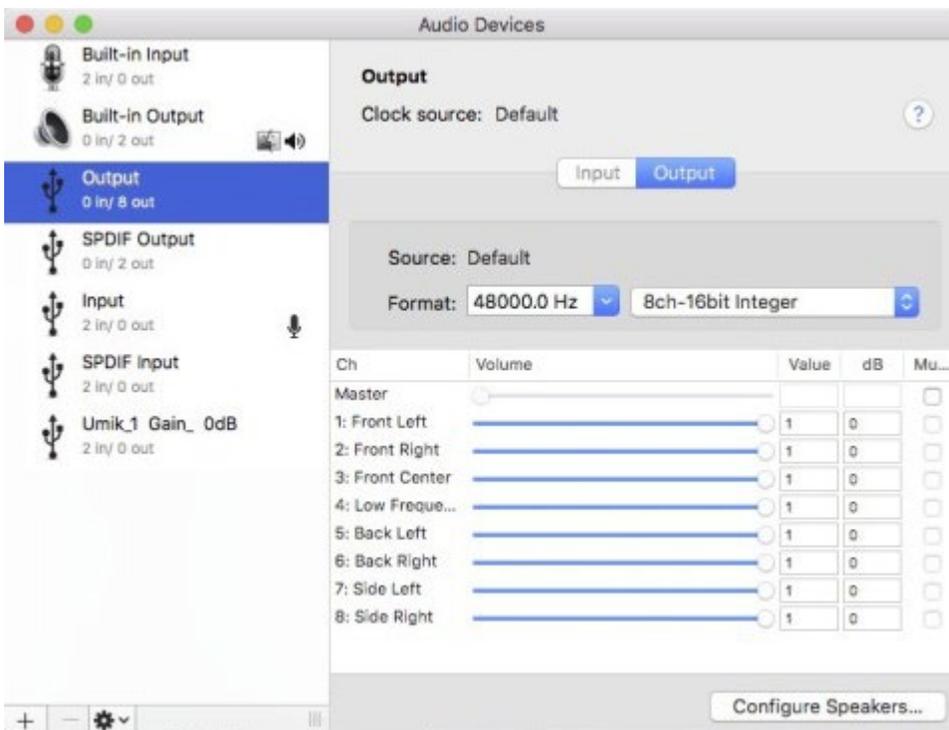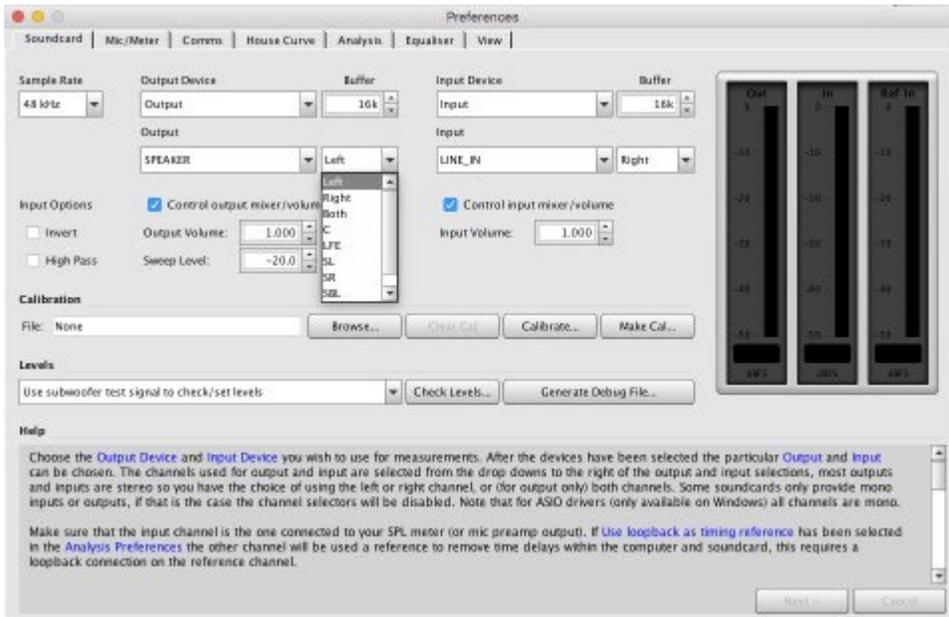
Sample Rate Selection

When using ASIO drivers sample rates up to 1536 kHz will be offered, if the interface supports them. Large buffer size (high latency) is recommended for ASIO, especially at higher sample rates. Reliable operation at high sample rates depends on the computer speed and RAM.

The Java drivers offer sample rates up to 192 kHz on MacOS and Windows or 1536 kHz on Linux, though devices may only be available for 44.1 kHz and 48 kHz depending on the OS and the devices. **Make sure**

the input and output devices are configured in the OS to operate at the rate selected in REW, otherwise the OS will resample between the selected rate and the rate at which the input or output device is actually running.

It is best to use 44.1 kHz or 48 kHz for acoustic measurements unless the test specifically requires measurement results above 20 kHz (studying tweeter resonances, for example). Higher sample rates increase memory use, slow down processing and do not improve accuracy.

Calibrating the Interface

This step is not applicable when using a USB microphone as the input, skip directly to Check levels

Once the audio input and output have been selected REW is ready to make a calibration measurement of the interface's frequency response. This is important to check that the interface is configured correctly, particularly to ensure that monitoring is not active - if any of the input signal is routed back to the output measurement results will not be valid. The calibration result can be used to remove the interface's response from measurements.

1. Connect the interface's line or headphone output directly to its line input - use the channel that will be used to make measurements, which should be the same one that has been selected in the **Input Channel** control.
2. Press the **Calibrate soundcard...** button on the Soundcard Preferences panel and follow the instructions in the help panel at the bottom
3. The measurement result should be fairly flat (varying by much less than 1dB over most of the range) but will roll off at the lowest and highest frequencies and will often have some ripple at the high end. If the result shows excessive variation between 20Hz and 20kHz a warning message will be shown, such a measurement should not be used. Below is a valid measurement from a laptop's internal card which rolls off fairly rapidly below 20Hz (graph vertical scale is 1dB/division):
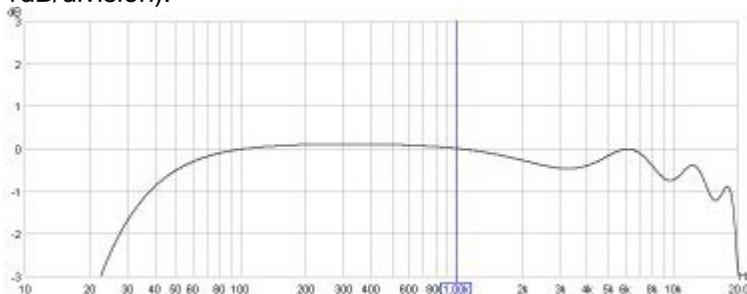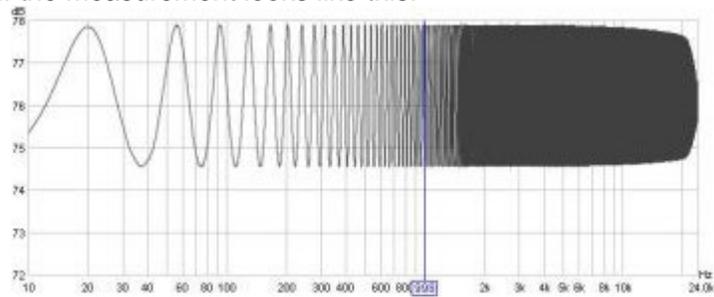

   This is a measurement from another laptop interface, showing better low frequency response and a smoother high frequency response:
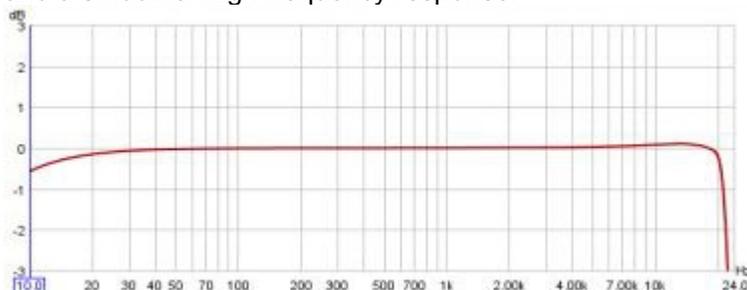
4. Some results from the measurement are shown in the notes:
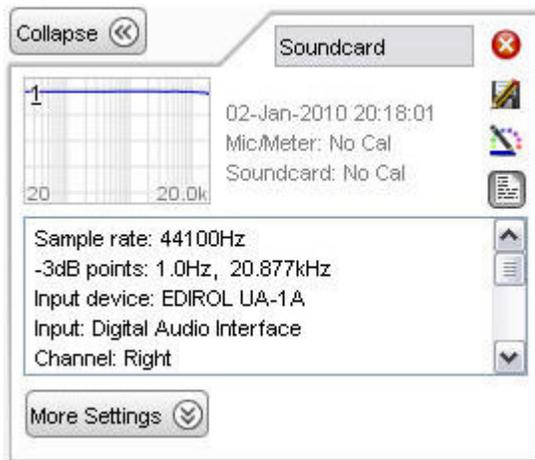
5. If the measurement looks like this:

    it is probably due to a feedback loop from the Line In to the output. This can happen if the interface has some feature for record monitoring - for example, there may be a "Monitor" feature for the Line In that must be turned off to get correct results, or there may be a setting for "Record without monitoring" that must be selected. Feedback can also happen if "Line In" is not muted in the Playback mixer for the interface, or if "Listen to this device" has been selected in the Windows Recording properties for the input.

6. Having obtained a measurement it needs to be saved as a calibration file. Press the **Make cal file...** button in the Soundcard Preferences and choose a name and location for the file. The file is saved and then automatically re-loaded as a calibration file to use for all subsequent measurements with that interface. On the next startup the file will be loaded automatically.

7. The calibration file will be applied to all new measurements made after it has been loaded. To apply or remove a soundcard calibration file for an existing measurement, use the **Change Cal... button** in the measurement panel.

8. Remove the loopback connection and connect the interface to the mic or SPL meter and AV processor/equaliser

Note that interface measurements made from the Soundcard Preferences panel use the full sweep range to half the interface sample rate, regardless of the sweep end frequency setting, and the soundcard calibration file is NOT applied to such interface measurements

Note also that the soundcard calibration file is only valid for the sample rate at which it was measured, if the sample rate is changed the interface should be re-measured at the new sample rate

The next step is to [Check levels](#)

[Help Index](#)

Check Levels

REW measurements are usually made at a level of about 75 dB SPL. This is not very loud, it is the same level aimed at by most receivers to adjust speaker trims. Using very loud test signals is likely to damage your speakers and your ears. Do not use test signal levels any higher than you would be comfortable listening to for long periods. Setting the signal level REW uses during measurement involves generating a pink noise calibration signal and adjusting the AV processor's volume control and/or the calibration signal level so that at the measurement point (usually ear height at your main listening position) your SPL meter shows a level of around 75 dB. Then, **if not using a USB microphone**, the interface's input volume needs to be adjusted to get a good signal level from the SPL meter or mic preamp when the cal signal is playing.

With a USB microphone the input volume control can be left at the unity gain (0 dB) setting, which is selected by default when the microphone is first plugged in (and set automatically by REW if the **Control input volume** box is ticked on the Soundcard preferences).

Check Levels Procedure

Open the Soundcard Preferences panel and choose whether to set the levels using your subwoofer or a main speaker, making the appropriate selection in the drop-down box in the **Levels** panel. This tells REW whether to use a subwoofer or speaker calibration signal. If you have connected the interface output directly to your subwoofer or to an equaliser that is connected to your subwoofer choose **Use Subwoofer to Check/Set Levels** here, if you are connected to an AV processor input you can use the subwoofer or main speaker settings.

Press the **Check Levels...** button and follow the instructions on screen. The test signal defaults to an RMS Level of -12 dBFS. If connected to an AV processor, start with the volume fairly low and increase it until the meter is reading around 75 dB. The exact level is not critical. If connecting directly to an equaliser such as the BFD, use the Sweep Level control to change the level of the generated signal. In either case, the final Sweep Level will be used for subsequent measurements - remember to use the same AV processor volume setting whenever measurements are made.

If input levels are low DO NOT KEEP MAKING THE TEST SIGNAL LOUDER. Input levels should be set through the volume controls on the input path, not the output, using very loud test signals is likely to damage your speakers and your ears.

Notes on Volume Controls

The output volume control may have no effect on the level of line outputs, but will affect a headphone output or the PC's speakers - if using a true line output for the measurement signal the output volume control should be muted to prevent the signal being heard from the PC's own speaker(s).

*The following notes are relevant if an output device and output have been selected in REW and the **Control output volume** box is enabled and ticked.*

- REW sets the interface's output volume to half scale.
- The range of the REW controls is 0..1, but the underlying gain controls in the interface are generally logarithmic. If you use REW to alter the volume control settings using the arrows on the spinners, the settings will change in increments corresponding to around 0.5dB, however the interface's own controls will usually have lower resolution, so changes may have no effect on the output until the next step in the interface's control is reached.
- Changes in volume settings outside REW (e.g. in the Windows mixer) are automatically detected and reflected in REW's controls.
- In the playback mixer (if there is one) it is important to mute the input used for the signal from the

mic or SPL meter (usually Line In), otherwise a feedback loop results as the signal picked up by the mic/meter is fed directly back to the output. Here is an example of correct playback mixer settings on a Windows XP system, only Wave is unmuted, the Wave volume is full scale and the output volume is half scale.
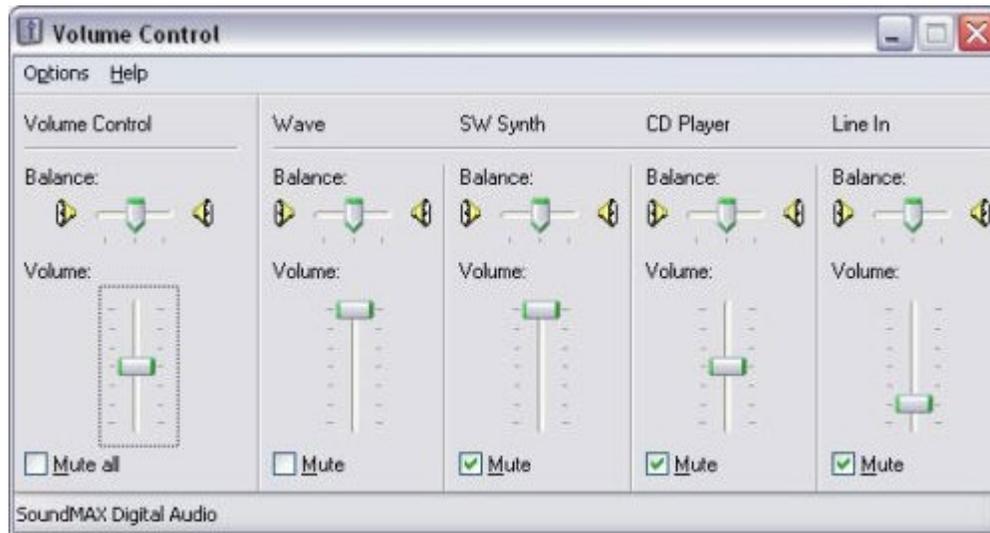
Trouble-Shooting Setting Input Volume

- Some interfaces do not provide an input volume control (e.g. some USB cards). If using an SPL meter or mic and preamp, as long as the signal RMS level is within the -30 to -12 dB range when the speaker/subwoofer cal signal is playing everything should work OK. If the level is below -30dB try reducing the meter range to raise the level, but be careful not to reduce the range to the point the meter is overloaded. If the level is above -12dB try increasing the meter range to lower the level.

- If using a USB mic the signal may be much lower, below -50 dB, this is normal.

- Some interfaces do not provide programmatic access to their input volume control so REW may not find a volume control or it may not be able to alter it. In that case you can use the controls in the interface's mixer or the OS audio level controls to make the required adjustments manually.

- On some interfaces REW may not be able to directly select the required input via its device and input selectors - for example, on Audigy 2 selecting the Line In is done by selecting "Analog Mix" in the Record panel of the Basic tab of the Creative Surround Mixer then going to the Source panel and muting all the sources except for Line In. If it does not seem possible to select the required input via REW's device and input selectors, or REW does not seem to be making the correct settings, leave the input device set to **Default Device** and make the input selection and input volume adjustments via the interface's mixer or the OS audio level controls. An example of suitable Windows XP volume control settings is shown below, with Line In selected
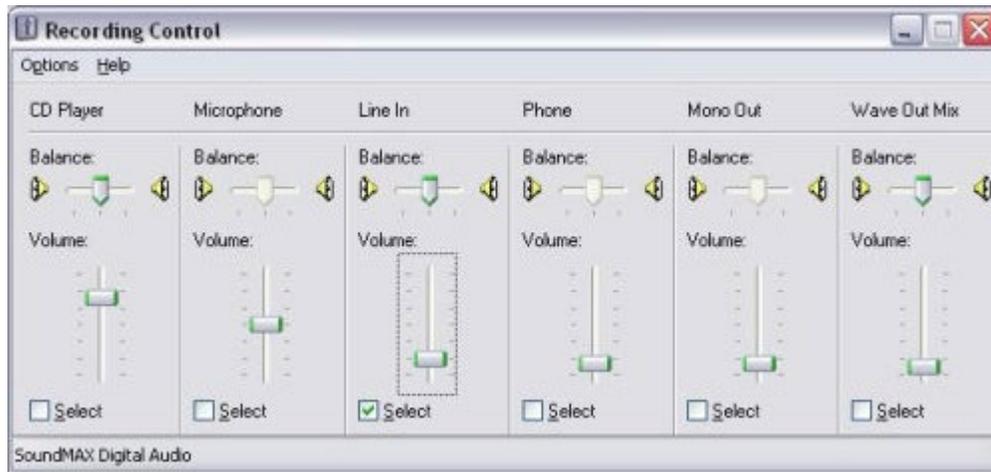
After checking levels the next step is to calibrate the SPL reading.
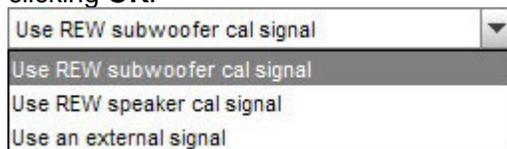
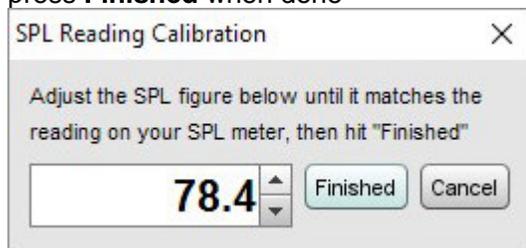Help Index

Calibrating the SPL Reading

Calibrating the SPL reading gives REW an absolute SPL reference by entering a reading from your SPL meter while a speaker or subwoofer calibration signal is playing. Alternatively a microphone calibrator may be used. This step is not required if you are using a USB mic with a cal file that contains a sensitivity figure.

SPL Calibration Procedure

1. Select the Cal files tab of the Preferences panel
2. If you are using an SPL meter as the microphone when making your measurements:
    o   Set the meter to C weighting and select the **Input device is a C Weighted SPL Meter** option in the mic cal file entry for the input being used
    o   Set the meter range to suit the measurement level used in the check levels process (the 80dB range is recommended for the Radio Shack meter)

3. If you are using a measurement microphone leave **Input device is a C Weighted SPL Meter** unchecked
4. Open the REW SPL meter by clicking the **SPL Meter** button in the toolbar then press the **Calibrate** button
5. On the dialog which appears choose whether to calibrate the meter reading using your subwoofer or a main speaker driven by a calibration signal generated within REW, or to use an external test signal you provide, making the appropriate selection in the drop-down box and clicking **OK**.

Use REW subwoofer cal signal	▼
Use REW subwoofer cal signal	
Use REW speaker cal signal	
Use an external signal	

6. Enter the reading from your external SPL meter (not the REW meter) in the calibration panel and press **Finished** when done

SPL Reading Calibration   ✕

Adjust the SPL figure below until it matches the reading on your SPL meter, then hit "Finished"

78.4   [Finished] [Cancel]

For more information about the REW SPL Meter see the SPL Meter help

SPL Calibration Notes

•   If you are using a USB mic with a cal file that contains a sensitivity figure SPL calibration is not required, REW will show this message:.

- If you switch between a number of analog microphones or interfaces there is a method to include sensitivity information in the mic cal file. Add a line at the start of the file of the form:

  ```
  Sensitivity -12.34 dBFS
  ```

  where "-12.34" is replaced by the REW dBFS input level when the microphone is driven by 94 dB SPL. When that cal file is used REW will use the sensitivity data to calculate the SPL calibration offset. Note, however, that this calibration will **only be valid when using the same mic interface gain and input volume that were used when it was measured**.
- At the end of the calibration process a message is displayed showing the maximum SPL that can be measured with the current input level settings. If higher SPL needs to be measured the input sensitivity must be reduced (by lowering the soundcard input volume setting or, if using an external microphone preamp, reducing the preamp gain or, if using an external SPL meter, increasing the range setting) and repeating the calibration
- A warning is shown if measurements are attempted before calibrating the SPL reading as the results of the measurement would not reflect the actual SPL. The SPL reading is shown in red until the SPL calibration procedure has been completed
- REW remembers the calibration setting for the next startup, so it will usually not be necessary to repeat this process, but check the REW SPL reading against your SPL meter or calibrator before starting a new set of measurements
- If the input volume setting is changed it will be necessary to repeat the calibration process
- If another input device is selected it will be necessary to repeat the calibration process for that device if it has not been done previously. If the device has previously been calibrated REW will automatically adjust calibration settings to suit
- if the signal generator is already playing a signal when you click **Calibrate** it will continue to play the same signal

REW is now ready to make measurements.

Help Index

Making Measurements

Once the audio input and output <u>have been chosen</u>, the interface <u>has been calibrated</u>, the levels <u>have been checked</u> and the SPL reading <u>has been calibrated</u> REW is ready to make response measurements.

Connections should be as explained in <u>Getting Started</u>, if connected to an AV processor, select the input to which the interface's output is connected.

Making a Measurement

- Press the Measure button (or Ctrl+M) to bring up the Measurement dialog

- Make sure the SPL/Impedance selector at the top of the Measurement panel is set to SPL (see <u>Impedance Measurement</u> for information about measuring impedance)
- Enter a name for your measurement. REW can automatically append a number or the date and/or time to the name according to the selection to the right of the name. To use just a number or just a date or time leave the name blank and select the desired option. The name that will be used is shown below the name you enter, the buttons to the right of the naming options provide further settings. Notes for the measurement can be entered in the box. The name and notes can be changed afterwards in the <u>Measurements Panel</u>

- Set the Start Freq to the lowest frequency for which you wish to see the response and End Freq

to the highest. The sweep will span the range from half the start frequency to twice the end frequency (with an overall limit of half the interface sample rate) to provide accurate measurement over the selected range

- Level controls the rms signal level at which the sweep is generated. The maximum value is -3 dBFS or its equivalent in dBu, dBV or volts, unless the View preference Full scale sine rms is 0 dBFS has been selected, in which case the maximum is 0 dBFS. Using the maximum value places the peaks of the signal at digital full scale. The default value is -12 dBFS. This control is normally preset to the Sweep Level established during the Check Levels process. If you will be comparing measurements from several speakers, or comparing a series of measurements from a speaker, make sure they are measured with the same Sweep Level.

- Measurement progress and input levels are shown below the Level control.
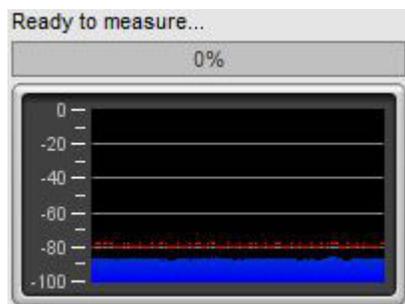
- Settings for the measurement method are on the right hand side. Length controls the length of the sweep, specifying the number of samples in the sequence. The default is 256k which works well for most purposes. Dividing the number of samples by the interface's sample rate gives the sweep duration in seconds, shown to the right. The overall duration includes silent periods before and after the sweep.

  If the input and output are on the same device and so share a common clock longer sweeps will provide higher signal-to-noise ratio (S/N) in the measurements. Long sweeps may be problematic when the input and output are on different devices, such as when using a USB mic, as their sample clock rates will differ. Over a long sweep a significant difference in clock rates could cause severe distortions in the shape of the impulse response and affect the phase response. This can be corrected by using the acoustic timing reference and the Analysis option to Adjust clock with acoustic ref or, if a loopback connection is being used as a timing reference, the Analysis option to Adjust clock with loopback.

  Each doubling of the sweep length improves S/N by almost 3 dB. However, the time required to perform the processing after each sweep will more than double. If REW is running on a computer which does not have a fast processor and a lot of memory, measurements will be much faster using the shortest sweep length (128k samples), at a small S/N penalty of about 3dB compared to the default. At least 4 GB of RAM and a fast processor are recommended if using the 1M sweep, at least 8 GB for the 4M sweep. Invalid measurements or out-of-memory errors may occur on computers which have insufficient RAM or processor speed.

- REW allows multiple sweeps to be averaged, although best results are generally obtained by using single, longer sweeps rather than multiple, shorter sweeps. Multiple sweeps are offered when not using a Timing reference or when using a loopback as a timing reference. **Do not use multiple sweeps if the input and output are on different devices** (for example, if the input is a USB mic). If Repetitions is more than 1 REW uses synchronous pre-averaging, capturing the selected number of sweeps per measurement and averaging the results to reduce the effects of noise and interference. The pre-averaging can improve S/N by almost 3 dB for each doubling of the number of sweeps. Averaging can be useful if the measurements are contaminated by interference tones, whether electrical or acoustic, as they typically will not add coherently in the averaging and hence will be suppressed by the process.

  Warning: some interfaces do not maintain sample synchronisation between the successive sweeps which produces a corrupt measurement that has multiple, closely-spaced peaks of approximately the same level in its impulse response, 1 peak for each sweep. This can also

happen if the input and output are on separate devices. **If the frequency response with multiple sweeps is significantly different from the response with a single sweep, stick with single sweeps**

- A Timing reference can be selected for the measurement, see Measuring with a timing reference for more on the settings when using a timing reference. When not using a timing reference there are 3 choices for the t=0 position of the measured impulse response: (1) **Set t=0 at IR peak** will place t=0 at the exact position of the peak of the IR, determined by interpolating an oversampled version of the IR; (2) **Set t=0 at IR start** will place the t=0 point at the sample before the IR first exceeds 10% of the peak level; (3) **Set t=0 from delay estimate** will use cross-correlation with a minimum phase version of the impulse to estimate and remove any time delay.

- There are two Protection mechanisms for sweep measurements. **Abort if heavy input clipping occurs** will abort the measurement if more than 30% of the samples in an input block are clipped. **Abort above SPL limit** will abort the measurement if the SPL exceeds the limit that has been set. Note that if the limit set is higher than the input can measure before clipping occurs it will not offer any protection, but in that case the first mechanism may be triggered.

- Playback selects whether the sweep is played by REW or will be played back from a Measurement Sweep file that has been generated and saved previously from the signal generator. Usually REW plays the sweeps, see Measuring with file playback for more on using a file.

- The Sample rate, Output and Input are selected below the Playback control. If a timing reference is being used the Ref output and Ref input may also be configured.

- Normally a single measurement is made, but a series of measurements can be made without further intervention by setting the total number required in the Measurements count. That may be useful when making a series of polar measurements, using the Delay control to specify how long REW will wait before starting each sweep to allow time to reposition a speaker turntable, for example. The Delay control can also be used for single measurements to provide time to move away before the measurement starts. When making multiple measurements a panel is shown with the number of measurements remaining, the name that will be assigned to the next measurement and a pause button that will stop the sequence when the current measurement completes and continue it when unpaused.
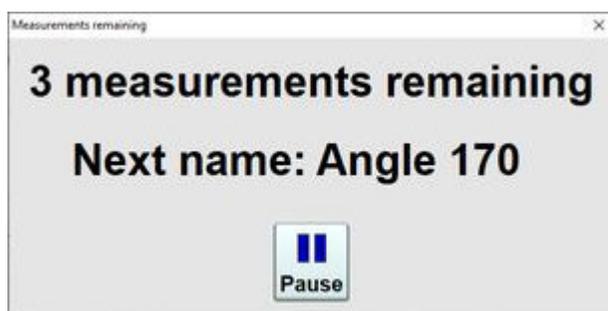
- If a multi-input measurement is being made there is a **Configure average...** button to select whether the individual inputs are weighted and/or have their SPL aligned before the average is calculated.

- The **Check Levels** button generates a few seconds of a pink noise signal that depends on the frequency range selected for measurement and checks that the input level is not too high or too low. Pressing **Cancel** while the pink noise signal is playing will turn it off (it turns off automatically after 3 seconds). The rms level that was measured is shown on the measurement panel, with a warning if the level is too high or too low

- Press **Start** to make a measurement. If a delay has been configured the time remaining before the sweep begins is shown

- When the sweep starts progress is shown on the measurement panel along with a display of the measurement headroom and, if a timing reference is being used, the headroom for the reference
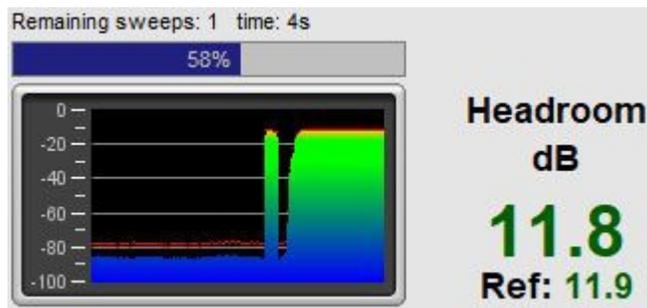
The result of the measurement is displayed in the graph area, information about the measurement appears in the Measurements Panel. The measurement name can be changed if desired by entering a new name in the box at the top of the measurement panel
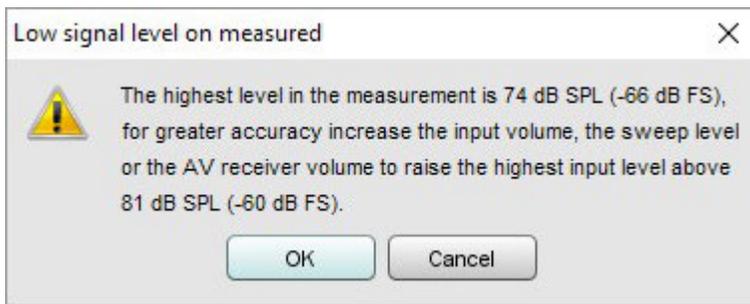

Notes relating to each measurement can be entered in the notes area, click the Notes button if the notes area is not visible
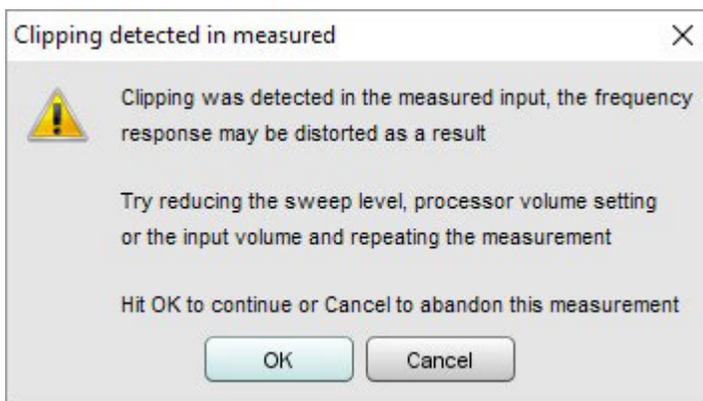

For details of the various ways of viewing the measured data, including averaging multiple measurements, refer to the Graph Panel help.

Measurement Headroom

The headroom figure on the measurement panel shows how far away the input is from clipping, and hence how much the sweep level could be increased before clipping would occur. The figure is red if there is less than 6dB of headroom (warning that the input is close to clipping), green between 6 and 40 dB (or between 6 and 60 dB for USB mics). A message is shown if the headroom is more than desirable, as increasing the Sweep Level or the AV processor volume would improve the signal-to-noise ratio in the measurement which in turn increases the accuracy of the impulse and frequency responses. Note that after making such a change subsequent measurements will be at a higher SPL on the graphs than those made before the change.

Low signal level on measured   ✕

⚠ The highest level in the measurement is 74 dB SPL (-66 dB FS),
    for greater accuracy increase the input volume, the sweep level
    or the AV receiver volume to raise the highest input level above
    81 dB SPL (-60 dB FS).

            [OK]   [Cancel]

If the room's resonances are very large the input signal level may exceed the input range and cause clipping. If this occurs a warning is displayed, as input clipping will cause errors in the derived frequency response. The sweep level, AV processor volume or input volume should be reduced and the measurement repeated. Note that after making the change subsequent measurements will be at a lower SPL on the graphs than those made before the change.

Clipping detected in measured   ✕

⚠ Clipping was detected in the measured input, the frequency
    response may be distorted as a result

    Try reducing the sweep level, processor volume setting
    or the input volume and repeating the measurement

    Hit OK to continue or Cancel to abandon this measurement

            [OK]   [Cancel]

If the signal levels are very low this may indicate a connection problem:

Very low signal level on measured   ✕

⚠ The highest level in the measured input is just 54 dB SPL (-87 dB FS)

    This is much lower than it should be. Check the input channel
    selection, connections, input volume and output volume/mute.

    The highest level should ideally be above 81 dB SPL (-60 dB FS)

            [OK]   [Cancel]
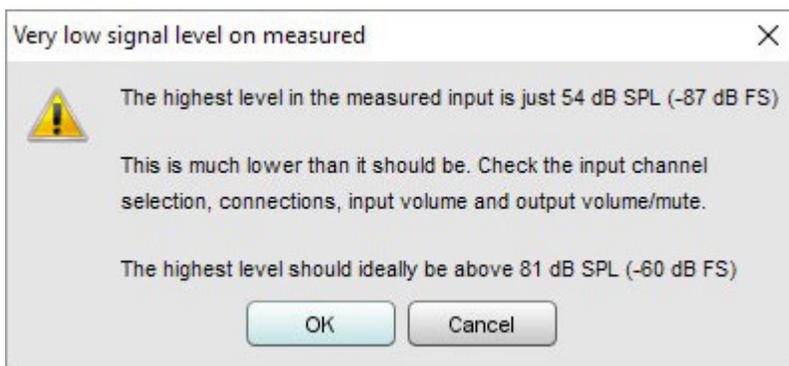
After measuring the response of a channel you can look at adjusting EQ immediately, or make other measurements first.

Note that some resonances which are very pronounced when measuring a speaker alone do not appear if a pair of speakers (e.g. Left and Right) are run together - this is because the positioning of the speakers in the room can prevent some resonances being excited (in particular, the odd order width modes will not be excited by content which is the same on Left and Right speakers if they are symmetrically placed across the width). Such resonances can often be left uncorrected, to identify them compare measurements from individual channels with those made with two channels driven at the same time (achieved on AV32R DP or AV192R by setting the `Repeat Sig.` entry in the TMREQ filter menu to Yes and selecting the channel which is to repeat the test signal, or on other processors by connecting both left and right interface outputs to the selected AV processor input or using a Y lead to drive two inputs at once).

Measuring with a timing reference

REW can make use of a timing reference when it measures, according to the setting on the measurement panel. The timing reference selection controls whether REW uses a loopback on the interface as a timing reference, or an acoustic timing reference, or no reference. Using a timing reference allows REW to eliminate the variable propagation delays within the computer and interface so that separate measurements have the same absolute timing. A **Timing offset** can be set that will shift the response relative to the timing reference, allowing the time of flight delay to be removed if using a loopback connection, for example. The timing offset can be set manually or from the figure calculated using the Estimate IR Delay function.

If a loopback is selected the reference channel signal must be looped back from output to input on the interface and measurements will be relative to the loopback timing. Usually this means measurements will have a time delay that corresponds to the time it takes sound to travel from the speaker being measured to the microphone, which can be removed for subsequent measurements using the **Timing offset**.

If an acoustic timing reference is used REW will generate a timing signal on the output that has been selected to act as the reference before it generates measurement sweeps on the channels being measured. The level of the timing reference is set relative to the measurement level using the Ref level trim control, it may need to be higher or lower depending on whether the speaker used as the timing reference is further away or closer than the speaker being measured. The timing signal is a high frequency sweep to allow accurate timing, a subwoofer cannot be used as the reference channel. Measurements will have a time delay that corresponds to the difference in their distance from the microphone compared to the distance of the reference speaker - if the reference speaker is further away the delay would be negative. The delay can be shifted using the **Timing offset**. When an acoustic timing reference is used individual measurements **taken from the same mic position** will have the same relative timing, allowing trace arithmetic to be carried out on the traces in the All SPL graph. **Note** that multiple sweeps cannot be used when using an acoustic timing reference.

If using a timing reference REW can calculate the delay through the system being measured relative to the reference and show it in the measurement Info panel as **System Delay** in milliseconds, with the equivalent distance in feet and metres shown in brackets. Any **Timing offset** is shown below the **System Delay**. For speakers the delay estimate is based on the location of the peak of the impulse response. Subwoofers have a broad peak and a delayed response due to their limited bandwidth so the delay is instead measured relative to the start of the impulse response. The start of the impulse response cannot be located as precisely as the peak, however, so delay values are less accurate for subwoofer measurements.

Measuring with file playback

It sometimes is not possible to connect to an input on the system being measured for it to play back the sweep. If the system can play back from files a sweep file can be used instead.
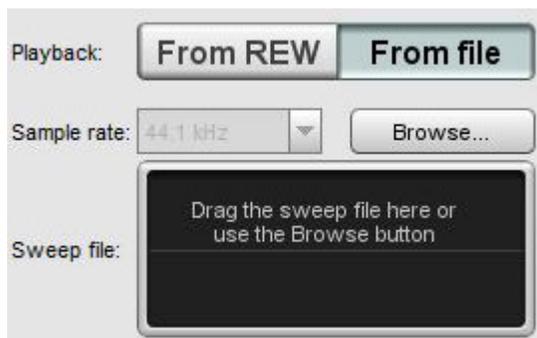

A Measurement Sweep file must be generated and saved from the REW signal generator. The sweep file must include the acoustic timing reference signal, REW uses that to synchronise input capture to playback

on the system being measured. A copy of the file is placed on the system being measured and the file is also loaded into REW so that it can configure the correct measurement frequency range, sweep settings and sample rate, which are retrieved from the file.
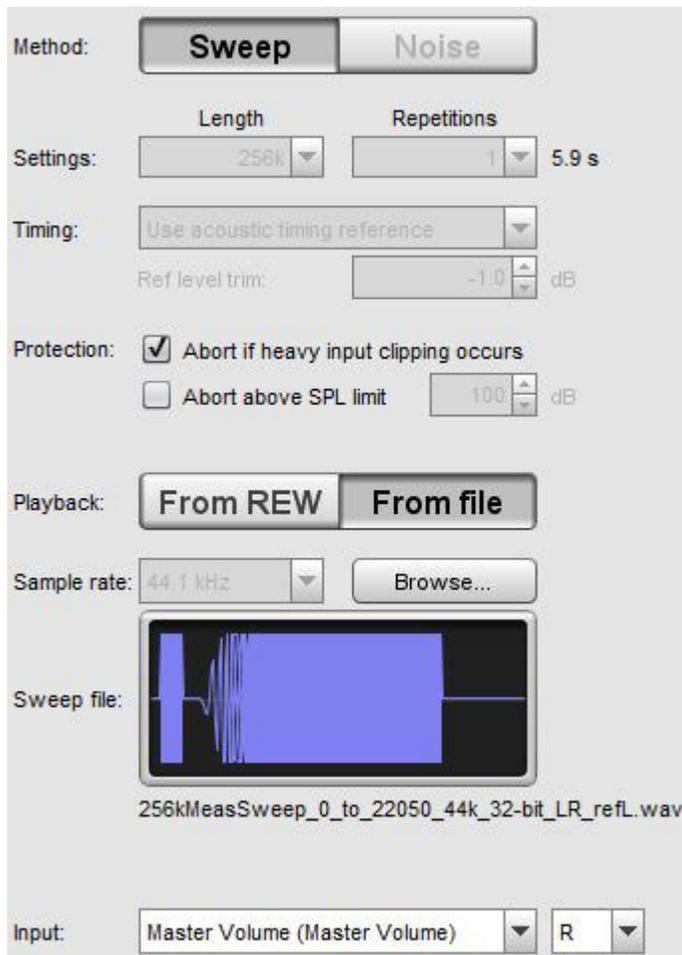
After loading the file press **Start**, REW will then capture the noise floor and pause, showing **Waiting for timing reference...**
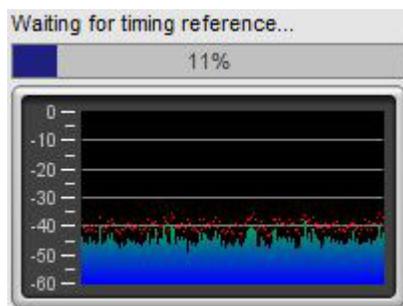
Now play back the file on the system being measured, once REW detects the timing reference it will capture the sweep and generate the measurement.

Number of measurements

REW can have up to 199 measurements loaded at once according the setting for Maximum measurements in the View preferences, the default is 30. If the maximum number has been reached when a new measurement is requested a warning is given as the first measurement would need to be removed to make room for the new one:

Help Index

Offline Measurements

Sometimes it isn't possible or practical to connect a computer to the system you wish to measure and the live File playback option is not suitable. In those cases REW's measurement sweep signal can be played back through the system and the response recorded by other means, the sweep signal and response recording can then be imported to generate a new measurement. This is done using the Import Sweep Recordings entry in the File menu.

Generating the sweep signal

The first step is to produce the sweep measurement signal that will be played back when making the response recording. It is generated from the Signal Generator.
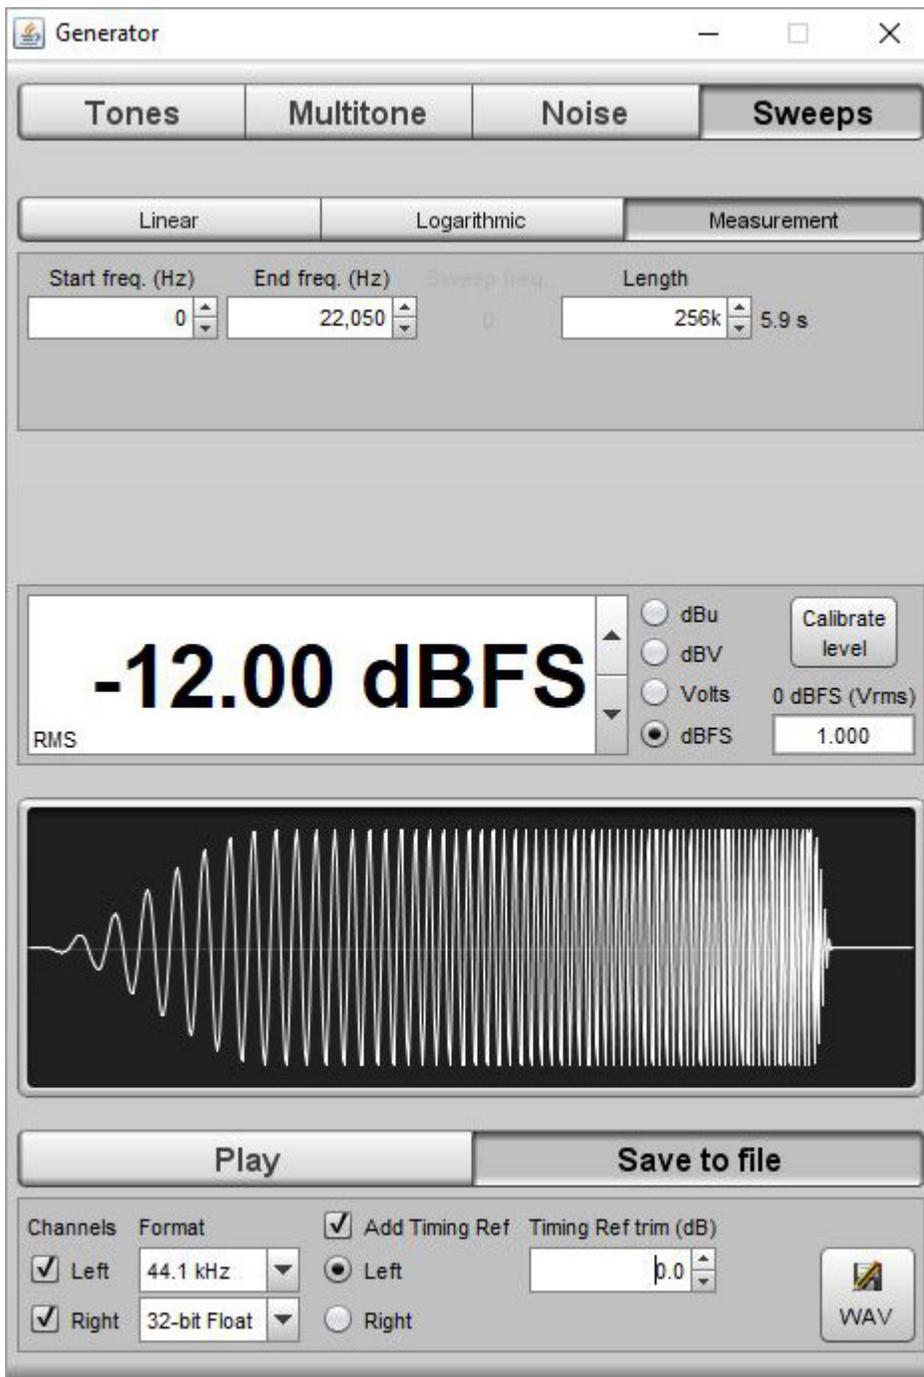
Using a timing reference signal is recommended for offline measurement. Set the signal level for the sweep and the timing reference trim (typically the trim will be zero) and choose the sample rate at which the sweep signal will be generated - this should be the same as the sample rate at which the response will be recorded, or an integer multiple or fraction of the recording rate. Choose 32-bit float as the WAV sample resolution if the system you are measuring supports that, if not try 24-bit.

Importing sweep recordings

Bring up the import dialog by using the Import Sweep Recordings entry in the File menu (or using the Ctrl+Shift+N shortcut).
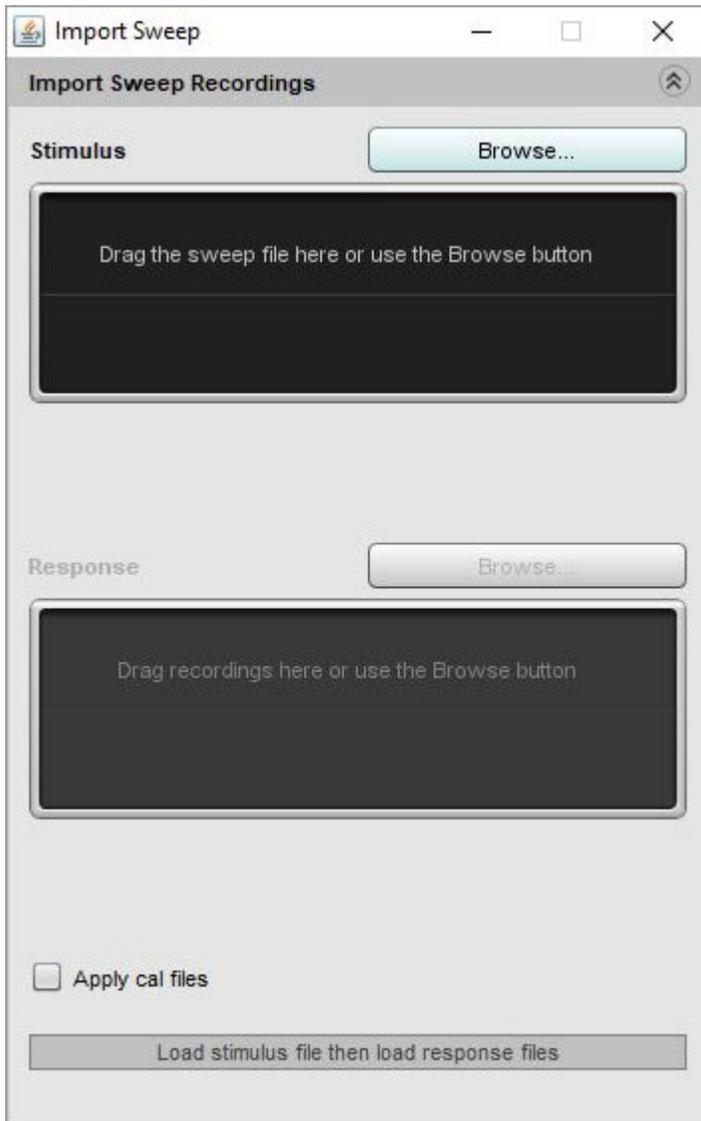
Load the measurement sweep file first, either browsing to the file location or dragging it onto the stimulus window. You will be asked to choose which channel to import, 1 is the left channel, 2 is the right. If the current mic and soundcard cal files specified in the REW cal files preferences should be applied to the files you are importing (for example, if you made the recordings with a calibrated mic and the mic cal file is loaded in REW) select the **Apply cal files** box. An image of the signal waveform is displayed.

Now load the recorded response file. You can browse to the file or drag the file onto the response window. If you have more than one response recording you can drag them all onto the response window at once and a separate measurement will be generated for each. If the recording sample rate is an integer multiple or fraction of the measurement sweep sample rate the recording will be resampled to match the sweep rate before processing.
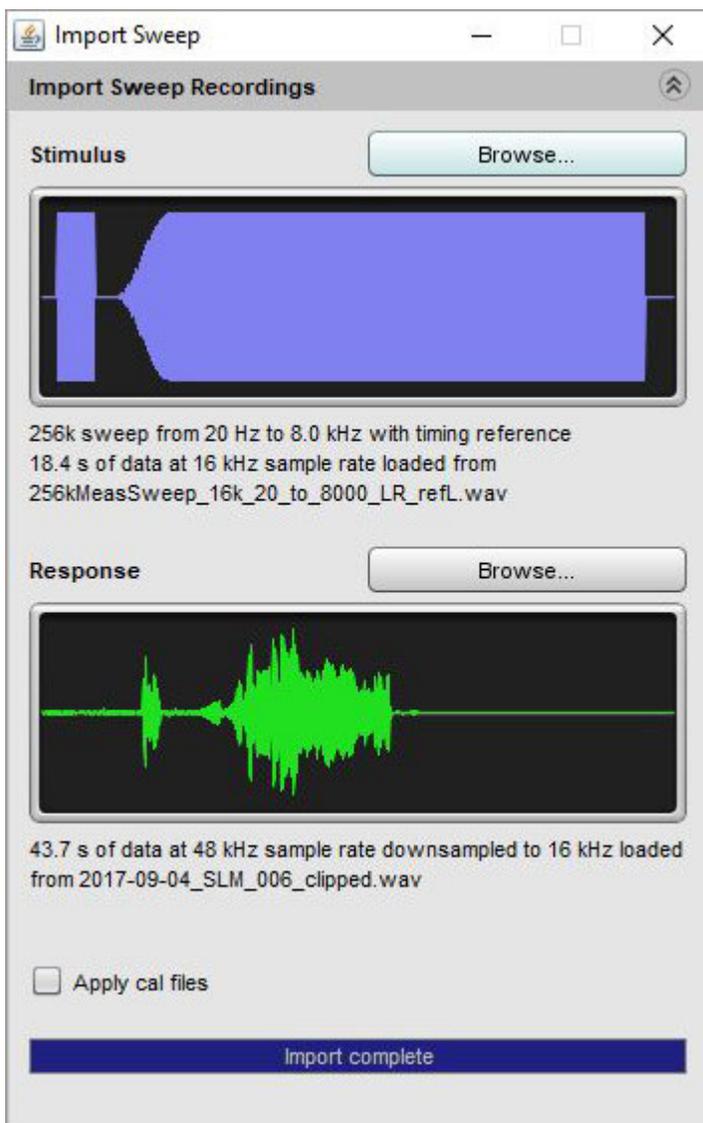
That's it, a new measurement will be generated from the recordings and can be processed in REW like any other measurement. The measurement notes will include the file names used to generate it and the measurement name defaults to the file name of the response recording.

[Help Index](#)

Impedance Measurement

REW can make measurements of impedances up to a few hundred ohms by using both inputs of the soundcard. Impedance measurements of drive units can be used to calculate the Thiele-Small parameters. The general connection arrangement for impedance measurements is shown below:
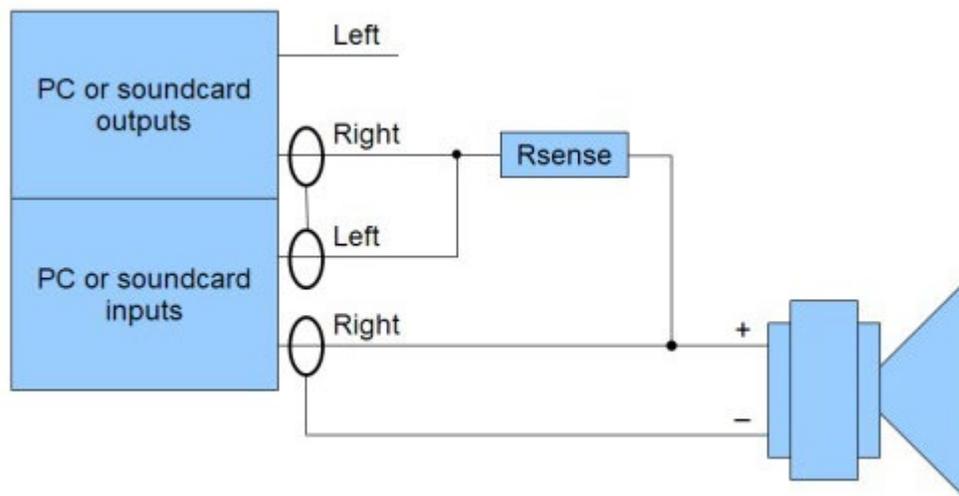
The sense resistor, **which must be non-inductive**, is used to measure the current flowing into the load, which will be (Vleft - Vright)/Rsense. The sense resistor should have a power rating sufficient to cope with a short circuit load and have a low temperature coefficient. The voltage across the load is Vright, so the impedance is nominally voltage/current = Rsense*Vright/(Vleft - Vright). Note that prior to completing the calibration process the accuracy of impedance measurements is only as good as the accuracy of the value entered for the sense resistor.

The load appears in parallel with the input impedance of the soundcard and the impedance of the test leads. Before calibration those effects are partially corrected for by using the parallel equivalent of the sense resistor and soundcard input resistance in the calculation, the soundcard input resistance figure is entered on the dialog as R_{INPUT}.

Good results can be obtained using a headphone output (or, even better, a good headphone amplifier or a device with a high power headphone output) to drive the load, with a 100 ohm sense resistor. If a line output is used the sense resistor typically needs to be larger as line outputs have high output impedance and limited drive capability, try 1 kohm but note that the results will have much higher noise levels and be more susceptible to background noise. If the device being used to drive the load can cope with lower loads a lower sense resistor, 47 ohm for example, will improve the results.

An alternative is to drive the load via a power amplifier, which can deliver the lowest noise levels and most accurate results, but **great care must be taken as the levels a power amplifier can generate can easily damage soundcard inputs.** If using a power amplifier the sense resistor can be much lower, 33 ohms or less, but the soundcard inputs should be connected via a resistive divider providing around 20dB of attenuation and ideally the inputs should also be protected by back-to-back zener diodes to clamp the input to less than 5V.

The soundcard input connected to the load must be the same one which has been chosen as the input in the REW soundcard settings. In the diagram above that is the right input, but if the left is being used simply swap left and right in the diagram. *If the left and right connections are the wrong way around the impedance measurements will show curves that are shifted up by approximately the value of the sense*

resistor.

The input channels should have the same gain. If the audio interface has individual channel gain controls adjust them to have the same gain before starting to measure. An easy way to do that is to connect the impedance rig with the test leads left open and use the signal generator to play a sine tone at 1 kHz at the intended measurement level while observing the input levels on the Level meters. Adjust the input gains so the input and reference channel levels match to within 1 dB and are not clipping.

Sweeps of 10 seconds or longer are recommended for impedance measurement to provide good signal to noise ratios in the results.

Calibrating the Impedance Rig

To get the most accurate results the impedance measurement rig must be calibrated. The calibration can remove the effects of small gain differences between the soundcard input channels, differences in the frequency responses of the channels, and the impedance of the test leads. Three measurements are needed to fully calibrate the rig, but as a minimum the first (open circuit) calibration measurement should be performed. REW automatically saves calibration data from each measurement in the REW log files directory.

Open circuit calibration

The first step is to calibrate the measurement rig with the test leads open. They should still be connected to the rig, but left open at the load end. This measurement compensates for gain differences between the input channels.

- Press the Measure button (or Ctrl+M) to bring up the Measurement panel, select the Impedance button and press the **Open circuit cal** button - the button text is red if impedance calibration has not been done. Disconnect the test leads at the load end before starting the measurement.

- Press **Start** to make the calibration measurement. The completed measurement shows the level of the measurement channel (usually right) compared to the reference channel, where a reading of 100 Ohms corresponds to 100%, 99 Ohms would be 99% etc. If the difference between the 2 channels is too large (more than 2 dB) the calibration is abandoned as it is likely there is a connection error or the input channels have different gain settings, re-check the connections and channel gains and try again.

- After the open circuit calibration gain differences between the input channels will be corrected.

Short circuit calibration

The second step is to calibrate the measurement rig with the test leads shorted. They should still be connected to the rig, but shorted together at the load end. The measurement compensates for the series impedance of the test leads, it particularly improves high frequency accuracy. If this measurement is **not** carried out the R_{LEADS} box can be used to compensate for the test lead resistance. After this measurement the R_{LEADS} box is hidden since it is no longer used.

- Press the Measure button (or Ctrl+M) to bring up the Measurement panel, select the Impedance button and press the **Short circuit cal** button. Short the test leads together at the load end before starting the measurement.
- Press **Start** to make the calibration measurement.
- After the short circuit calibration the series impedance of the test leads will be calibrated out.

Reference resistor calibration

The third and final step is to make a measurement of a known reference resistor. The reference resistor must be non-inductive and its precise value must be known, errors in the reference resistor will translate directly into impedance measurement errors since it is used to scale all subsequent measurements. It should have a value around that of the impedances to be measured, for speaker impedance measurements a reference of 100 ohm or lower is recommended. The reference resistor measurement compensates for frequency response differences between the input channels. It particularly improves accuracy at the frequency extremes. After this measurement the R_{INPUT} box is hidden since it is no longer used.

- Press the Measure button (or Ctrl+M) to bring up the Measurement panel, select the Impedance button and press the **Reference cal** button. Connect the reference restor to the test leads before starting the measurement.
- Press **Start** to make the calibration measurement.
- After the reference resistor measurement calibration is complete.

The calibration data is saved in the REW log files folder and automatically loaded when REW starts up. If the test leads are changed (swapped for longer or shorter leads, for example) or the sample rate is changed the calibration steps should be repeated.
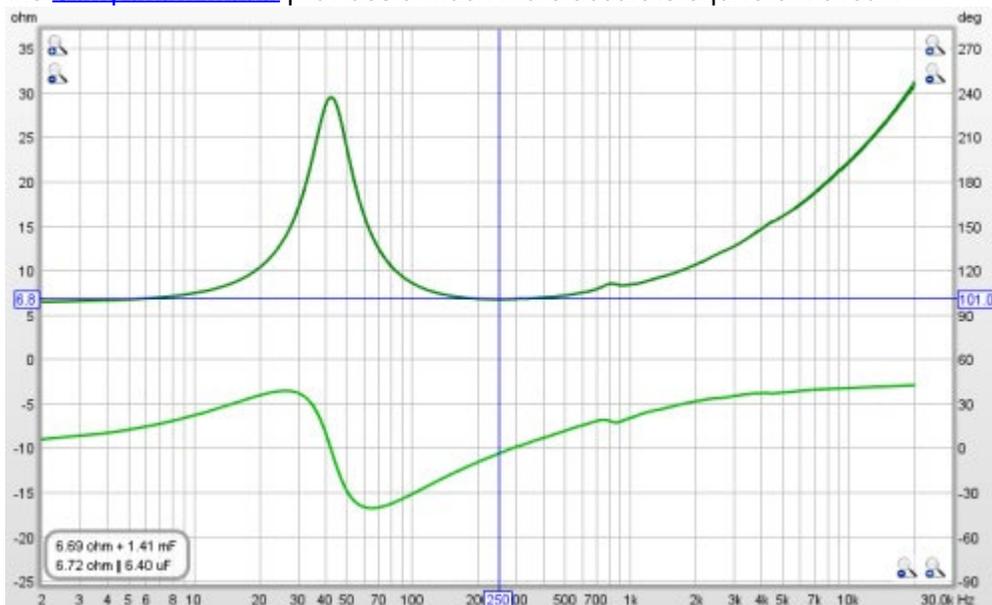
Making an Impedance Measurement

- Press the Measure button [icon] (or Ctrl+M) to bring up the Measurement panel and select the Impedance button

- Enter the **exact** value of the sense resistor. This must be measured accurately with a good quality, calibrated multimeter or impedance bridge, or a very high precision resistor (0.1% or better) should be used. Any error in the value of the sense resistor directly affects the measurement results.

- Set the Start Freq to the lowest frequency for which you wish to see the response and End Freq to the highest. If measuring a drive unit to determine its Thiele-Small parameters measure up to 20 kHz. The sweep will span the range from half the start frequency to twice the end frequency (with an overall limit of half the soundcard sample rate) to provide accurate measurement over the selected range.

- Level controls the rms signal level at which the sweep is generated. The maximum value is -3 dBFS or its equivalent in dBu, dBV or volts, unless the View preference Full scale sine rms is 0 dBFS has been selected, in which case the maximum is 0 dBFS. Using the maximum value places the peaks of the signal at digital full scale - some soundcards may distort at the maximum level. If you are using a power amplifier beware of excessive levels. Loudspeaker impedance is typically best measured at a terminal voltage of around 100 mV to 200 mV at most, the output voltage to achieve that depends on the sense resistor value. For a 100 ohm sense resistor and an 8 ohm driver the output for 100 mV at the terminals would be around 1.35 V.

- Settings for the measurement method are on the right hand side. Length controls the length of the sweep, specifying the number of samples in the sequence. The default is 256k. Dividing the number of samples by the soundcard's sample rate gives the sweep duration in seconds, shown to the right. The overall duration includes silent periods before and after the sweep.

- If Sweeps is more than 1 REW uses synchronous pre-averaging, capturing the selected number of sweeps per measurement and averaging the results to reduce the effects of noise and interference. The pre-averaging improves S/N by almost 3dB for each doubling of the number of sweeps. Averaging is particularly useful if the measurements are contaminated by interference tones, whether electrical or acoustic, as they typically will not add coherently in the averaging and hence will be suppressed by the process. **The output and inputs must be on the same device if using multiple sweeps.**

- The **Check Levels** button generates a few seconds of a pink noise signal that spans the frequency range selected for measurement and checks that the input level is not too high or too low. Pressing **Cancel** while the pink noise signal is playing will turn it off (it turns off automatically after 3 seconds). The rms level that was measured is shown on the measurement panel, with a warning if the level is too high or too low.

- Press **Start** to make a measurement, progress is shown on the measurement panel along with a display of the measurement headroom.

The result of the measurement is displayed in the graph area, information about the measurement appears in the Measurements Panel. Measurements are given a default name of the date and time at which they are made, a more appropriate name can be entered in the box at the top of the measurements panel.

For impedance measurements when the mouse cursor is within the graph panel the equivalent series resistance + inductance or resistance + capacitance and parallel resistance||inductance or resistance||capacitance of the impedance at the cursor position is shown at the bottom left corner of the graph. This may be useful when making measurements of inductors or capacitors to check their value, but the Component Model provides a much more accurate equivalent circuit.

For details of the various ways of viewing the measured data, including averaging multiple measurements, refer to the Graph Panel help.

Input Channels Swapped

If the input channels have been connected the wrong way around the impedance measurements will be too high by approximately the value of the sense resistor, make a test measurement of a resistor (of less than 100 ohms) to check that everything is wired correctly. Swapped channels will also shift the phase by about 180 degrees.

Impedance Measurement Quality

The main source of measurement noise is acoustic noise and vibration during the measurement. Loudspeakers act as microphones, generating small voltages in response to sounds and vibrations that are picked up as part of the load voltage. To minimise this effect use long sweeps, low sense resistor values, avoid noisy environments and isolate the loudspeaker from vibration. Using a power amplifier to drive the speaker provides a much lower drive impedance, which reduces the effect of noise, and allows a low series resistor to be used.

REW provides a **Noise filter** option which applies a bandpass filter to the captured sweep, synchronised to the sweep frequency, to reduce the influence of noise. The **High** setting is suitable in most circumstances, but it may have a slight smoothing effect on very sharp resonances. If that is observed try the **Medium** or **Low** settings, or turn the filter off. When the filter is being used the sweep duration should be at least 10 seconds, a warning is shown if a shorter sweep is selected. Longer sweeps improve signal to noise ratios in the results with or without the noise filter.

[Help Index](#)

Multi-input capture

Multi-input capture is a Pro upgrade feature that allows REW to capture multiple inputs simultaneously for measurement or for the RTA. Each input can have its own calibration file. Individual SPL metering is available for the first four inputs, the level meters can show the input level for the first 16 inputs. Measurements made with multi-input capture will show an rms average of all the captured channel magnitudes with an underlay showing the span of the individual channels and can optionally include individual measurements for each channel.
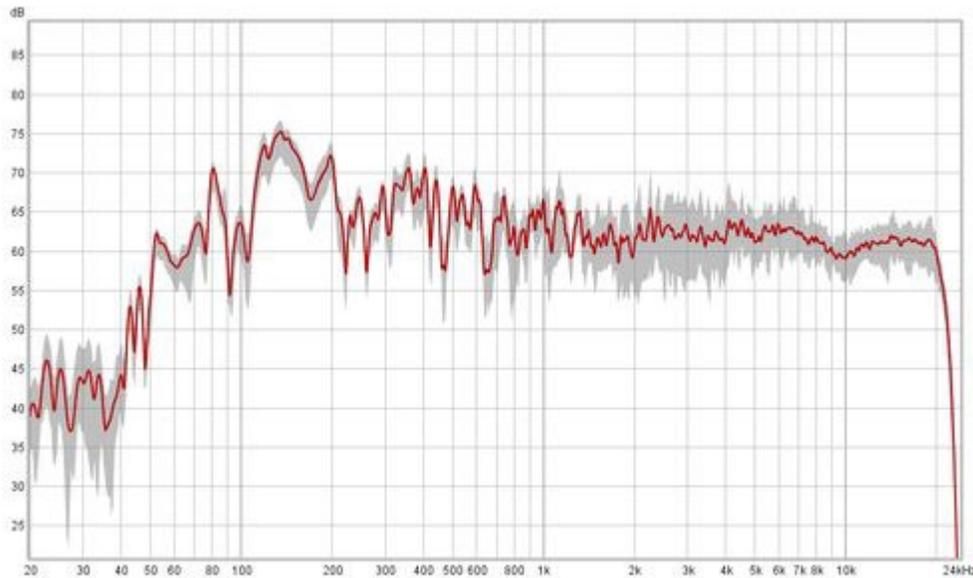
Weighting and SPL alignment

The weighting of each input in the average can be configured and inputs can be aligned to a common SPL in a chosen frequency range before being averaged, useful for spatial averaging. Weighting and SPL alignment can be configured before the measurement is made (using the **Configure averaging**) button on the Measure dialog) or adjusted afterwards (using the **Adjust average** button in the SPL & Phase graph controls). Four preset weighting configurations can be saved.

Adjust average for UMIK-X 16 inputs

SPL alignment

- ◉ Use average SPL of measurements
- ○ Use SPL specified below

Alignment level (dB):   85.0

Available range: 1.83 Hz to 23.98 kHz

Alignment centre: 1 kHz ▼

Alignment span (octaves):   1 ▼

[Align SPL and update average]

[Revert to measured]

Input weightings

| [Load Preset 1] | [Load Preset 2] | [Load Preset 3] | [Load Preset 4] |
| [Save Preset1] | [Save Preset2] | [Save Preset3] | [Save Preset4] |

1: 1.00	2: 1.00	3: 1.00	4: 1.00
5: 1.00	6: 1.00	7: 1.00	8: 1.00
9: 1.00	10: 1.00	11: 1.00	12: 1.00
13: 1.00	14: 1.00	15: 1.00	16: 1.00

[Apply weighting and update average]   [Reset]

Multi-input RTA measurements

RTA measurements with multi-input capture show an rms average of the individual channel magnitudes and individual traces for the first 16 inputs. The maximum number of input traces to show is configured in the View preferences, the processing and graphics capabilities of the computer may mean fewer than 16 can be displayed. RTA measurements are saved in the same format as sweep measurements allowing the same weighting and SPL alignment adjustments to be made to them after measurement have been captured. Weighting and SPL alignment can be configured before or during capture using the **Configure average** button on the RTA input selection dialog.
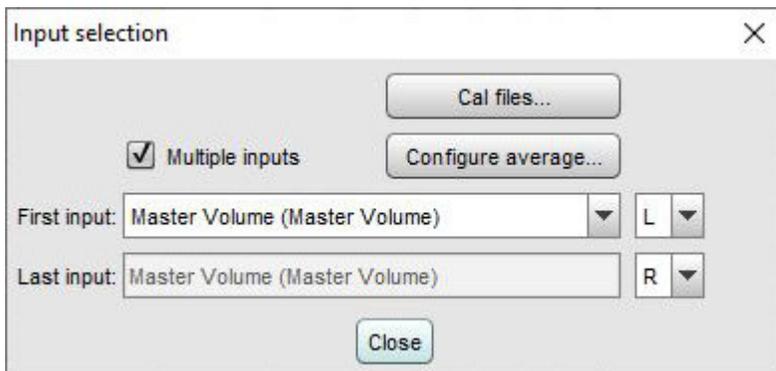
Input selection

The first and last channels to be captured are selected on the Soundcard preferences, Measure dialog or RTA input dialog. The example below shows selection using ASIO drivers, a similar selection is offered for Java drivers. Note that Java drivers on Windows are limited to two channels, ASIO must be used for multi-channel devices. On macOS or Linux Java drivers offer as many channels as the interface supports.
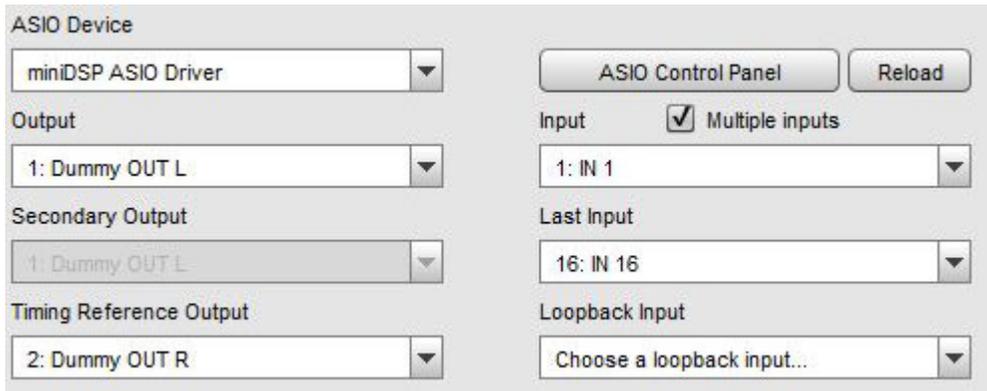
Microphones

Multi-input capture would typically be used with a mic array (such as the miniDSP UMIK-X) or a multi-channel interface with multiple analog microphones connected to it. REW captures a contiguous range of input channels. Whilst it is possible to use this feature with multiple USB microphones on Windows when using an ASIO wrapper, this is **NOT** recommended because:

- REW would only be able to access the input volume setting for the first USB mic and would use that to adjust SPL values for all the mics, so all the USB mics would need to have the same OS volume setting
- USB mics typically provide a stereo pair of inputs, each carrying the same data, so there would be two identical input captures for each mic

When using the miniDSP UMIK-X there is an option on the Measure dialog to produce an average of the four mics on each UMA4 mic array as well as an overall average of all the selected inputs. The array average measurement names will be prefixed with "A1-" for the first array, "A2-" for the second and so on.

[Help Index](#)

Thiele Small Parameters

The Thiele-Small Parameters window is used to calculate the parameters for a drive unit from measurements of its impedance. Three methods are offered: sealed box, added mass and dual added mass. The dual added mass method if from the paper "An Added-Mass Measurement Technique for Transducer Parameter Estimation" by Jeff Candy and Claus Futtrup, JAES Volume 65 Issue 12 pp. 1005-1016; December 2017. It gives the most accurate results, but all the methods are dependent on the quality of the measurements.

Each method requires a "free air" measurement, with the driver rigidly supported, without any baffle, and mounted vertically (i.e. so that the cone is firing horizontally as it would be in a typical loudspeaker installation). Additional measurements are with either mass added to the cone or with the unit in a sealed (air tight!) enclosure (ideally with a volume a little less than the expected Vas).

For single added mass measurements the mass added should be approximately half the expected Mms figure. For dual added mass make one measurement with masses that total a little less than the expected

Mms, then one with half the masses removed, then finally the free air measurement. Some pre-conditioning of the unit with signals at medium levels helps to stabilise the behaviour and suspension compliance and reduce memory effects in the suspension from periods of storage or lack of use. Quiet conditions are important for good measurements, drive units act as microphones and pick up noise and vibration, affecting the results. The measurements should be made up to 20kHz so that the lossy inductance of the voice coil can be accurately modelled, and the impedance calibration step should be carried out before making measurements.

Added mass measurements require very precise figures for the masses to get accurate results. The more accurate the scale used, the better. It is recommended that the overall mass to be added is split into 4 pieces, of approximately equal mass, placed equally spaced around the cone near the voice coil and firmly attached. Blu-tack works well, either on its own or with a nut embedded in it when more mass is required. Do not place the masses near the outer edge of the cone, that will adversely affect the cone behaviour. For the second measurement in a dual added mass run a pair of the masses are (gently!) removed, leaving two on opposite sides of the cone.

Compensate for leakage losses and **Compensate for Air Load** check boxes are shown for sealed box measurements, they take into account the leakage loss of the sealed box (which would be shown as Q_l in the results) and the air mass load due to the sealed box. These compensations use the Carrion-Isbert method described by Claus Futtrup in the documentation for his *Driver Parameter Calculator* application at http://www.cfuttrup.com/

An Example Run

The results show the process for TS parameter calculation of a mid-bass drive unit. The plots below show the impedance measurements made in free air and then with the added masses. A least squares fit of an electrical model of the drive unit impedance is carried out on the free air measurement to determine the model parameters. Another least squares fit is carried out on the secondary measurement to determine the modified motional parameters and the TS parameters are then calculated.
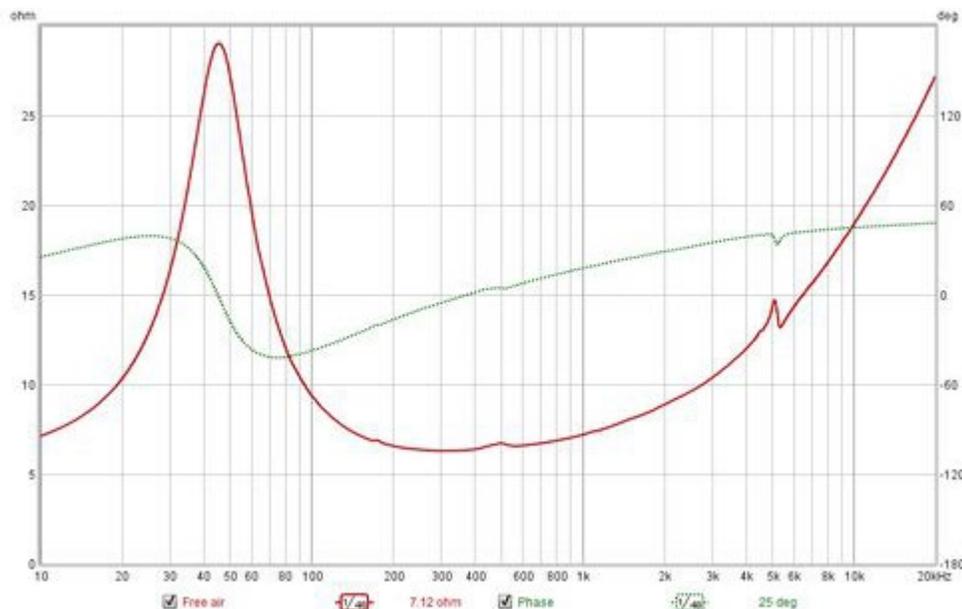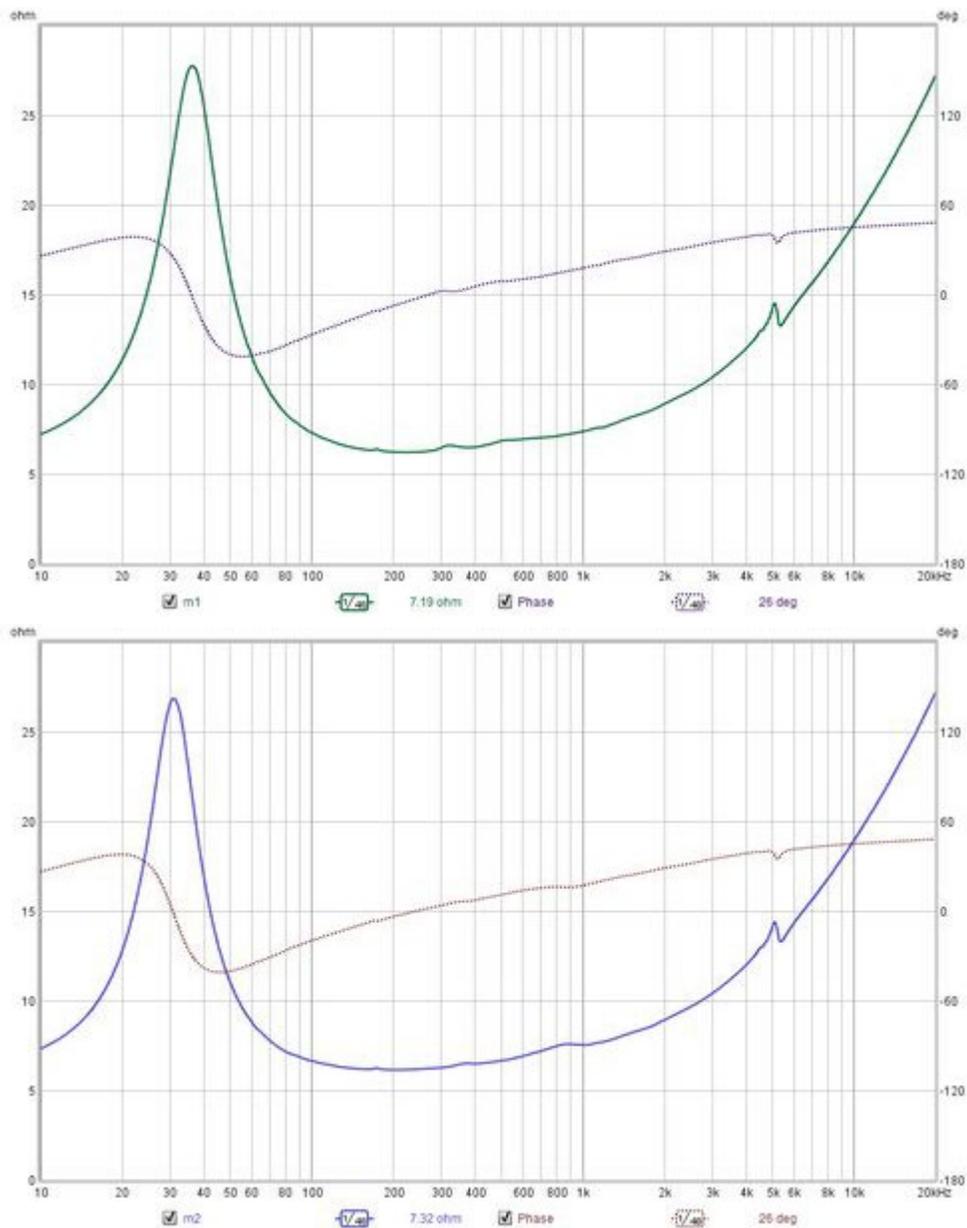
To calculate the TS parameters the measurements are selected and the required values are entered:

- the DC resistance of the voice coil (R_{DC}) in ohms. Accurate measurement of low resistances is not easy, but the impedance model REW uses can compensate for a DC resistance which is not precise.
- the effective area in square centimetres, most driver data sheets include an effective area figure but if this is not available REW can calculate the figure for you given the effective diameter, which is the diameter of the cone plus a proportion of the surround, typically 1/3 to 1/2, just click the calculator icon on the left hand side of the effective area box
- the air temperature in degrees Celsius
- the air pressure in millibar
- the the additional masses in grammes or volume of the sealed box in litres.

The **Calculate Parameters** button is then clicked, with the following results.

Thiele-Small Parameters

Measurement method

Dual added mass ▼

Free Air Measurement

1: Free air ▼

Added mass measurements

2: m1 ▼

Added mass 1 (g): 8.017

3: m2 ▼

Added mass 2 (g): 16.048

Motional Impedance (Ritter 3PC)

R_0 (ohm):	27.95
C_{MES} (uF):	303.7
L_0 (mH):	38.845
β:	0.0926
ω_0:	430.7

Blocked Impedance (T-F)

dR (ohm):	0.062
L_{EB} (uH):	31.1
L_E (mH):	0.773
R_{SS} (ohm):	100000.00
K_E (S-H):	0.0702

Manually Entered Values

Voice Coil DC Resistance (ohm):	5.721
Effective Area (cm^2):	104.0
Air Temperature (Celsius):	20
Air Pressure (mB):	1013.25

Calculate Parameters

Write Parameters to File

Simplified Model Parameters

R_E	5.783 ohm	L_E	134.5 uH
R_{ES}	23.13 ohm	R_2	9.88 ohm
C_{MES}	303.7 uF	L_2	196.9 uH
L_{CES}	41.32 mH	R_3	1.54 ohm
		L_3	361.0 uH

R_E	5.783 ohm	f_S	44.9 Hz	M_{MS}	14.91 g
Z_{min}	6.296 ohm	Q_{MS}	2.396	C_{MS}	0.841 mm/N
f_{min}	313 Hz	Q_{ES}	0.496	R_{MS}	1.757 kg/s
f_3	2,038 Hz	Q_{TS}	0.411	V_{AS}	12.91 litres
$L_E (f_3)$	0.339 mH	F_{TS}	109.4 Hz	Bℓ	7.008 Tm
Dd	11.51 cm	L_P	85.79 dB (1W/1m)	Eta	0.23 %
Sd	104.0 cm^2	m1	8.017 g	m2	16.048 g

Secondary measurement:
Air temperature 20.0 C, pressure 1,013.25 mB giving density 1.2041 kg/m^3, c 343.2 m/s

The first column of results in the bottom of the window show the loudspeaker resistance R_E, which is generally a little higher than the DC resistance; the minimum impedance Z_{min} after the peak and the frequency f_{min} at which it occurs; f_3, which is the frequency at which the impedance has increased to sqrt(2)*Zf_{min}; the inductance at f_3; the effective diameter and the effective area. The second column shows the resonant frequency f_S; the mechanical (Q_{MS}), electrical (Q_{ES}) and total (Q_{TS}) Q-factors and the F_{TS}

figure (f_S/Q_{TS}). These parameters can also be calculated for any single measurement, without requiring a secondary measurement to be selected. The L_P figure and the M_{MS}, C_{MS}, R_{MS}, V_{AS}, $B\ell$ and Eta (efficiency) figures in the third column can only be calculated using extra measurements.

The results can be copied to the clipboard by right-clicking on the results area, or written to a text file using the **Write Parameters to File** button. When writing to file the separator between values, labels etc is as defined in the File → Export menu.

The plot below shows the modelled impedance traces overlaying the measured values for the free air measurement.
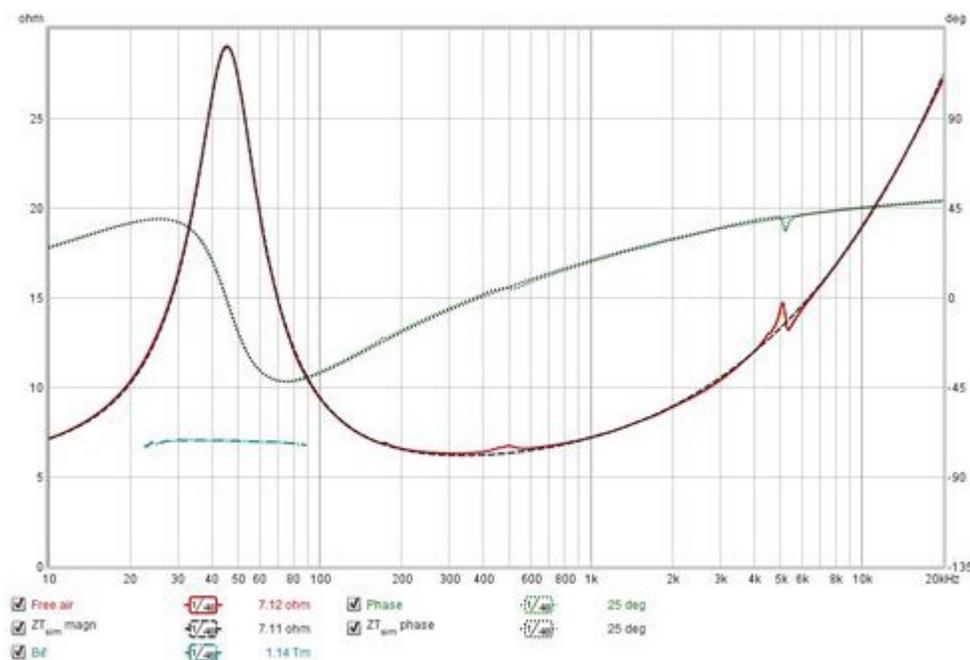
Dual added mass measurements include a $B\ell$ curve (cyan in the image), it should be close to horizontal around the resonant frequency if the measurements are good.

When TS parameters have been calculated the derived and simulated motional and blocked impedance magnitudes and phases can be plotted in addition to the total impedance traces, along with the ZM* magnitude and phase for dual added mass results. The simulated traces are produced using the model parameter values, the derived traces are produced by subtracting the model values from the measured values (for example, derived motional impedance is produced by subtracting the modelled blocked impedance from the total measured impedance).
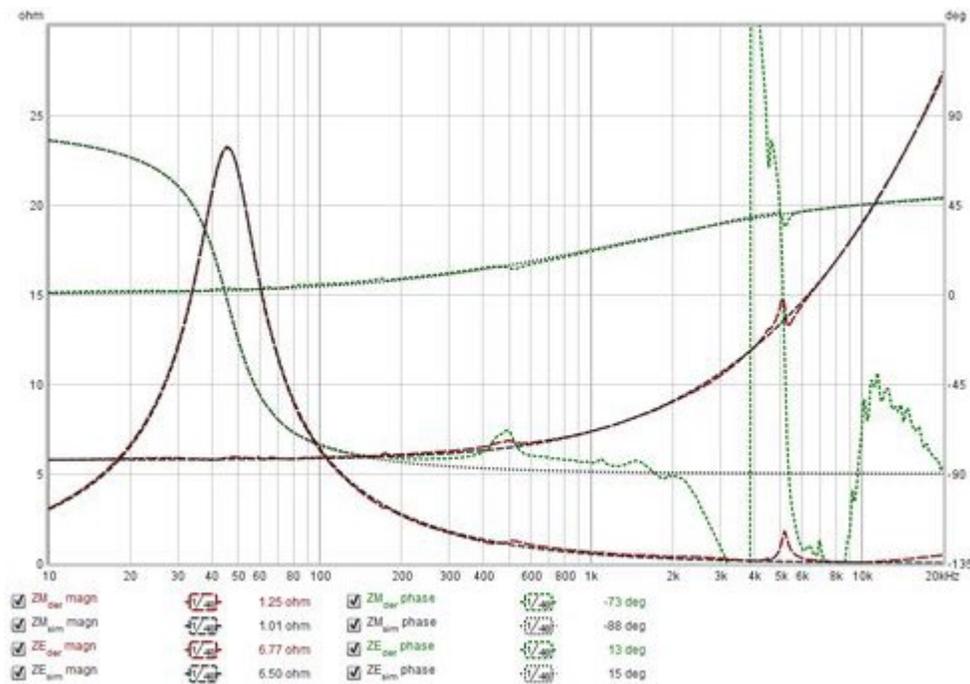
Electrical impedance model

REW uses a model for the electrical impedance component of the driver based on "Electrodynamic Transducer Model Incorporating Semi-Inductance and Means for Shorting AC Magnetization" by Knud Thorborg and Claus Futtrup, JAES Volume 59 Issue 9 pp. 612-627; September 2011. The diagram below shows the components of the electrical impedance model.
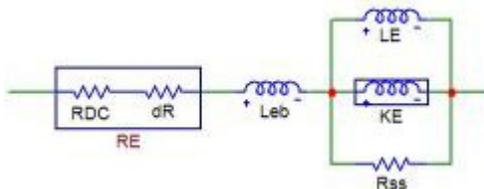
The model begins with a drive unit resistance R_E which is the DC resistance R_{DC}. REW has an additional resistance dR which allows for the difference between the actual RE and the measured resistance, which may differ due to measurement errors or coil temperature. It is followed by a series inductance L_{EB} and then a parallel combination of an inductance L_E, a semi-inductance K_E and a resistance R_{SS}. L_E represents the inductance of the part of the voice coil located inside the motor gap. L_{EB} represents the part of the coil outside the motor gap. The semi-inductance K_E has an impedance that varies with the square root of omega. It models the effects of eddy currents and skin depth in the pole piece. The parallel combination of L_E and K_E models the transition of the coil's behaviour from largely that of a conventional inductor at low frequencies to a semi-inductor at high frequencies. The R_{SS} component models the effect of electrically conductive material in the magnet system. The parameter values REW determines may be modified if desired and the effect on the modelled impedance and phase traces viewed on the graph, but the TS parameters which have been calculated will not be altered.

Mechanical impedance model

The mechanical impedance model incorporates elements that cater for the frequency-dependence of compliance. It uses the LOG model of viscoelasticity, from "Low-Frequency Loudspeaker Models That

Include Suspension Creep" by Knudsen and Jensen, JAES Volume 41 Issue 1/2 pp. 3-18; February 1993, with the added Retardation Spectra function as described in "Modeling Viscoelasticity of Loudspeaker Suspensions Using Retardation Spectra" by Agerkvist and Ritter, AES Convention: 129 (November 2010), Paper Number: 8217.

Simplified Model

As frequency-dependent component values are not supported by many circuit simulators REW also calculates values for an alternative blocked impedance model using two parallel resistor-inductor pairs in series, labelled R2-L2 and R3-L3, and the conventional R_{ES}, C_{MES}, L_{CES} motional impedance model without the frequency-dependent damping and compliance. The values of these components are shown in the "Simplified Model" box. This diagram shows the simplified model components.
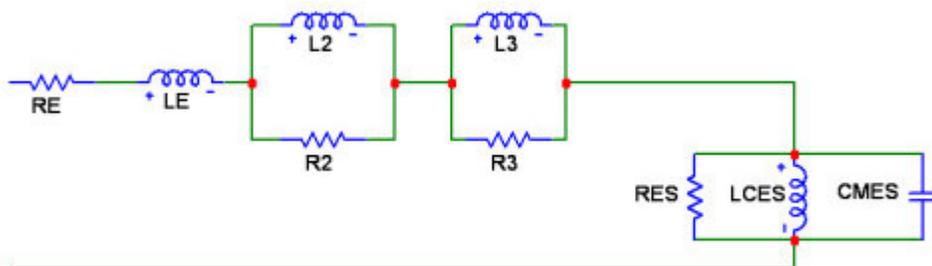
[Help Index](#)

Hints and tips

Here are a few hints and tips for using REW

- REW's mdat files can be opened by double-clicking on them or by dragging them onto the REW window
- Impulse response files in .wav or .aiff format can be imported by dragging them onto the REW window
- The measurement list can be navigated by using Alt+UP and Alt+DOWN to go up and down the list, or any of the first 9 measurements can be selected directly using ALT+1 or 2 or 3…9
- A measurement can be moved up or down in the list by selecting it then pressing and holding the left mouse button and dragging the measurement up or down to where you want it
- The graphs can be navigated by using Ctrl+RIGHT and Ctrl+LEFT to go to the next or previous graph, or Ctrl+1 or 2 or 3…9 will go to the 1st, 2nd etc graph
- Moving the cursor over the name of a trace in the graph legend will highlight that trace in the graph
- If a measurement's name is shown in blue on the measurements panel it means it has been modified since it was loaded
- Use Ctrl+Shift+3 to apply 1/3rd octave smoothing, repeat to remove the smoothing. Ctrl+Shift+X applies variable smoothing
- If you accidentally delete a measurement use the Restore Last Removed entry in the file menu to recover it
- You can clear or select all traces in an overlay graph by right clicking in the legend area
- The graph can be moved by pressing the right mouse button and dragging
- The mouse wheel can be used to zoom in or out on a graph
- Variable zoom is available by pressing and holding the right mouse button, then while the right button is down pressing the left button, then drag the mouse
- You can measure between points on a graph by holding down the Ctrl key, pressing the right mouse button and dragging the mouse
- There are also keyboard graph controls: the cursor keys move the graph cursors; shift+cursor keys move the graph; x and shift+x zoom out/in on the x axis, y and shift+y zoom out/in on the y axis
- There are keyboard shortcuts for most actions, listed in the shortcuts section
- The space bar can be used to press the currently highlighted button instead of the mouse
- To make more space for the graph (e.g. at lower screen resolutions) use the expand/collapse button above the measurement panels to reduce the width they take up and turn off the graph scroll bars using the button above the graph. You can also remove the labels under the toolbar and graph buttons by unchecking boxes in the View preferences to make more vertical space. You can remove the toolbar by unchecking the **Show toolbar** box in the View preferences (everything in the toolbar has a corresponding menu entry and a keyboard shortcut).

Help Index

Measurements Panel

The Measurements Panel shows the measurements which have been made or loaded and information about them. The tabs on the panel are used to select individual measurements, they include a thumbnail of the frequency response. The text next to the thumbnail shows the name of the file the measurement was loaded from or has been saved to (if it has been saved), the date and time the measurement was made and the mic/meter and soundcard calibration files that were used (hover the mouse over that text area to bring up a tooltip that shows the full text width). The order of the measurements in the list can be changed by clicking on the currently selected measurement and dragging it up or down to a new position in the list.

Collapse

The measurement panel can be made narrower to provide more screen area for the graph, click the **Collapse** button to narrow the panel.
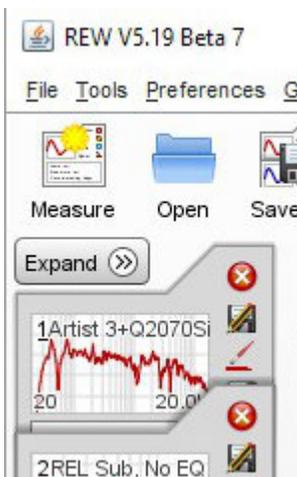
Measurement Controls

The text box at the top of the measurement panel is used to change the name of the measurement. The length of the name is limited to the width of the box. If a blank name is entered, "No Description" is used. The name is shown in blue for new measurements or measurements with unsaved changes, otherwise it is black.

The buttons in the column next to the name box are used to delete the measurement (deleting a measurement removes it from REW, saved measurement files are not affected), save the measurement, set the trace colour and expand or collapse the notes area. Trace colours can be reset to their defaults in the View Preferences.

Notes relating to each measurement can be entered in the notes area, they are saved with the measurement. Right-clicking in the notes area brings up a Cut/Copy/Paste menu.

Change Cal

The **Change Cal** button brings up a dialog to load, update or clear the calibration data files for the mic/meter or soundcard or to select whether a C-weighted SPL meter was used for the measurement. Note that changing the calibration data or C-weighted setting will clear any waterfall, decay or spectrogram plots that have been generated as they would no longer show valid data. Other affected plots will be regenerated using the new cal data.
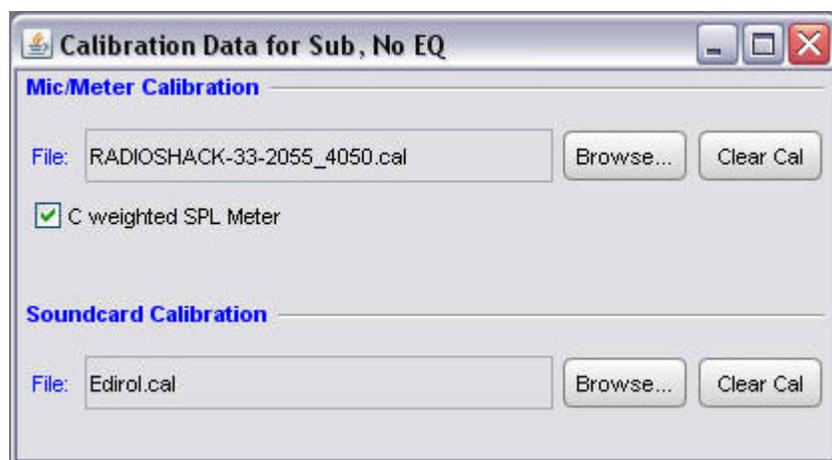
Help Index

Impulse Responses

Interpreting impulse responses is an important part of acoustic analysis. An impulse response measurement can tell us a great deal about a room and the way sound will be reproduced within it. It can show us what kinds of treatment will be helpful and whether treatments have been correctly applied to achieve the best results. This page explains impulse responses, the information that can be extracted from them and how REW can measure and analyse such responses.

What is an impulse response?

Before we can get very far in interpreting an impulse response we need to understand what an impulse response is. The impulse response is in essence a recording of what it would sound like in the room if you played an extremely loud, extremely short click - something like the crack of a pistol shot. The reason for measuring the impulse response (by more subtle means than firing a gun in the room) is that it completely characterises the behaviour of the system consisting of the speaker(s) that were measured and the room they are in, *at the point where the measurement microphone is placed.* An important property of an impulse, not intuitively obvious, is that it if you break it up into individual sine waves you find that it contains all frequencies at the same amplitude. Strange but true. This means that you can work out a system's frequency response by determining the frequency components that make up its impulse response. REW does this by Fourier Transforming the impulse response, which in essence breaks it up into its individual frequency components. The plot of the magnitude of each of those frequency components is the system's frequency response.

When an impulse response is measured by means of a logarithmically swept sine wave, the room's linear response is conveniently separated from its non-linear response. The portion of the response before the initial peak at time=0 is actually due to the system's distortion - looking closely, there are scaled down, horizontally compressed copies of the main impulse response there - each of those copies is due to a distortion harmonic, first the 2nd harmonic, then the third, then the fourth etc. as time gets more negative. The initial peak and its subsequent decay after time=0 is the system's response without the distortion.
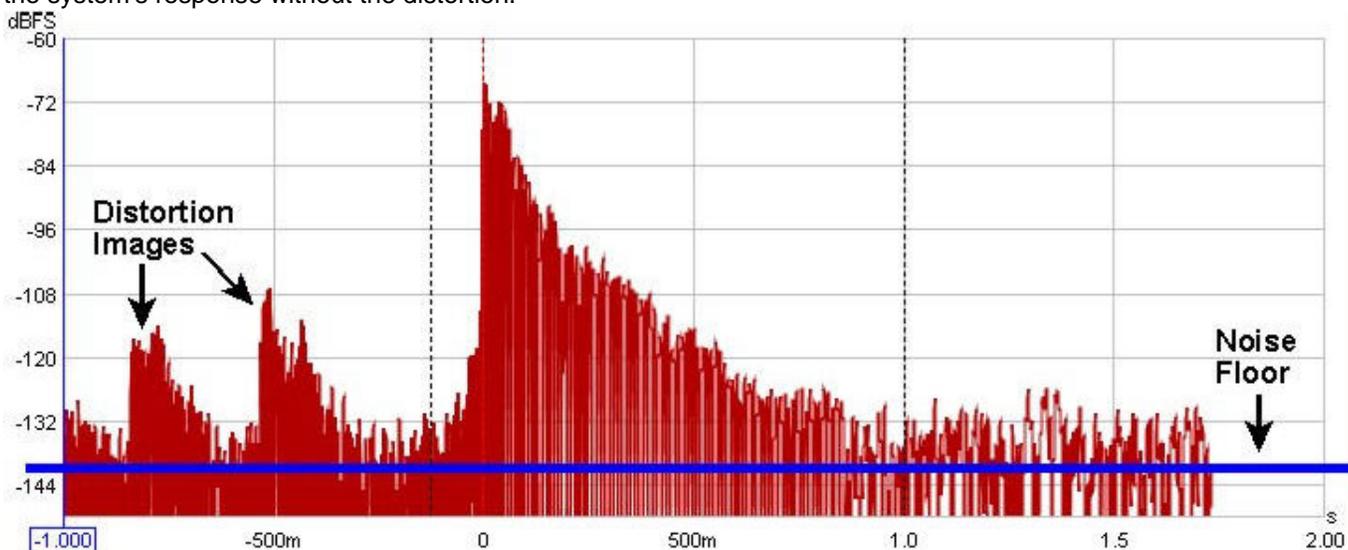

In a perfect system of infinite bandwidth with totally absorbent boundaries, the impulse response would look like a single spike at time 0 and nothing anywhere else - the closest you get to that is measuring the soundcard's loopback response. In a real system, finite bandwidth spreads out the response (dramatically so when measuring a subwoofer as its bandwidth is very limited). Reflections from the room's boundaries add to the initial response at times that correspond to how much further they had to travel to reach the microphone - for example, if the microphone were 10 feet from the speaker and a sound reflection from a wall had to travel 15 feet to reach the microphone, that reflection would contribute a spike (smeared out depending on the nature of the reflection) about 5 ms after the initial peak, because sound takes about 5 ms to travel that extra 5 feet.

When measuring full range responses from loudspeakers (rather than subwoofer responses) the reflections are easier to spot as the higher bandwidth of the full range system keeps the spike of the impulse (and the reflections)

quite narrow, but you need to zoom in on the time axis to see them. They are easier to spot with a linear Y axis (set to %FS instead of dBFS) and also show up more readily with the ETC smoothing set to 0.

Impulse Response Windows

After capturing a sweep, FFT processing is carried out to derive the system's impulse response and the corresponding frequency response. There are controls to adjust the position and widths of left and right windows that define the portion used to derive the frequency response, these controls may be accessed by pressing the "IR Windows" button on the toolbar.
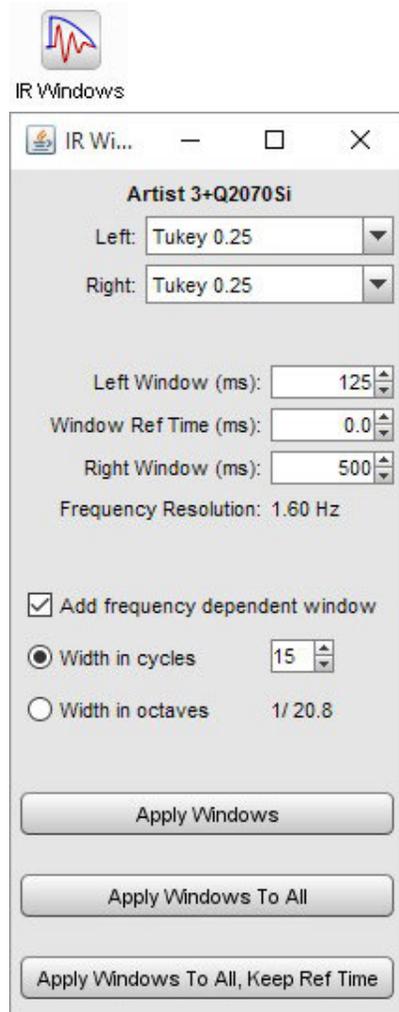
The windows and the region of the impulse response they cover can be viewed on the Impulse graph by selecting the "Window" and "Windowed" traces. The reference position for the windows is usually the impulse peak
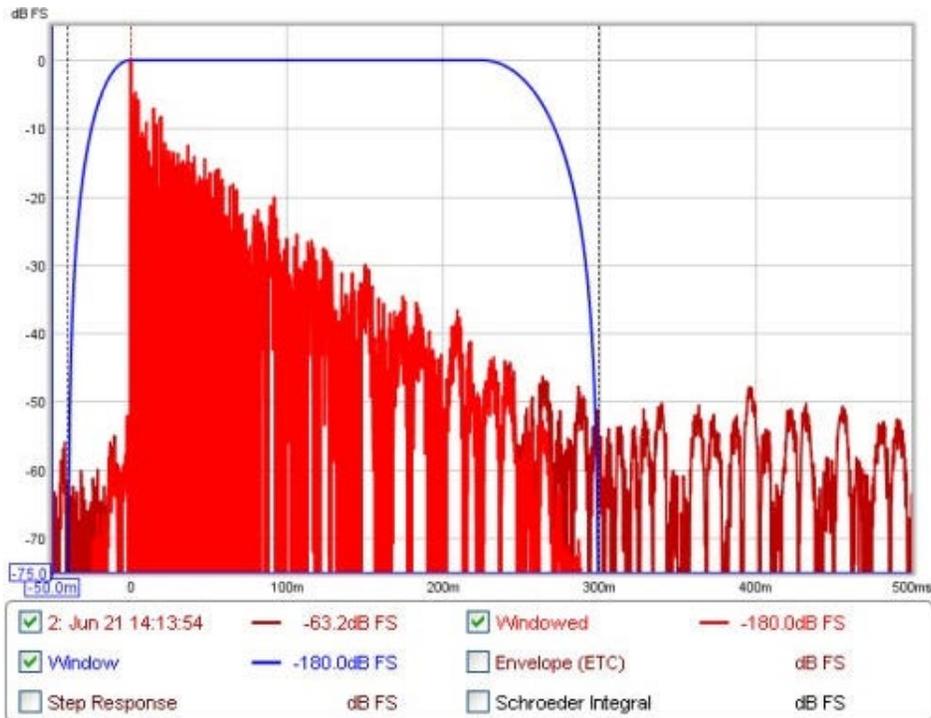
The default settings for the windows will usually be suitable. In smaller rooms it may be necessary to use a shorter right-side window duration, around 300-500ms - if the frequency response plot appears noisy and jagged try reducing the right window period and hit "Apply Windows" to recalculate the frequency response. In very large rooms the window can be increased to improve the frequency resolution.

The frequency resolution corresponding to the current total window duration (left and right combined) is shown above the Apply Windows button - the longer the duration, the higher the resolution. Alternative window shapes may be selected independently for the left and right windows.

In addition to the left and right windows a frequency-dependent Gaussian window can be applied. This is a window whose width varies inversely with frequency, getting progressively narrower as frequency increases. The width of the window can be specified as a number of cycles or an octave fraction. If the width is in cycles then the width (between the half amplitude points of the window) at any frequency will be the number of cycles times the period of that frequency, for example a 15 cycle window will have a width at 1 kHz of 15 * (1/1000) = 0.015 s or 15 ms. The corresponding octave fraction has an effect similar to applying a smoothing of the same octave fraction, except the variable window excludes progressively more of the late arriving sound as frequency increases rather than just averaging it out - this has similarities with the way the ear increasingly picks out the direct sound from the speaker at higher frequencies.

If **Add frequency dependent window** is selected the frequency dependent window is applied after first applying the selected left and right windows. The FDW is centred on the window reference time - for best results this should be at the peak of the impulse.

Help Index

Minimum Phase

In discussions of equalisation, and particularly equalisation applied to attempt to improve the acoustic response in a room, "minimum phase" will often crop up - generally in the context of whether or not EQ can successfully be used to address a response problem. So what is "minimum phase", and why should we care?

There are rigourous mathematical and systems theory definitions of what constitutes a minimum phase system, but I will not repeat them here. In the context of acoustic measurements a system which is minimum phase has two important properties: it has the lowest time delay for signals passing through it and it can be inverted.

Minimum Phase and Time delay

The "lowest time delay" property refers to the amount the frequency components of a signal are delayed *while delivering the measured frequency (SPL) response*. We can see the delay characteristics directly in the Group Delay plot of the system. Given a measured frequency response we cannot tell from the SPL response alone whether what we measured has this "minimum delay" characteristic. If there was a time delay in the overall system somewhere, such as the time taken for sound to travel from the loudspeaker to the microphone, that delay would render the system non-minimum phase (in the strictest sense of the term) but would not alter the SPL response we measured.

A time delay causes a phase shift that increases with frequency - for example, a delay of just 1ms results in a phase shift of 36 degrees at 100Hz but 3,600 degrees at 10kHz, because 1ms is 1/10th of the 10ms period of a 100Hz signal but is 10 times the 0.1ms period of a 10kHz signal, and each period is 360 degrees. The phase shift caused by a time delay is linear with frequency, meaning the 1ms example would give 36 degrees of phase shift at 100Hz, and twice that delay at twice the frequency or three times the delay at three times the frequency etc. If the frequency axis is set to linear the phase plot of a time delay looks like a straight line droppping down as frequency increases - how steeply it drops depends on how large the delay is.

Whilst constant time delays make it difficult to interpret the phase response, they can be removed from our measurements and they do not cause any problems with applying EQ. However, just removing time delays (or their effects) is not enough to make a system minimum phase, there is more to it than that.

Minimum Phase and Invertibility

Minimum phase systems can be inverted, which means that a filter can be designed that, if applied to the system, would produce a flat response *and* correct the phase response at the same time. That is clearly a nice property to find if we want to apply EQ. If we apply EQ to a system that is not minimum phase, or more particularly in a region where it is not minimum phase, the EQ will not produce the results we would like. It may still be possible to achieve a flat response, but correcting the phase response would elude us. It is simply not possible.

A simple example of something that renders a response non-minimum phase is reflections that are as large or larger than the direct signal (reflections along paths that are different but the same length can combine to produce higher levels, or a curved surface can focus a reflection). In the simple case of a reflection that is exactly the same amplitude as the direct signal, we would find there were regularly spaced frequencies at which the reflection is 180 degrees out of phase with the direct sound. When those signals combine, the result is zero amplitude at those frequencies (an extreme example of the comb filtering often seen in acoustic measurements). That zero level cannot be restored to what it should have been by any amount of EQ, as the EQ affects the direct and reflected signals equally so the signals still cancel. If a response has regions in it that are zero it cannot be inverted and it is not minimum phase. If the reflection is larger than the direct sound the problem is equally tricky, as although we no longer have a zero level we would end up with a situation where the corrections the EQ is applying would have to keep getting larger to counter the ever larger reflection and we would quickly run out of headroom.
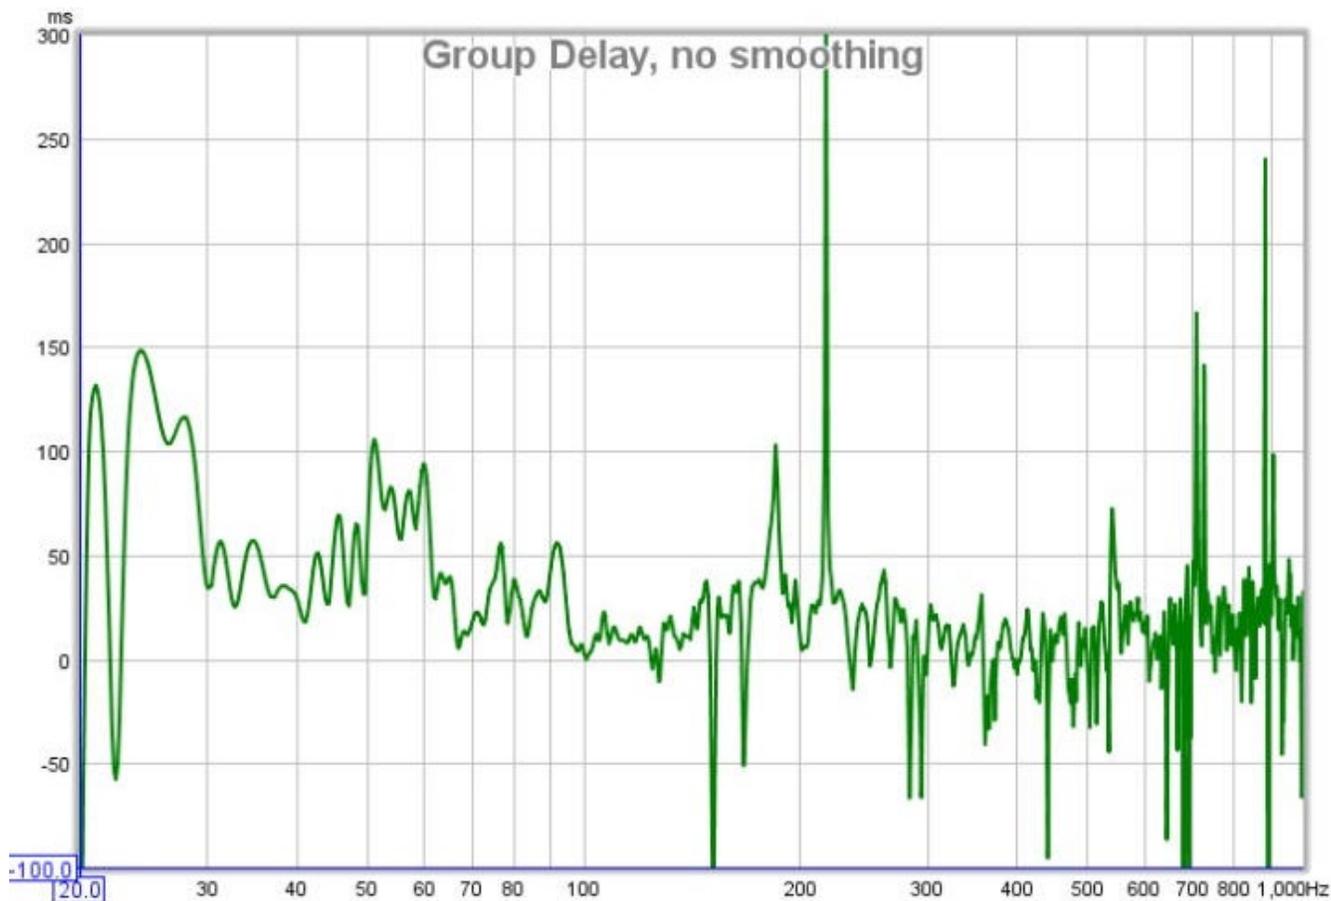
Identifying Minimum Phase Regions

Room responses are mixed phase, meaning there are some minimum phase regions and some regions that are not minimum phase. The minimum phase regions tend to be at lower frequencies, but we cannot simply say a response is minimum phase below some specific cutoff. It is not possible to identify minimum phase regions from looking at the

wrapped phase response, especially if the measurement has any time delay. The unwrapped response gives some more clues, plotted against a linear frequency axis, but often covers such a huge span that it is impractical to use. Even if we remove any time delay in the measurement the phase response alone still doesn't let us easily identify the minimum phase regions. There is a straightforward method, however. Here is a measurement of a sub+main speaker in-room:

We might hazard a guess that this is largely minimum phase below the room's transition frequency, and non-minimum phase above, but to avoid the guesswork we can look at group delay. The group delay plot shows us how much each frequency is being delayed - mathematically, it is the slope of the unwrapped phase plot, so anywhere that phase is dropping linearly corresponds to a constant group delay region (i.e. that region is delayed by a constant time). Here is the group delay plot for the measurement:

That gets us a bit closer, we can speculate that the places where there are particularly wild swings in group delay are not minimum phase, but it still doesn't let us easily identify the minimum phase regions. To do that, we need to compare the measurement with a system that has the same amplitude response but *is* minimum phase and look at the measurement's *excess* group delay. The minimum phase response is generated by using the measurement amplitude and calculating the corresponding minimum phase from it, using a mathematical relationship between the two that holds for minimum phase systems. By looking at the difference between the measured and minimum phase (the excess phase) and measuring the slope of that difference to find the excess group delay, we get this plot:
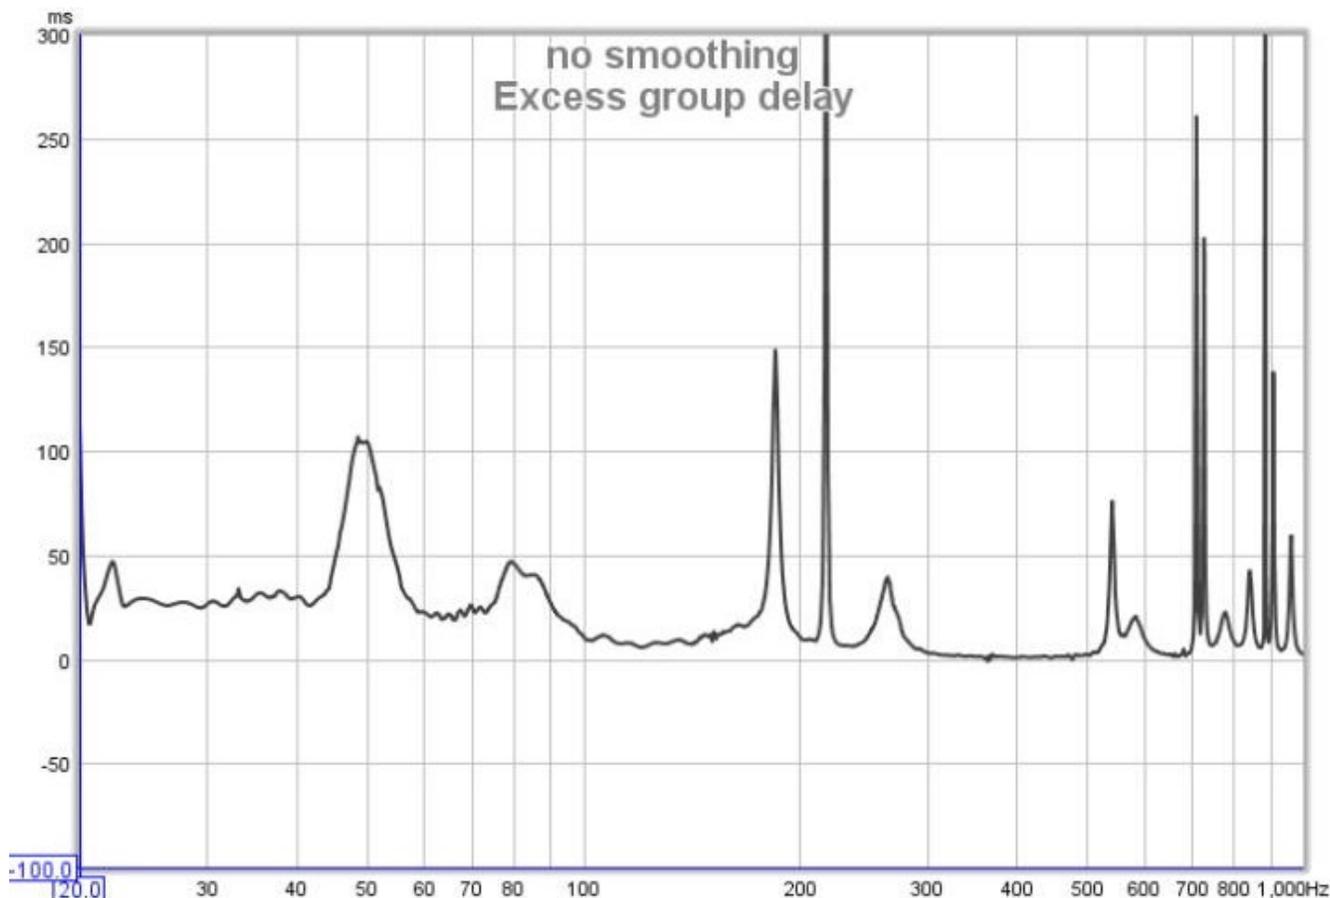
Now we have something we can work with. Anywhere the excess group delay plot is flat is a minimum phase region of the response. We can see there are regions even at very low frequencies where the response is not minimum phase, between about 44 and 56Hz for example. These will usually correspond to regions where there are sharp dips in the response, and underline the poor results which are often found when trying to lift such regions with EQ. Low frequency peaks on the other hand are usually in minimum phase regions, the plot is fairly flat in the region of the 28Hz and 60Hz peaks, which bodes well for attempts to apply EQ to them. In general, the peaks in a response are a result of features that are correctible through equalisation (speaking technically, they are due to the poles of the response and the equaliser can place zeroes that cancel the poles).

There are regions at relatively high frequencies which are minimum phase, such as 300 to 500Hz, despite the wild variations of the response in that area, so it would be possible to apply EQ there. However, we need to remember that the measurement is only valid for the microphone location at which it was made, and as frequency increases the response changes more rapidly as the microphone moves. EQ that looks good at the original measurement position may give worse results at other positions, so it is important to check wherever listeners will be. Narrow bandwidth EQ adjustments should not be used outside the modal range, the higher the frequency the broader the EQ adjustment needs to be to stand any chance of being useful outside a very small region.

As an aside, the excess group delay plot also clearly shows there is a time offset between the subwoofer and the main speaker, the sub being about 25ms delayed, which is not so obvious from the overall group delay plot. Excess group delay is a useful plot for time aligning speakers.

A common cause of non-minimum phase behaviour in rooms

If minimum phase systems are cascaded (connected in series) the overall system remains minimum phase - the individual transfer functions of the systems are multiplied together and this retains the minimum phase characteristics. In terms of the paragraph about invertibility above, the minimum phase systems will not have zero amplitude anywhere and multiplying non-zero values together will not generate a zero value. However, *adding* the responses of minimum phase systems gives a result which is typically *not* minimum phase throughout its response. If

there are any areas where the responses of the systems we are adding are equal in magnitude but opposite in phase, their sum will be zero. Here we see the problem for room responses, because the room response we measure is the summation of many different responses due to the sound radiating into the room and reflecting from its surfaces. This also applies even at the lowest frequencies, as we can see in the following.

Axial modes in a rectangular room

To provide a simple example of how the summation of the signals in a room can make it non-minimum phase, even at low frequencies, we can look at the behaviour of axial modes in a perfectly rectangular room. Such results are easily simulated (in this case by REW's own simple modal simulation tool), giving us a well controlled set of responses to study. For the responses below the room dimensions are 7.00 x 6.86 x 3.43m, giving length modes every 24.5Hz, width modes every 25Hz and height modes every 50Hz. The source is against the front wall, 0.25m from the left wall and 0.15m from the floor. The mic is 1.5m from the rear wall, 4.28m from the left wall and 1m from the floor. The room surfaces have uniform absorption of 0.20 at all frequencies.

The first plots show the individual SPL and phase responses of each axis. All are perfectly minimum phase, so the excess phase (the black line) is flat and remains at zero. A linear frequency axis is used to more easily see the modal effects, which are linearly distributed in frequency.
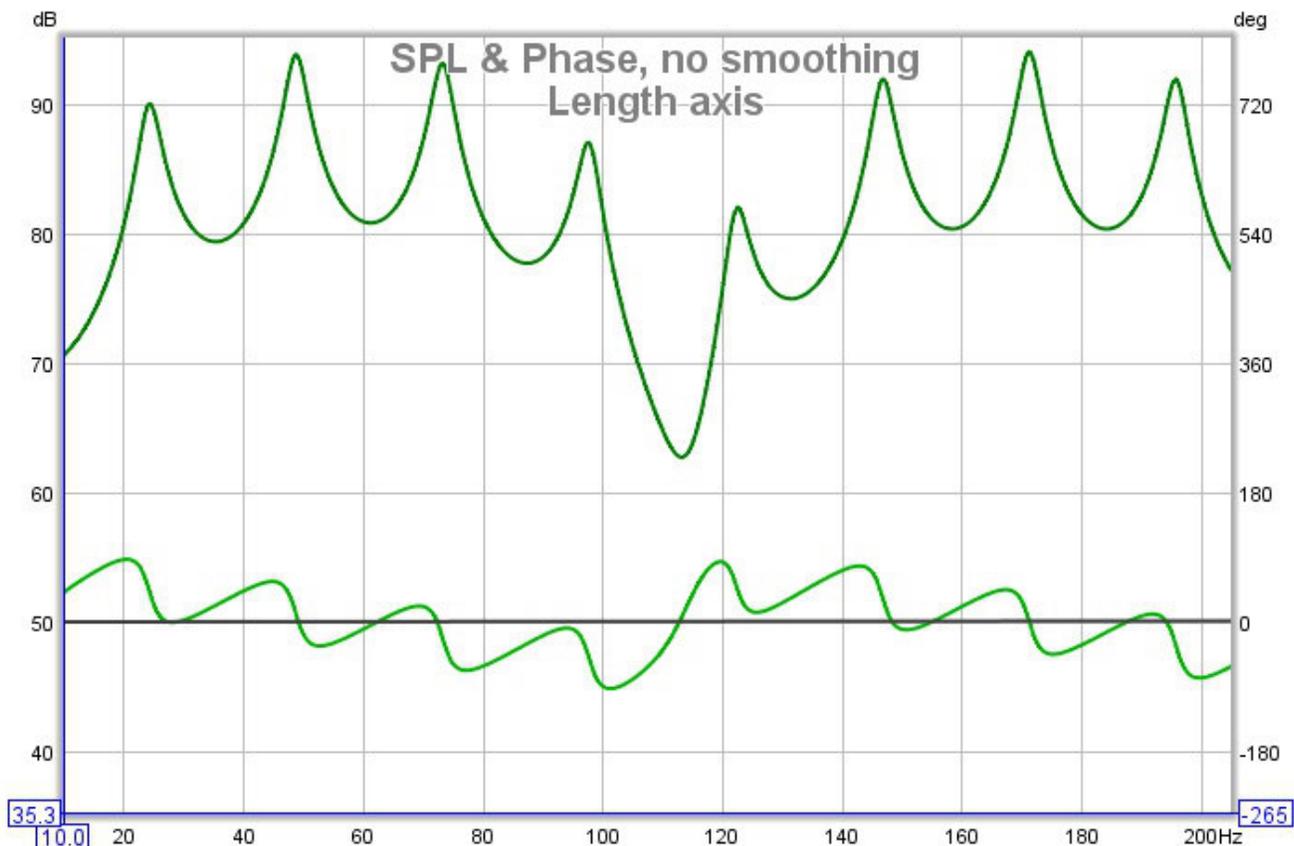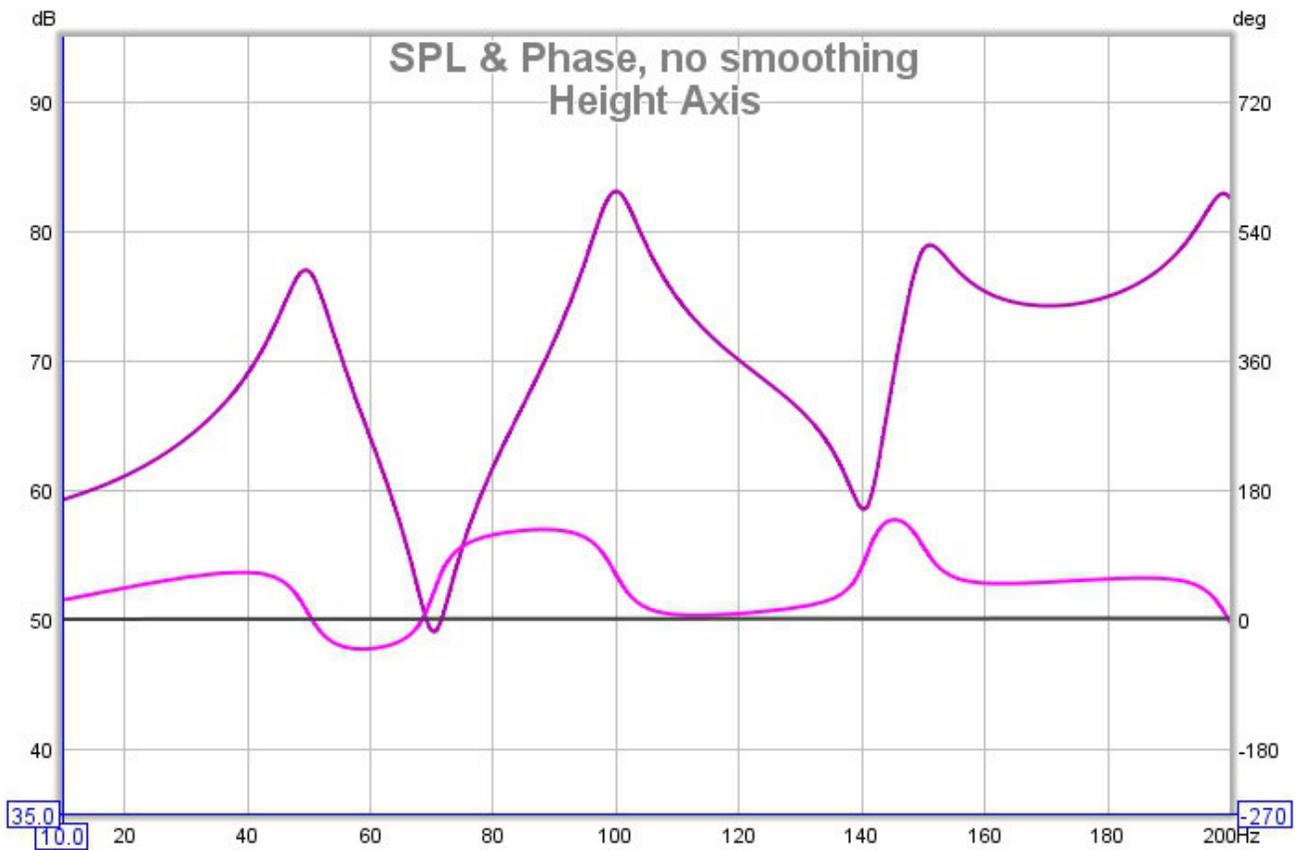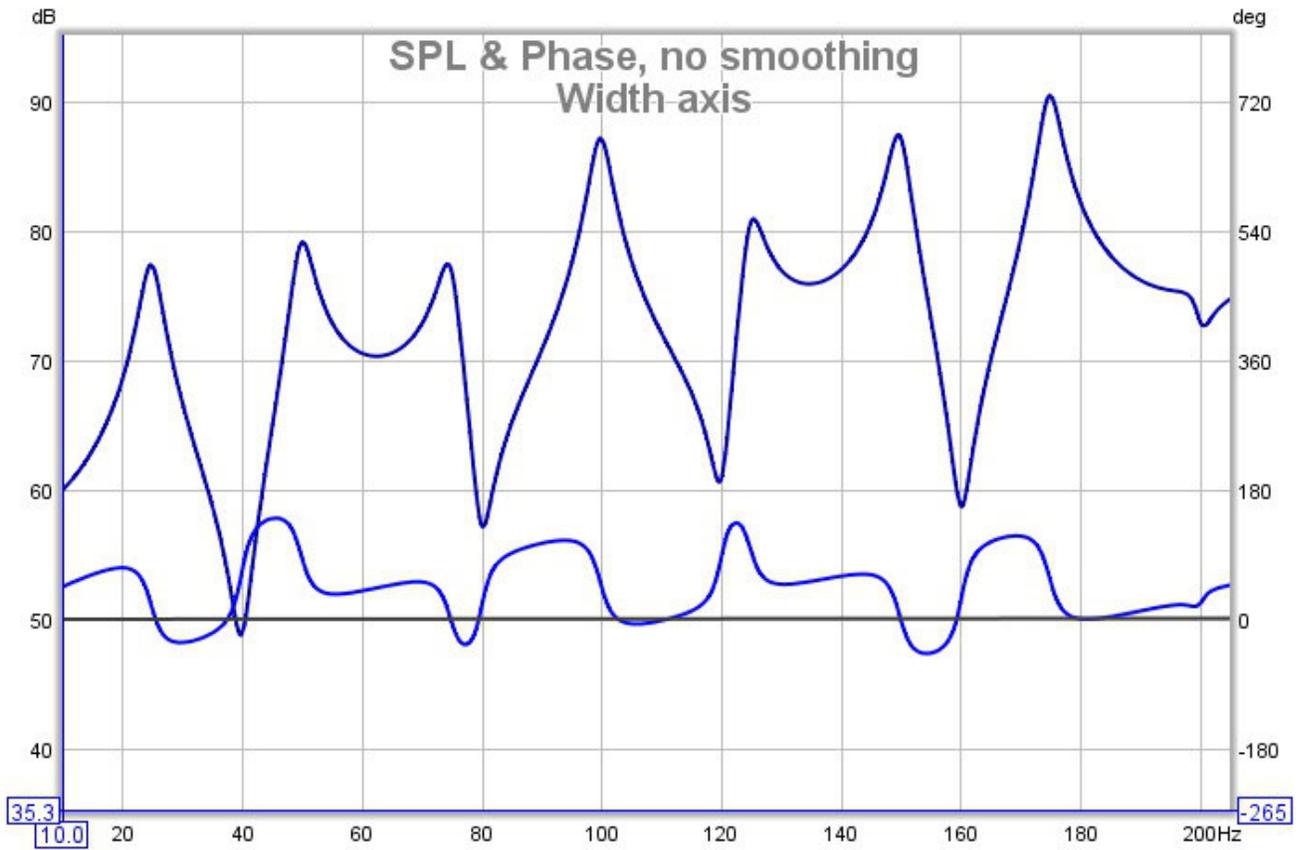
Now for the combined response, which shows the minimum phase in grey and the excess in black, followed by the excess group delay plot:
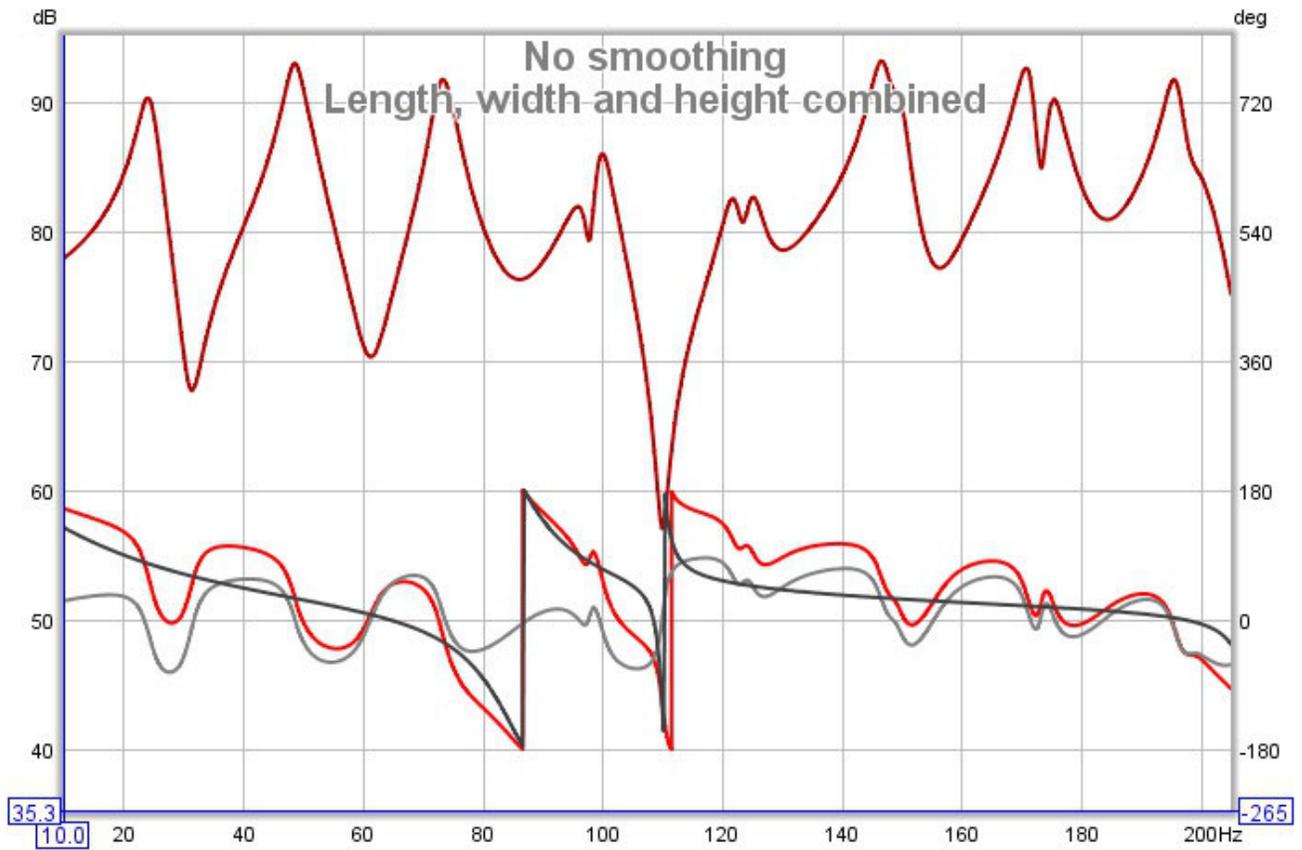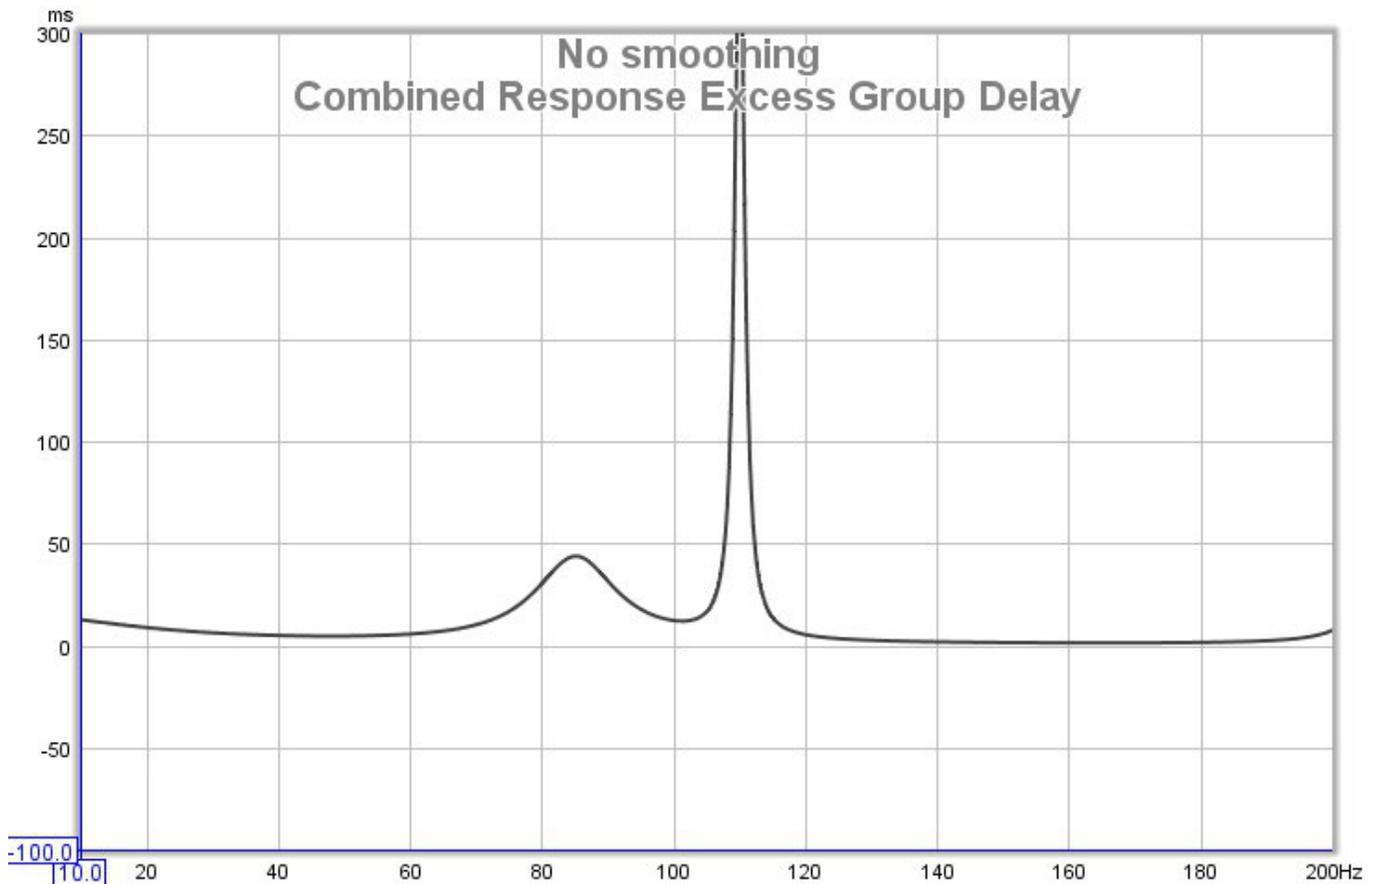
The response is no longer completely minimum phase anywhere in the span, as we can see from the excess phase, but it deviates dramatically in the 70-120Hz region. At 110Hz, where there is a sharp dip in the response, there is a sharp peak in the excess group delay. Attempting to EQ the response to flat in this region would be foolish. Regions where the response is far from minimum phase would typically not give the results we might expect and they are best left alone from an EQ perspective. Non-minimum phase regions are also likely to show greater variation with position and to be more affected by changes within the room, as a change that affects any of the signals that sum to the response in the minimum phase region can greatly alter the behaviour there. On the plus side, broadband acoustic treatments in the room are effective regardless of the room's minimum or non-minimum phase behaviour.

Note that the predicted EQ results REW shows in its EQ window are obtained by applying the chosen filters to the measured impulse response, and include the effects of non-minimum phase behaviour, so they accurately portray the results that would be obtained at the point the measurement was taken.

Help Index

Why Can't I Fix All my Acoustic Problems with EQ?

This topic deals with an important question: why isn't EQ enough to sort out acoustic problems? There are plenty of products that claim to be able to correct room responses, so why would anyone need to bother with acoustic treatments and bass traps and absorbers and all that stuff? Technology to the rescue, right?

Those are important questions, and understanding the answers to them can help a lot with better understanding acoustics in general. There are a few places where the answer gets a little technical, but for the most part the explanation is fairly easy to follow. Along the way to answering the questions above, we will touch on the answers to two other questions:

Why does phase matter?

Why should I look at time domain signals instead of just frequency responses?

What does EQ do?

As a starting point we need to take a look at what an equaliser can do for us. The basic function of an equaliser is to alter the frequency response. We can use it to try and make all the frequencies in the response equal - the clue is definitely in the name there! Particular equalisers are sometimes described as operating in the frequency domain, or operating in the time domain, or operating in both. In fact all equalisers, **without exception**, operate in the time *and* frequency domains and have effects in both.

What are the limits of applying EQ?

Location, location, location

Before we start adjusting an EQ to alter the frequency response, we need to see a response to adjust, so we need to make a measurement. This brings up the first limitation. The measurement is made at a single position, and the frequency response of that measurement is only valid at that position, moving the mic elsewhere and making another measurement will produce a different frequency response. It may be a little different, or it may be (and usually is) a lot different. The changes made by an equaliser in the path to the speaker are the same no matter where we are in the room, so since the response is changing in different positions and the EQ isn't, it stands to reason that the EQ is only going to be good in places where the frequency response is the same as the one we used when setting the EQ.

Reading some of the advertising blurb for EQ products you could be forgiven for thinking that some clever guys somewhere have figured out a way around this. They haven't. The best you can do is to look at the frequency responses measured at many positions in the area where you need the correction to work, figure out which bits of them are sufficiently common, and come up with a compromise EQ setting that helps somewhat in most places and doesn't do too much harm elsewhere. It can help, but it is no magic bullet.

What if I only listen in one spot?

So if the EQ is only good for one position, and I only sit in one position, what's the problem? The problem is very small movements make big differences. At high frequencies the wavelength of sound is very short. At 20kHz it is just 17mm, about 5/8". The frequency response varies dramatically at high frequencies over very short distances, so even if you only sit in one spot, and sit very still, the best you could hope for is an EQ setting that would work up to a few kHz. For a more reasonable range of motion a few hundred Hz is more likely.

Resolution

So we are prepared to make some compromises. One sweet spot will do, and fixing the response up to a few hundred Hz would actually help a lot, it's usually all over the place down low. So let's break out the EQ and start adjusting. The next problem we run into is the adjustments don't seem to be working right. Say the frequency response shows a 6dB dip at 100Hz. We put in 6bB of boost there and tweak the width so it matches the dip we saw. But the frequency response hardly moved, especially in the middle of the dip. What's going on? The problem is probably with the resolution of the measurement. If you have used a 1/3 octave RTA, for example, to measure the response, the bar at 100Hz actually spans the range from about 89Hz to 112Hz. That 6dB dip is probably due to a much deeper but very narrow dip within that 23Hz span. You have to make a high resolution measurement to see what is going on, an RTA isn't going to cut it for this work.

Headroom

The RTA has left the scene and we are making high resolution measurements. And they look awful. There are big peaks and some huge, narrow dips. The 6dB dip we saw at 100Hz actually turns out to be a 17dB dip at 98Hz. Never mind, the EQ allows up to 24dB gain. But listening with the fix in place reveals massive distortion. We have run right out of headroom, with clipping all over the place. Even after playing around with the levels to get rid of the clipping the result even slightly out of the sweet spot is much, much worse. Sharp dips in the response are very sensitive to position, even at very low frequencies. The sensitivity to position and the headroom problems mean we cannot do anything about them with EQ. The best we can do is deal with the broad, shallow dips and work on the peaks.

Minimum phase and all that

We know most of the limitations of the equaliser now. We moved things around a bit and used a few absorbers and got rid of the worst of the dips. After a lot of painstaking tweaking of the EQ the frequency response is actually pretty flat. But it still sounds awful. So now what is going on?

The next few paragraphs get a little more technical, but it is worth sticking with it. Equalisers are, with a few exceptions, minimum phase devices (some are linear phase, but that doesn't help with the problem facing us). When we make an adjustment to the frequency response on the EQ, we also change the phase response, an often ignored part of the measurement we made. We need to take a short diversion to look at why we should care about the phase.

Why does phase matter?

Measurement software measures the Transfer Function of the system it is hooked up to. The transfer function has two parts, the familiar frequency response, and the phase response. Systems can have the same frequency response but actually have totally different effects on signals passed through them - the difference lies in their phase responses. As a simple example of how big a difference phase can make, consider the results from measuring two very different signals: an impulse and a period of periodic noise. Both of these signals have perfectly flat frequency responses, looking at the frequency responses we could not tell them apart. The time signals obviously look completely different, so what happened to that difference when the signal went through an FFT to make the frequency response? It is all in the phase responses. The impulse has zero phase at all frequencies. The periodic noise has random phase. Just as looking at frequency response alone cannot tell us what a signal looks like, looking at the frequency response of a transfer function alone cannot tell us what the system does to signals that pass through it, we have to look at the phase response as well.

So the answer to why our system, with its nicely flattened frequency response, still doesn't sound right lies in the phase response. Room responses are, for the most part, not minimum phase. The technical explanation of that probably would not help with our understanding of the problem we are faced with, but the outcome is this: we can do almost what we like with the frequency response (within the limits we have discussed already) but the phase response is beyond the reach of our EQ. Anything we do in the EQ's frequency

response adjustments will have a corresponding effect in the phase response, and while the frequency response adjustment we make can be equal and opposite to the room's frequency response, the same is not true of the phase. That is what it means for the room not to be minimum phase, it has done things to the phase of the signal that we cannot mirror in our EQ. Fixing the frequency response but not the phase response means we cannot make the time signal look like it did before the room got hold of it, however much time we spend fiddling with the EQ. We have hit the limit.

The value of looking at time signals

That brings us to another item I said we would touch on, the value of looking at time signals and not just at frequency responses. The frequency response is only half of the description of what the system is doing to signals that pass through it, the phase response is the other half. Trying to understand systems by looking at the frequency response alone is like trying to understand a book by reading only the even numbered pages. To really understand you need to look at both. That is a bit problematic, however. The frequency response is fairly easy to understand, but the phase response doesn't give up its secrets quite so easily. To properly use it we end up looking at various quantities derived from it, such as group delay or phase delay. It gets complicated. But there is an alternative.

The systems we measure can be described in two ways: in the frequency domain by their transfer function (frequency and phase responses) or in the time domain by their impulse response. They are two views of the same system, the transfer function is the FFT of the impulse response and the impulse response is the inverse FFT of the transfer function. To study how the system behaves and what it does to signals, we can look at both. The impulse response has the benefit that it captures all the information in one signal, which puts it one up on the transfer function, though it is not as immediately intuitive as a frequency response. It readily gives up information that is less easily spotted in the transfer function though, such as early reflections or the slow decays of room modes. It is well worth taking some time to become familiar with the impulse response and some of the quantities derived from it, such as the impulse response envelope (aka ETC).

Does EQ help or not?

Given all the limitations we have uncovered, and with the problem of non-minimum phase on top, we might wonder whether the equaliser is any good to us. All is not lost, however. The non-minimum phase behaviour of the room is connected to the dips in the response. It means we are even less able to deal with them, but there wasn't a lot we could do about them anyway, so we are really not much worse off. On the plus side, the peaks of the response are caused by features that lie firmly in the region our minimum phase equaliser can handle. We can use the equaliser to help tame the peaks, and the lower down they occur, the better the results we are likely to get - a nice complement to our acoustic treatments, since they start to struggle (or we start to struggle with the size of them!) at low frequencies. EQ is a useful tool to keep handy when trying to fix our acoustical problems, but it can only ever be a small part of the solution.

Help Index

SPL Meter

The SPL Meter is an integrating, logging sound level meter that displays sound pressure level, equivalent sound level or sound exposure level based on the RMS level of the input channel. It offers A, C and Z weightings, fast or slow exponential filters, a high pass filter to suppress wind noise, and records minimum, maximum and unweighted peak levels.

The meter takes into account both the soundcard and microphone calibration files and corrects its readings accordingly, allowing IEC class 0 performance when used with a calibrated microphone and SPL calibrator. Note that the maximum boost resulting from the calibration files can be limited by a setting in the Analysis Preferences to prevent excessive boosting of the noise floor. Data recorded by the meter can be logged, graphed and saved to a text file.

The meter displays either sound pressure level (SPL), time-average equivalent sound level (L_{eq}) or sound exposure level (L_E) according to the selection made on the buttons below the display. The SPL reading is filtered with either a "Fast" (125ms) or "Slow" (1s) time constant, selected via the F/S buttons. For general use the "Slow" setting is best. When the **HP** button is pressed a high pass filter is applied that eliminates content below approximately 8Hz.

Meter Weighting

SPL measurements use weighting curves to shape the signal they receive and emphasise those regions that are of interest for certain requirements. The A and C-weighting curves are shown in the figure below.

C-weighting gives the broadest response (apart from the "Flat", "Zero" or "Z" weighting), with -3dB points at 31.5Hz and 8kHz. A-weighting has a much more pronounced low frequency roll-off. It is modelled on the sensitivity of the ear to low level sounds (about 40dB SPL). A-weighting has the same high frequency roll-off shape as C, but the curves do not align at high frequencies because they are adjusted to both have zero gain at 1kHz, which shifts the A-weighting curve up relative to C-weighting.

Meter Display

The meter display shows the currently selected measure, the level and an overload ("OVER") indicator which is lit if the soundcard input range is exceeded. The OVER indicator can be reset by clicking in the display area or using the **Reset All** button. When SPL is selected the display shows dB, the selected weighting in brackets (A, C or Z) and either F for Fast or S for Slow. When measuring equivalent sound level it shows LAeq, LCeq or LZeq according to the selected weighting. SEL is displayed when measuring sound exposure level. The time over which the equivalent sound level or sound exposure level figures have been calculated is shown in the Elapsed Time display at the bottom of the meter. Note that equivalent sound level is useful for making measurements of subwoofer levels using your processor or receiver's internal calibration signal, usually difficult due to the large fluctuations of the level. Equivalent sound level displays a result averaged over the time since **Reset All** was last pressed; simply start the test signal, press **Reset All** and wait for the reading to stabilise to get a precise level.

The SPL meter window can be resized as required, the main SPL display digits automatically scale to suit the available space. The meter can also be set to full screen.

Below the main row of controls a level meter shows the current soundcard input level in dBFS, the peak is shown by the red bar while the numeric indicator and the coloured bar show the rms level.
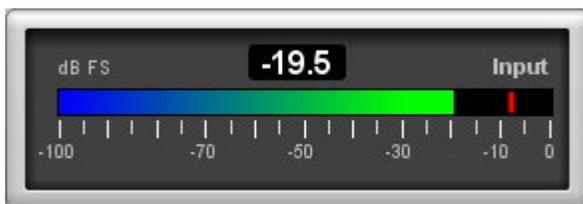
Below the soundcard level meter is a **MinMax** button to display min, max and peak values in the SPL display and a **Reset All** button to reset the elapsed time, min, max and peak values, equivalent sound level, sound exposure level and the overload indicator. The **Calibrate** button starts the meter calibration process, whilst the record button turns the meter on or off. When MinMax is selected the values are displayed alongside the main reading.

Meter Input Selection

The soundcard input channel which is measured is selected in the Soundcard Preferences. Setting the audio input is described in Getting Started.

Calibrating the SPL Reading

Valid readings are not displayed until SPL Calibration has been carried out. The SPL reading may be calibrated against an external SPL meter or SPL calibrator by pressing the **Calibrate** button. The text on the button and the SPL meter digits are red if the meter has not been calibrated. If the soundcard input is altered in the REW soundcard preferences, or if the input levels are altered, it will be necessary to re-calibrate the SPL meter reading.

SPL Data Logging

The **Logger** button opens the SPL Logger graph window. The record button in the top right corner of the SPL Logger graph starts or stops logging of SPL values. When logging is in progress the SPL meter on/off, calibrate, weighting, filter time constant and high pass filter buttons are disabled. The logger records the SPL (with the currently selected weighting and time constant), the minimum and maximum values, the unweighted and uncalibrated peak value, the equivalent sound level and the sound exposure level. A **Save** button above the SPL Logger graph allows the data recorded to be saved to a text file using the text delimiter set in the REW File menu, log files can be loaded using the **Open** button.

SPL Logger Controls

The control panel for the SPL logger graph has these controls:
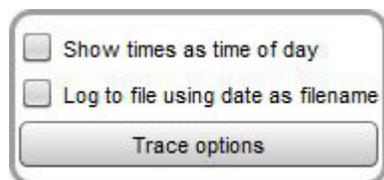
If **Show times as time of day** is selected the time axis shows the time of day of the logged data, otherwise it shows the elapsed time since logging started. Selecting **Log to file using date as filename** saves logged data to files in the REW logs directory automatically whenever the logger is running (the directory location is shown in the **About REW** dialog box). Filenames are formatted as SPL-YYYY-MMM-dd.txt, e.g. SPL-2016-Apr-24.txt. If a file for a day already exists a number will be appended, e.g. SPL-2016-Apr-24-1.txt, SPL-2016-Apr-24-2.txt etc. A new file is created if the logging continues past midnight, with the new date. Note that the files can become very large if logging continues for a long period, growing at about 18 kB per minute. If the option to log to files is selected remember to check the REW log directory to remove files you no longer need.

The **Trace options** button brings up a dialog that allows the colour and line type of the graph traces to be changed. If a change is made it will be used for all measurements shown on this graph. Traces can also be hidden, which will remove them from the graph and from the graph legend.

The format of the logged data is shown below, in this example the text delimiter is comma and the SPL meter was set to Z weighting.

```
SPL log data saved by REW V5.19
Note: includes Leq over last minute and last 10 minutes
Soundcard cal:
Mic/meter cal:
SPL cal offset: 104.85347938537598 dB
High Pass Filter: Off
Date: 16-Jul-2018 17:43:10
Start: 1531762990157
Log length: 7061 entries over 1311.6s
Log interval: 0.18575963718820862 seconds

Time[s], LZS, LZSMin, LZSMax, LZpeak, LZeq, LZE, LZeq1m, LZeq10m
```

```
0.186, 84.7, 84.7, 84.7, 95.1, 84.6, 77.3, 84.6, 84.6, 17:43:10.186
0.372, 84.7, 84.6, 84.7, 96.1, 84.7, 80.4, 84.7, 84.7, 17:43:10.372
0.557, 84.6, 84.6, 84.7, 96.1, 84.5, 82.0, 84.5, 84.5, 17:43:10.557
0.743, 84.7, 84.5, 84.7, 96.1, 84.6, 83.3, 84.6, 84.6, 17:43:10.743
0.929, 84.7, 84.5, 84.7, 96.1, 84.6, 84.3, 84.6, 84.6, 17:43:10.929
```

[Help Index](#)

Signal Generator
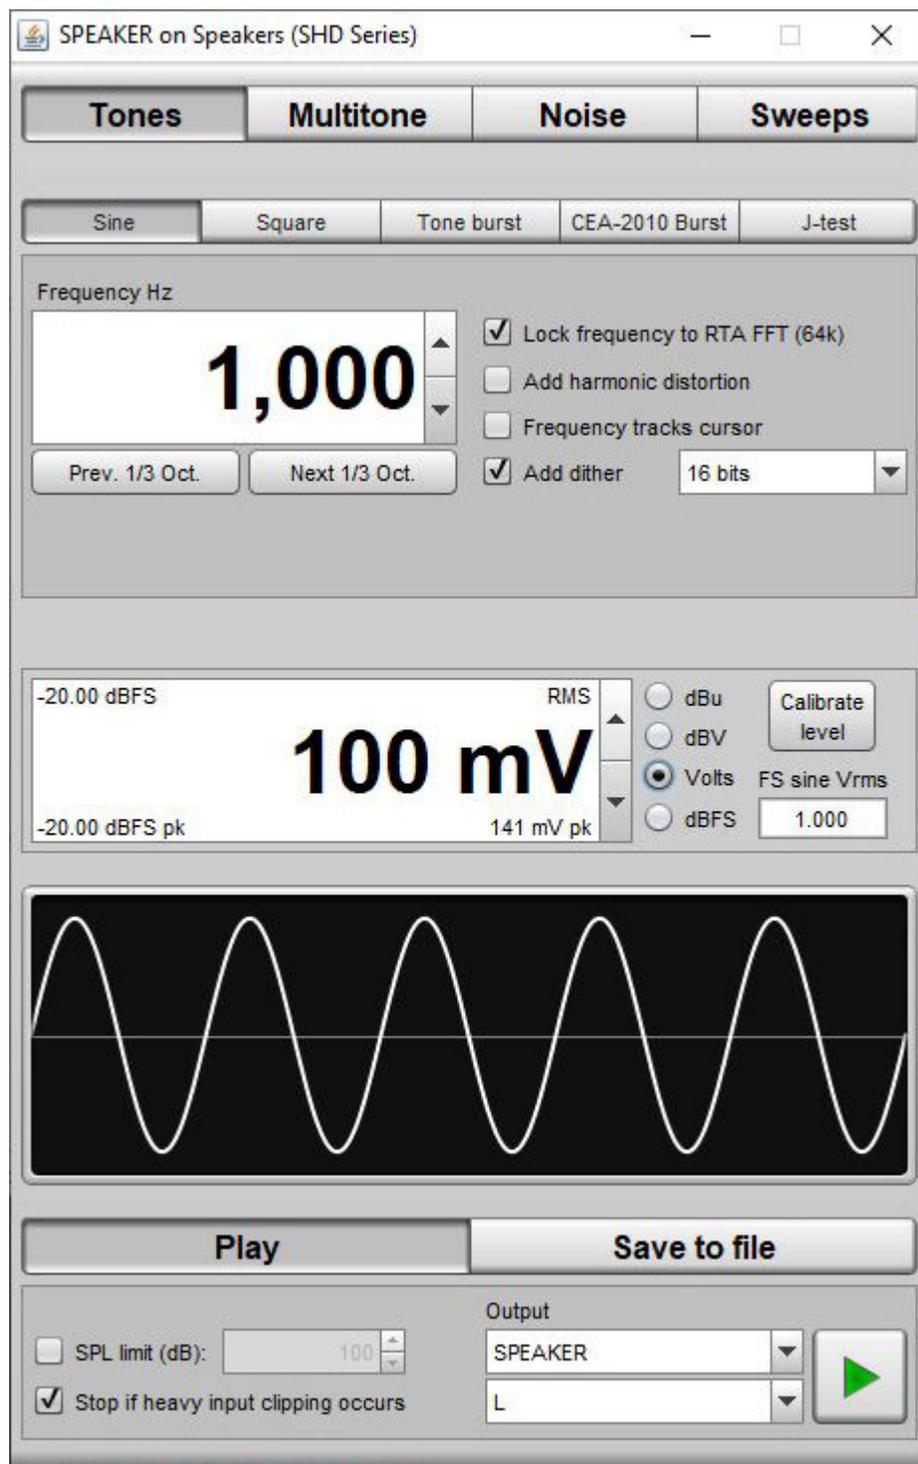
The Signal Generator divides its signals into four groups, Tones, Multitone, Noise and Sweeps. Within those groups it offers the following signal types:

- [Sine Waves](#)

- Optionally including harmonic distortion

- [Square Waves](#)
    - Variable duty cycle

- [Tone Burst](#)
    - Windowed tone burst of 0.5 to 100 cycles
    - Looping to repeat the burst

- [CEA-2010 Burst](#)
    - 6.5 cycle Hann-windowed tone burst
    - Looping to repeat the burst

- [J-test](#)
    - Jitter test signal, main component at one quarter sample rate

- [Dual Tone](#)
    - SMPTE, DIN, CCIF and Custom

- [Triple Tone](#)
    - Borberly, Cordell, Klingelnberg and Custom

- [Multitone](#)
    - Linear, octave, decade and 'no interharmonic distortion' spacings

- [Pink and White Random Noise](#)
    - Full range (spectrum to below 10 Hz)
    - Custom filtered

      and for pink only:

    - Speaker Calibration
    - Subwoofer Calibration
    - CTA-2034 filtered

- [Pink and White Periodic Noise](#)
    - Length to suit RTA FFT length
    - Full range (spectrum to below 10 Hz)
    - Custom filtered

      and for pink only:

    - Speaker Calibration
    - Subwoofer Calibration
    - CTA-2034 filtered

- [Sine Sweeps](#)
    - Linear
    - Logarithmic (synchronised, harmonic have same phase as fundamental)
    - Looping to repeat the sweep

- [Measurement Sweep](#)
  - REW's logarithmic measurement sweep

RMS Signal Level

The RMS signal level can be set for any of the signal types with 0.01 dB resolution relative to digital full scale. Display units of dBFS, dBu, dBV and volts are offered. The arrow buttons on the RMS Level spinner change the value in steps of 0.1dB, or any required value can be typed directly into the level box. The value will be converted as required, for example typing 775mV into the level control when in dBu mode will convert to 0 dBu (or -2.2 dBV). In modes other than dBFS the dBFS value is shown at the top left of the level control.

The [View setting for full scale sine](#) determines whether a sine wave whose peaks reach digital full scale is assigned an rms level of 0 dBFS (the AES definition) or -3.01 dBFS (mathematically more correct).

The output voltage levels are calibrated using the **Calibrate level** button. To calibrate the level a sine wave signal should be played and the RMS voltage that is produced measured using a multimeter or oscilloscope and entered in the calibration dialog. Multimeters tend to be most accurate with signals that are around mains frequency. When the voltage reading has been entered the corresponding full scale voltages (rms and peak) are displayed. If the output scaling of the device being used is known (taking into account any volume control settings) it can be entered directly into the text field below the **FS sine Vrms** label.
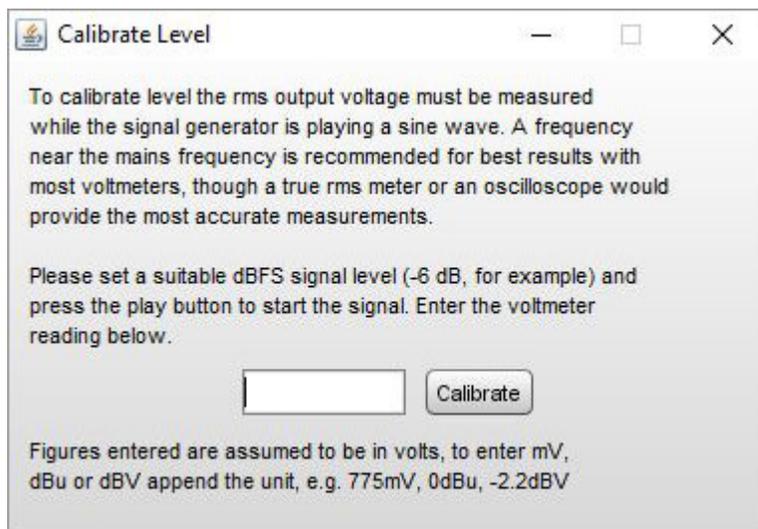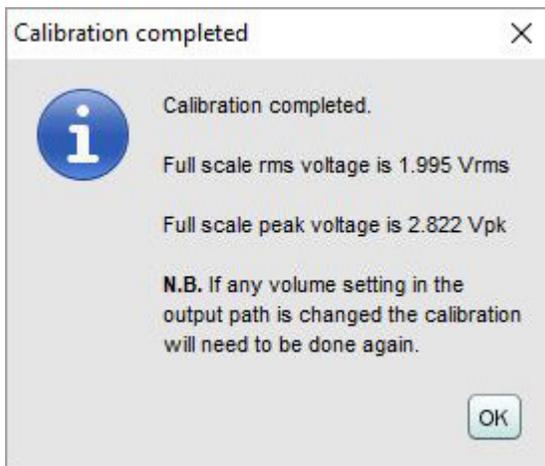
If the output scaling of the device being used is known (taking into account any volume control settings) it can be entered directly into the text field below the **FS sine Vrms** label. Clicking in the triangle in the upper left corner brings up a list of preset values that may be entered, the labels and values can be changed as required.
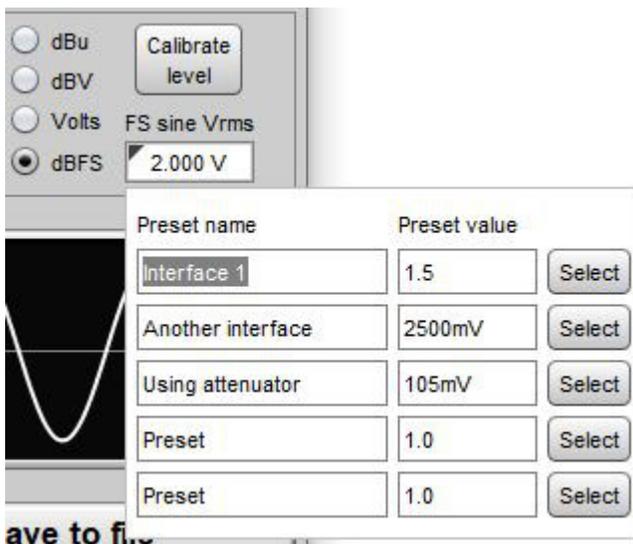
Waveform preview

A preview of the signal waveform is shown below the level control. The preview auto-scales to the range of the output, if the output is close to digital full scale the full scale limits will be shown as red lines at the top and bottom of the display.
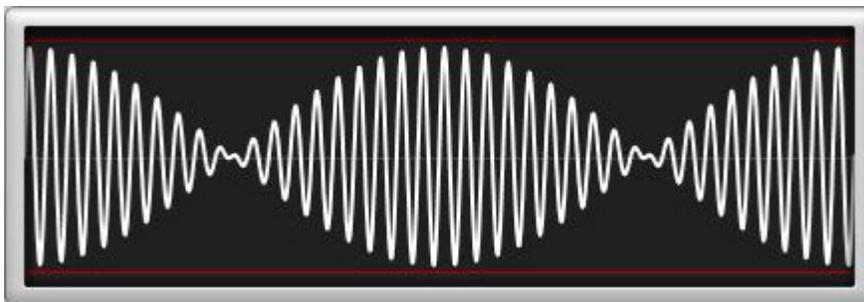
Signals which would exceed digital full scale cause the waveform display to turn red and display a warning.

Output

Setting the audio output device is described in <u>Getting Started</u>. The output on the chosen device can be selected in the signal generator within the **Play** panel. When using Java drivers the signal can also be generated on either or both channels of the output if it is stereo. If a loopback connection is being used as a timing reference (via the timing reference selection on the Measure dialog) the generator output will be sent to both the selected output and the timing reference output.
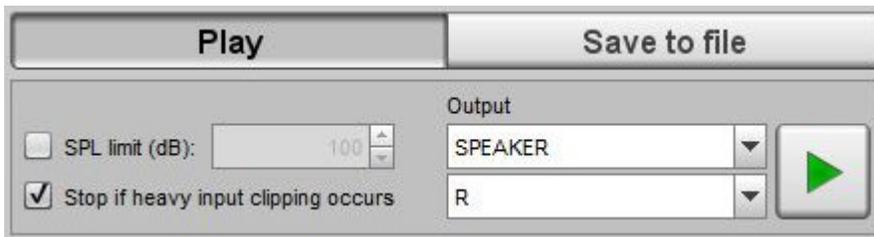
Protections

There are two protection mechanisms for the signal generator, aimed at avoiding excessive levels being generated. If **SPL Limit** is selected the generator will be stopped if the input SPL exceeds the figure set in the spinner. If **Stop if heavy input clipping occurs** is selected the generator will stop if more than 30% of the samples in an input signal block are above the clipping threshold. **Note that** the REW input must be running for these mechanisms to function, which means the SPL meter, level meters or RTA must be active.

Saving signals to WAV files

The signals can be saved to a stereo wave file containing the selected signal in either or both channels. The options offered depend on the signal type - file duration for a sine wave, for example, or number of repetitions for a tone burst. The signal level in the file is per the RMS Level setting of the signal generator. The sample rate and bit depth are as selected in the panel. If saving a measurement sweep there is an option to include a timing reference, which is recommended. Measurement sweep files include 2 to 3 seconds of 16-bit dither before the test signals start to allow playback devices to lock to the source. Note that the REW measurement sweep may change between REW versions so always save sweeps from the REW version being used for analysis of the captured response.
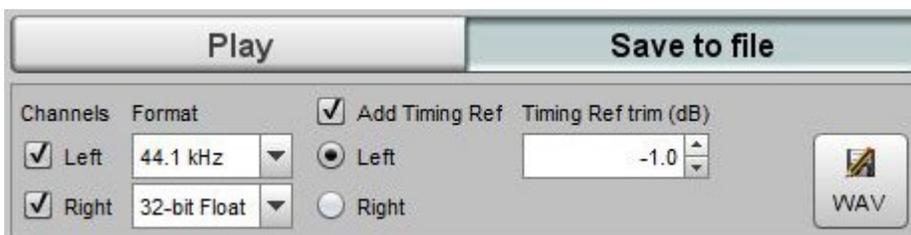
Sine Wave
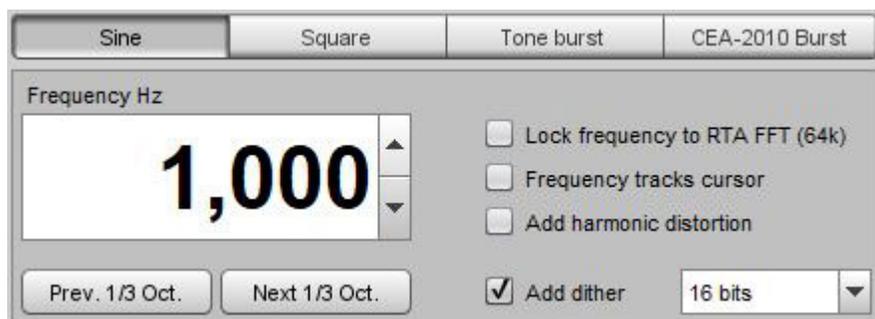
Sine waves can be generated with frequencies between 1.0 Hz and half the soundcard sample rate, e.g. 24 kHz for a soundcard operating at 48 kHz. The frequency is controlled by entering a value in the Frequency box, or using the arrow buttons to increment or decrement the value in steps of 0.5 Hz for frequencies below 200 Hz and steps of 1 Hz thereafter. If the option to lock the frequency to the RTA FFT has been selected the actual frequency that has been generated is shown at the top right corner of the frequency display when the generator is running as it will typically differ slightly from the entered frequency. The **Prev. 1/3 Oct.** and **Next 1/3 Oct.** buttons move the generator to a one-third octave centre frequency.

Lock frequency to RTA FFT

When the **Lock frequency to RTA FFT** box is checked the generator frequency is adjusted to the nearest FFT bin centre for the current RTA FFT length (meaning the signal is periodic within the FFT length). This allows a rectangular FFT window to be used for maximum spectral resolution of the RTA plot. The FFT length is shown in brackets. The exact frequency that has been generated is shown at the top right corner of the frequency display when the generator is running.

Frequency Tracks Cursor

The frequency can also be controlled via the graph cursor by checking the **Frequency tracks cursor** box. When this box is checked the generator frequency is linked to the position of the graph cursor and will change to follow the cursor frequency as it is moved - the changes are smooth with no phase discontinuities.

Add harmonic distortion

When the "add harmonic distortion" box is checked a window is displayed allowing control over the levels and phases of harmonics from the 2nd to the 9th. Each harmonic can be enabled or disabled using the check box. The levels are adjusted in dB relative to the level of the fundamental, with the equivalent percentage value shown in a label alongside. Above the controls is a display of the maximum signal level that can be used with the current distortion settings without clipping the output. Note that only harmonics that fall within the range supported by the current sample rate are generated. The signal preview includes the effects of any distortion selected.
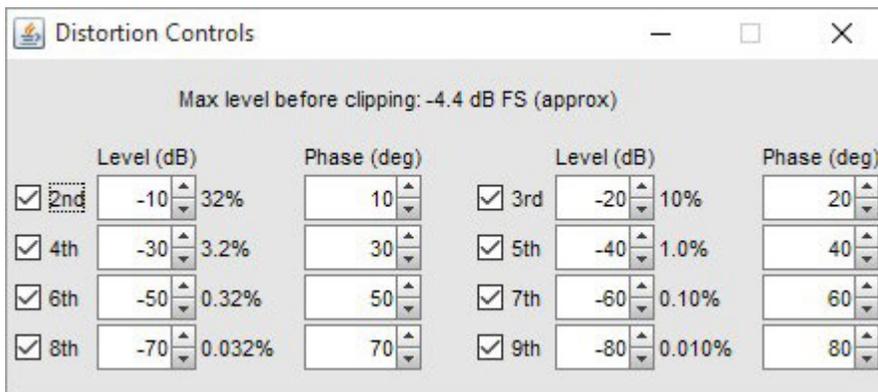
The harmonics dialog can also be used to synthesise signals that use harmonics of a fundamental, such as approximations to a square wave (odd harmonics) or the polarity test signal below, which is generated using just the fundamental and a 0 dB 2nd harmonic at -90 degrees phase.

Add dither

When the **Add dither** box is selected the generator adds 2 lsb pk-pk triangular dither to the output to remove quantisation noise spikes. The level at which the dither is added is controlled by the sample width selector to the right of the check box. **N.B. When using the Java drivers audio data is limited to 16 bit precision on Windows.** Dither is beneficial if making very precise distortion measurements of an electronic device such as a receiver, processor or equaliser. It is usually not required when making acoustic measurements as the quantisation artefacts it removes are far below the acoustic noise floor. The Graphs below show the effect of the dither option during a loopback test of a soundcard playing a 1 kHz tone at -6 dBFS. The first plot is without dither, the second plot is with dither. Addition of dither cleans up much of the noise that was apparent below -120 dBFS, especially at high frequencies, making the true harmonic distortion levels more visible.
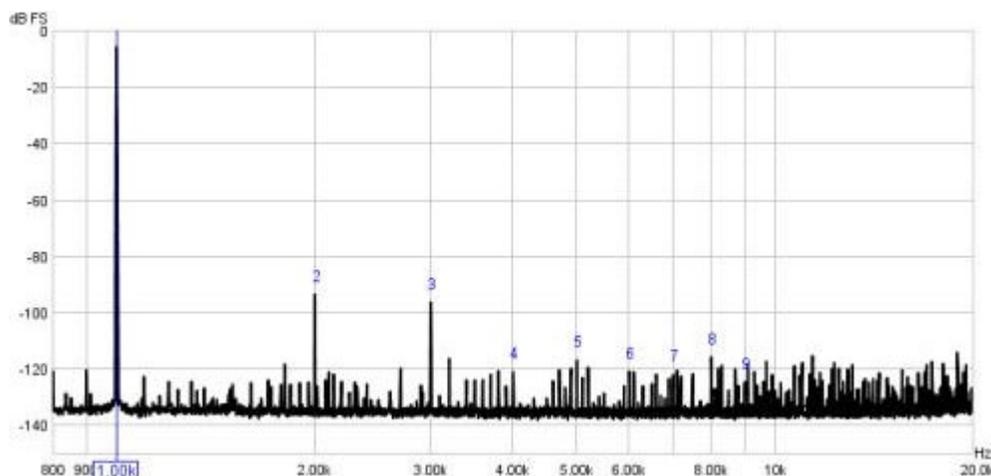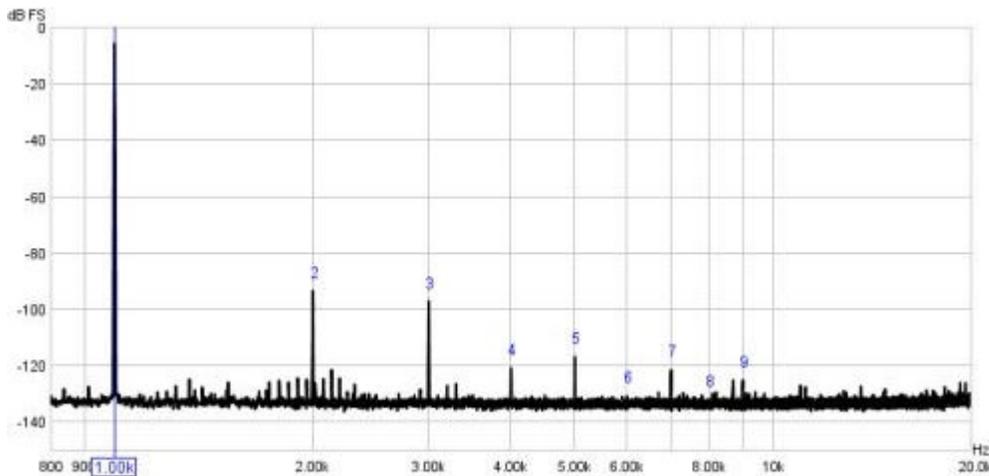
Square Wave

The square wave generator allows duty cycles between 1% and 99% in 1% steps. The frequency that is generated is adjusted to ensure there is an even number of samples in the period, so that the spectrum of a 50% duty cycle square wave will only have odd harmonics. The actual frequency that has been generated is shown at the top right corner of the frequency display when the generator is running, at higher frequencies this can be significantly different to the frequency that was entered.

Tone Burst

The Tone Burst generator produces a windowed tone burst of the selected number of cycles at the selected frequency.

If the **Repeat the burst** checkbox is selected the burst will be repeated at the selected period. The period may be specified in cycles, ms or as a number of samples. Using a number of samples equal to the RTA FFT length allows a stable display of the signal spectrum if a Rectangular RTA window is used. The **Previous 1/3rd octave frequency** and **Next 1/3rd octave frequency** move the generator to a 1/3rd octave centre frequency.

CEA-2010 Burst
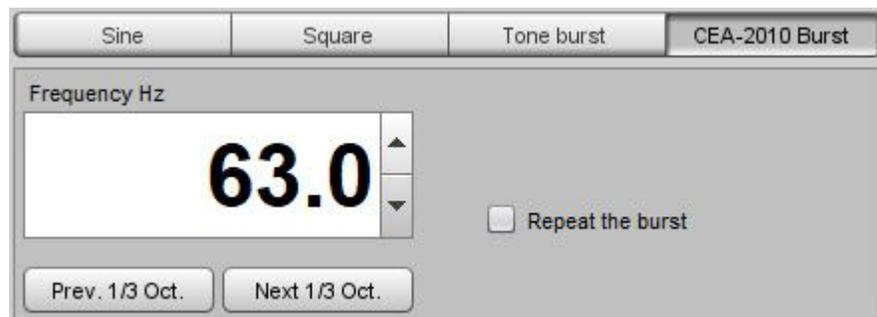
The CEA-2010 Burst generator produces a 6.5 cycle, Hann-windowed tone burst at the selected frequency. This signal is used for testing the maximum power output of subwoofers by using an RTA in Spectrum mode to observe the levels of the distortion components produced when the signal is playing, usually testing at 63, 50, 40, 31.5, 25 and 20 Hz. The signal is also used when testing maximum SPL of loudspeakers under CTA-2034-A, at 1/3 octave centre frequencies over the range from 20 Hz to 5 kHz. The CEA-2010 limits for the distortion components are shown in the table below, where f_0 is the test signal frequency.

Start Freq (Hz)	End Freq (Hz)	Limit (dB)	Comment
16	$1.59 * f_0$	0	Fundamental
$1.59 * f_0$	$2.52 * f_0$	-10 (32%)	2nd harmonic
$2.52 * f_0$	$3.78 * f_5$	-15 (18%)	3rd harmonic
$3.78 * f_0$	$5.61 * f_0$	-20 (10%)	4th and 5th harmonic
$5.61 * f_0$	$8.50 * f_0$	-30 (3.2%)	6th - 8th harmonic
$8.50 * f_0$	10 k	-40 (1%)	Higher order harmonics

If the **Repeat the burst** checkbox is selected the burst will be repeated at intervals of not less than 1 second (the actual interval is chosen to align with the RTA block length). The **Prev. 1/3 Oct.** and **Next 1/3 Oct.** buttons move the generator to a one-third octave centre frequency.

The highest level of the fundamental for which none of the harmonic limits are exceeded is the maximum output level at that test frequency. The reference level for the limits is the maximum level within 3 Hz of the test frequency for CEA-2010 or within the 1/3 octave span of the test frequency for CTA-2034-A. When the CEA burst signal is playing the RTA shows the limits and the peak level at the test frequency as an overlay, provided the RTA is in Spectrum mode. The peak level is shown in red if the limits are exceeded.
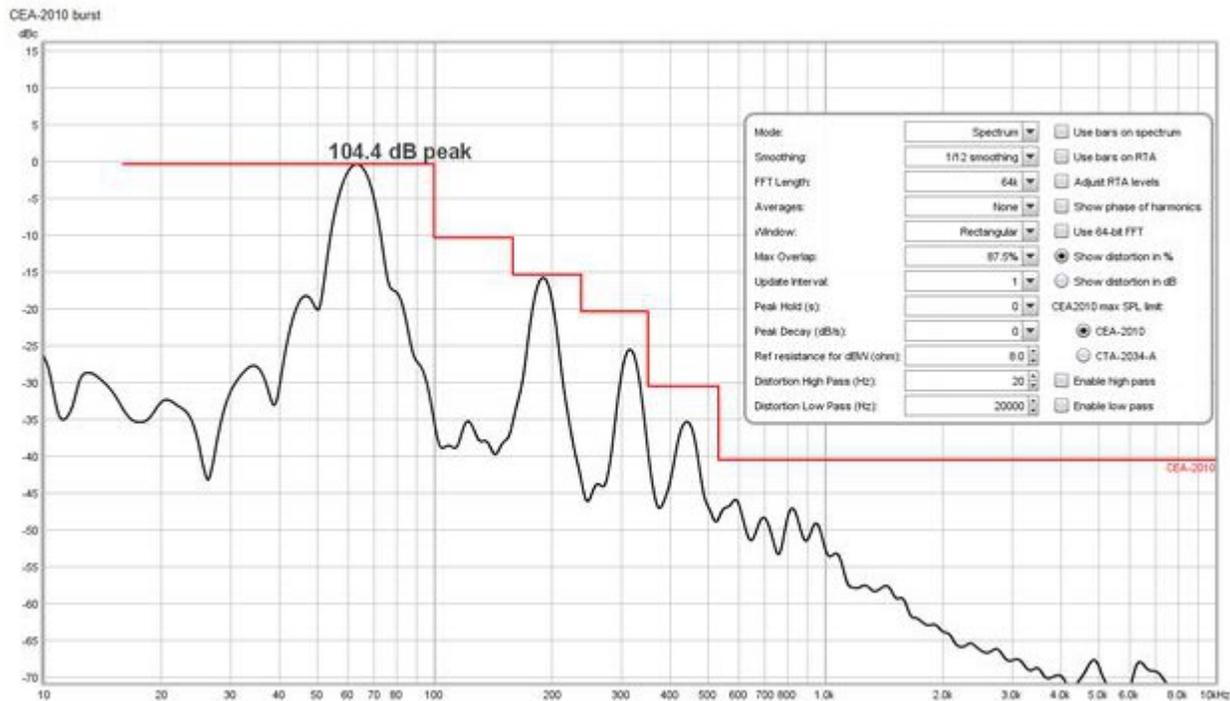
If the RTA data from a CEA-2010 signal is saved as a measurement the peak level and limits overlay will be shown on the SPL & Phase graph and the CEA-2010 test frequency and peak level figure will be recorded in the measurement notes.

The recommended RTA settings for 44.1 kHz or 48 kHz sample rate are:

For 88.2 kHz or 96 kHz use FFT length of 128k (131,072). Refer to the CEA-2010 or CTA-2034-A standards for details of the measurement procedure, or search for guides on the Internet.

J-test
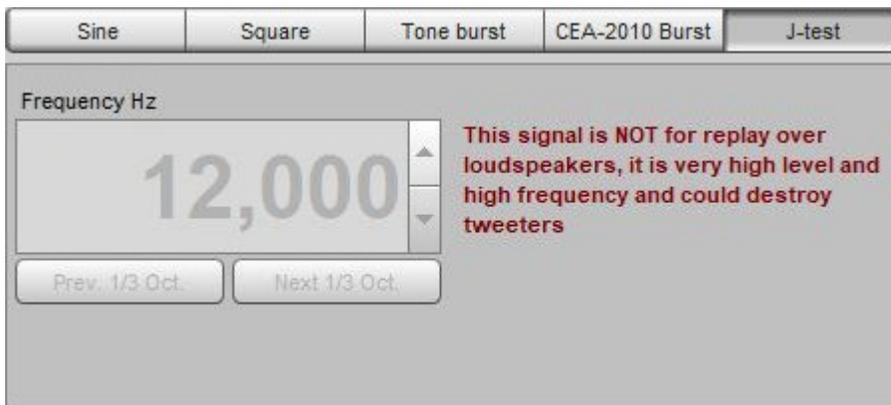
The jitter test signal consists of a square wave at one quarter of the sample rate at half full scale amplitude modulated by a square wave at 1/192 of the sample rate that alternates between 0 and -1 lsb at the 24-bit level. **This signal is NOT for replay over loudspeakers, it is very high level and high frequency and could destroy tweeters**. Note that Java drivers on windows only support 16-bit data and the OS adds dither on the input, so analysing the signal on Windows requires ASIO drivers.

Dual Tone
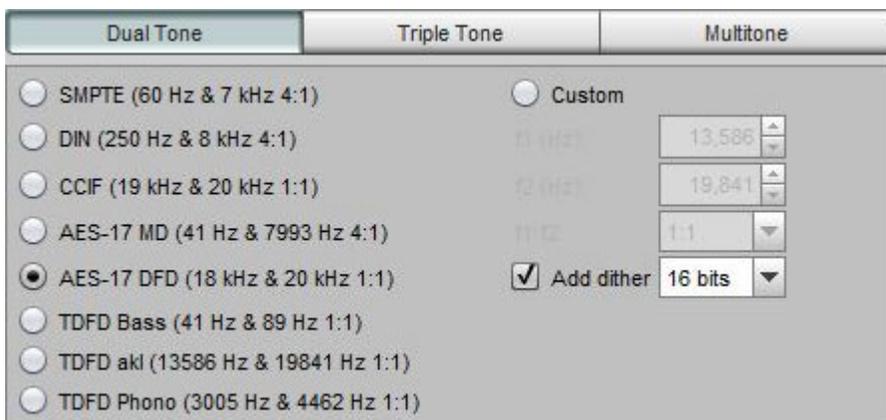
The dual tone generator is to facilitate intermodulation distortion measurements. It has presets for AES17-2015, SMPTE, DIN, CCIF and TDFD signals and allows custom signals to be generated with ratios of 1:1 or 4:1. Note that for valid IMD results with custom signals f2 must be higher than f1. The rms level of the combined signal is the signal generator rms level, but dual tone signals have a higher crest factor than a single tone so the signal will clip at a level 3 dB lower than a sine wave for 1:1 signals (crest factor 6 dB) or 1.7 dB lower for a 4:1 signal (crest factor 4.7 dB).

When the dual tone signal is being used the RTA can show figures for the intermodulation distortion and the levels of the various intermodulation components.

Triple Tone
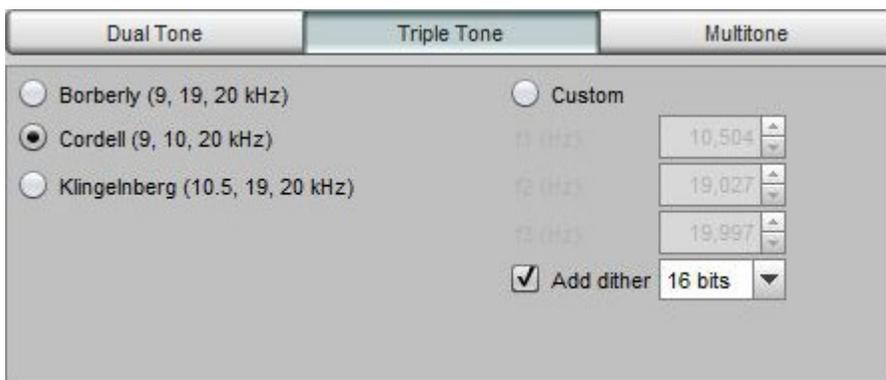
The triple tone generator provides an additional stimulus for intermodulation distortion measurements. It has presets for versions according to Borberly, Cordell and Klingelnberg and allows custom signals to be generated. All tones have the same level. The rms level of the combined signal is the signal generator rms level, but triple tone signals have a higher crest factor than a single tone so the signal will clip at a level 4.8 dB lower than a sine wave (crest factor 7.8 dB). For a discussion of these signals see Klingelnberg, Arndt: Non-linear distortion revisited, 29th Tonmeistertagung, VDT International Convention, November 2016.

When the triple tone signal is being used the RTA can show figures for the Total Distortion + Noise (TD+N) percentage.

Multitone
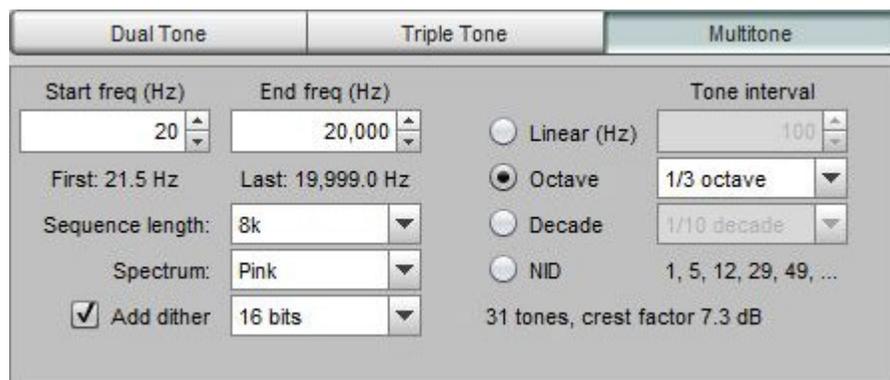
The multitone generator produces multiple tones over a defined frequency span. The tones can be spaced linearly, logarithmically at a selected fractional octave or fractional decade interval or in a sequence that places the tones such that they do not correspond to the low order harmonic or intermodulation products of other tones ('No Interharmonic Distortion' or NID). Octave fraction spacing uses frequencies from the preferred list. In all cases the tones are placed at the bin centres of an FFT of the selected sequence length, so that the behaviour of a system fed by the tones can be observed on an FFT with at least that length using a rectangular window.

The phases of the tones in the multitone sequences are adjusted to minimise the crest factor of the signal. For linearly spaced tones that will typically produce sequences with crest factors below 5 dB. Log spaced and NID tones may have crest factors of 12 dB or more. The crest factor is shown on the panel. Note that this is the crest factor of the signal as generated, the crest factor may increase during D to A conversion. The maximum RMS level before clipping is 3 minus the crest factor, e.g. -6 dBFS for a 9 dB crest factor (with the View option **Full scale sine rms is 0 dBFS** selected, 3 dB lower otherwise) but a lower setting may be required to avoid clipping on the receiving device.

The multitone sequence can be configured to have a white (equal amplitude) or pink (amplitude falling at 10 dB/decade) spectrum. Pink spectrum sequences will typically have much lower crest factors and are recommended for testing speakers. White spectrum is the norm for testing electronic devices, it has much

more energy at high frequencies than pink noise and is not recommended for use with speakers, it may damage tweeters at high levels.

When the multitone signal is being used the RTA can show a figure for the Total Distortion + Noise (TD+N) percentage and, if the FFT is two or more times the signal length, a signal to noise ratio (SNR) figure.

Pink & White Random Noise

The Pink Noise generator uses white noise filtered through a -10 dB/decade filter generated from a weighted sum of a series of first order filters, as devised by Paul Kellet circa 1999. Stated accuracy is within 0.05 dB above 9.2 Hz at 44.1 kHz sample rate.

The Full Range option outputs the filtered noise directly, giving the widest bandwidth and the greatest low frequency content. The Speaker Cal option applies 2nd order (40 dB/decade, 12 dB/octave) filters at 500 Hz and 2 kHz, producing a signal with its energy centred on 1 kHz. Sub Cal applies filters at 30 Hz and 80 Hz. Both are broadly in line with the THX test signal recommendations. CTA-2034 filtering applies shaping in accordance with ANSI/CEA-2034-A Standard Method of Measurement for In-Home Loudspeakers. The octave and 1/3 octave options filter the signal at the selected centre frequency with 1 octave or 1/3 octave total bandwidth. Custom Filtered allows low and/or high cut filter frequencies to be set arbitrarily, subject to a minimum bandwidth of 1/3 octave. The filters are Butterworth high pass and low pass with a choice of filter order from 2nd (12 dB/octave) to 8th (48 dB/octave).

REW automatically adjusts the signal levels for the various options and filter settings so that the signal RMS value reflects the setting in RMS Level. Note that as Pink Noise has random variations some clipping of peaks will occur at RMS levels above approximately -12 dBFS (or -9 dBFS if 0 dBFS is configured in the View preferences to be the level of a full scale sine).

White noise has much more energy at high frequencies than pink noise and is not recommended for use with speakers, it may damage tweeters at high levels.

Pink and White Periodic Noise

Periodic Noise (PN) sequences are ideally suited for use with spectrum and real time analysers (RTA's). They contain every frequency the analyser can resolve in a sequence length that matches the length of the analyser's FFT. Their great benefit is that they produce the desired spectrum shape without requiring any averaging or windowing, so the analyser display reacts much more rapidly to changes in the system than it would if testing with pink or white random noise, making them ideal for live adjustment of EQ filters. The PN sequences REW generates are optimised to have a crest factor (ratio of peak level to rms level) that does not exceed 6 dB for full range sequences, narrower sequences should have crest factors less than 6.5 dB. The CTA-2034 filtered noise has a crest factor of 12 dB, per the specification. Use Pink PN when measuring with an RTA or White PN with a Spectrum analyser.

The Sequence Length control must be set the same as the length of the FFT used by the analyser. If it is shorter than the analyser FFT there will be notches in the analyser display, as the periodic noise will not contain some of the frequencies the analyser is looking for. If it is set longer the extra frequencies will give a noisy display requiring more averaging. When using the REW RTA the sequence length is automatically set the same as the FFT length. Here are example RTA settings for a 64k FFT length:

The images below show the effect of correct and incorrect settings of the PN length for a loopback measurement with 1/48 octave RTA that is using an FFT length of 65536 (64k).
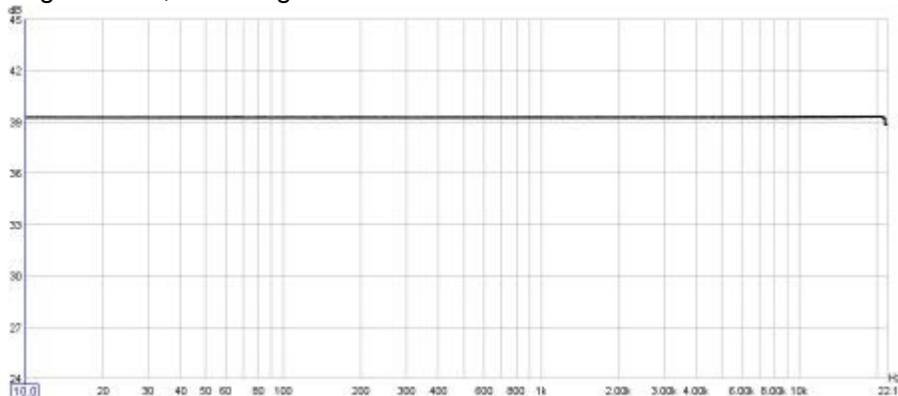
Length 32768, shorter than FFT
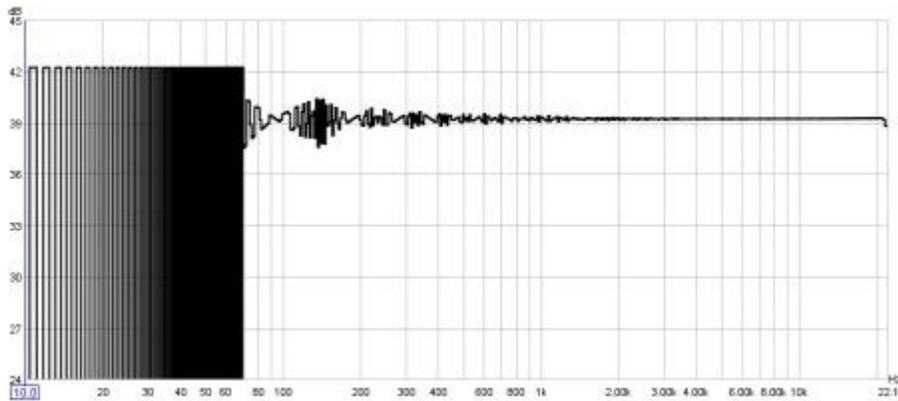
Length 131072, longer than FFT (no averaging)
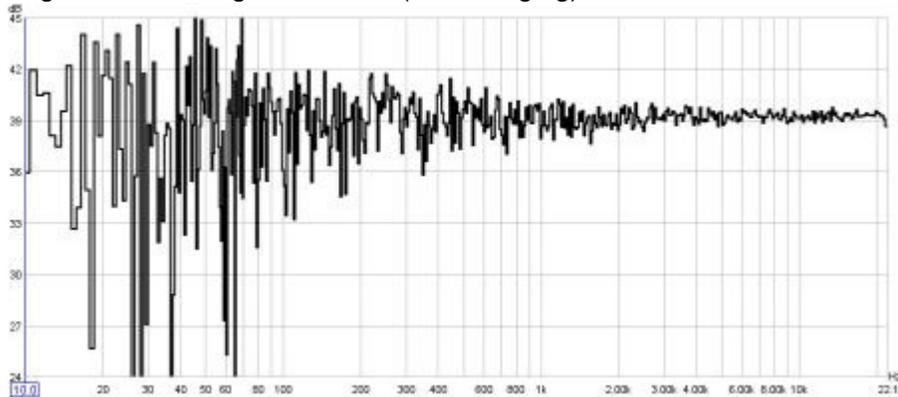
Length 65536, matching FFT

The Full Range option generates noise over a span from the lowest frequency for the selected FFT length to half the sample rate, giving the widest bandwidth and the greatest low frequency content. Pink full range has a flat spectrum below 10 Hz (like the pink random noise) to avoid excessive energy at the very lowest frequencies. The Speaker Calibration option generates noise from 500 Hz to 2 kHz, producing a signal with its energy centred on 1 kHz. Subwoofer Calibration generates noise from 30 Hz to 80 Hz. CTA-2034 filtering applies shaping in accordance with ANSI/CEA-2034-A Standard Method of Measurement for In-Home Loudspeakers and has a 12 dB crest factor. The octave and 1/3 octave options filter the signal at the selected centre frequency with 1 octave or 1/3 octave total bandwidth. Custom Filtered allows low and/or high cut frequencies to be set arbitrarily, subject to a minimum bandwidth of 10%. The Sub cal and Spkr cal signals are brick wall filtered. Other filtered signals offer choice of Butterworth high pass and low pass with filter order from 2nd (12 dB/octave) to 8th (48 dB/octave) or brickwall.

REW automatically adjusts the signal levels for the various options and filter settings so that the RMS values reflect the setting in RMS Level. Clipping of peaks will occur at RMS levels of -6 dBFS or above.

The Pink PN sequence can be saved to a wave file and used to generate a test file or disc to be played on a system whose response is to be measured. Make sure that the selected sample rate corresponds with the format of the disc to be made - for example, 44.1kHz should be used if generating a CD, or 48kHz for a DVD. When measuring the system the sample rate and FFT length must be the same as used for the test file or disc.

White noise has much more energy at high frequencies than pink noise and is not recommended for use with speakers, it may damage tweeters at high levels.

Linear Sweep, Log Sweep

The Signal Generator can produce sweeps with configurable start frequency, end frequency, duration and linear or logarithmic progression. Sweep duration can be up to 60 seconds. If the "Loop" box is checked the sweep will repeat continuously. The sweeps have a configurable raised cosine fade in and fade out (which can be set to zero for no fade). When a sweep is playing the current sweep frequency is shown on the panel.

Measurement Sweep

The Measurement Sweep signal is used by REW when measuring system response. It consists of a logarithmic sweep from half the start frequency to twice the end frequency (with an overall limit of half the interface sample rate) to provide accurate measurement over the selected range. If the start frequency is below 20Hz the signal begins with a linear sweep from DC to 10Hz, followed by a logarithmic sweep from there to the end frequency. The sweep duration is set using the Length control. This signal is selected automatically to make sweep measurements. Measurement sweeps can be saved to WAV files to use for offline measurement. Note that the measurement sweep may change between REW versions so always use a sweep generated with the REW version being used for analysis of the captured response.

Help Index

Scope

The Scope emulates a two channel oscilloscope with an additional math channel. It provides a live view of the audio data on the measurement and reference channels. The channels are colour coded, blue for CH1 (the measurement channel), yellow for CH2 (the reference channel) and green for the math channel. REW captures input data in blocks of 8k samples, so the scope receives blocks of samples at a time which makes it behave a little differently to a hardware scope.

The top of the screen shows the currently selected channels, their voltage scale, the timebase and the trigger settings.
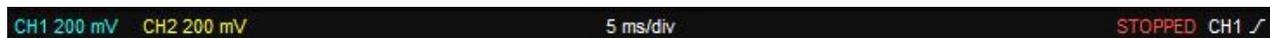
The bottom of the screen shows the capture buffer, with the currently displayed portion highlighted. The position of the trigger is shown as white line. When the scope is stopped the highlighted portion can be dragged to change the region displayed in the main scope area. For finer control of the displayed region the main display can be dragged by clicking and dragging it.
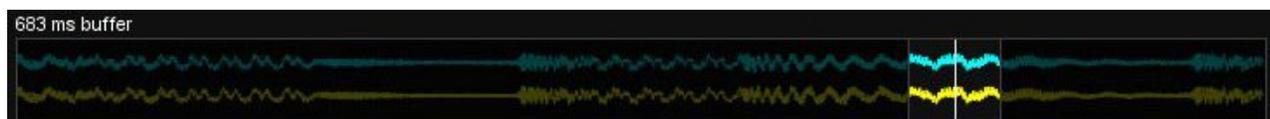
The trigger level and position are shown by white triangles at the left and bottom of the display. They can be adjusted by clicking on them and dragging them to the desired position. The trigger position is shown on the trace by a crosshair symbol.
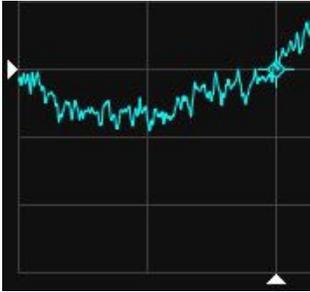

The control panel has channel selectors and voltage scale controls for each channel, a screen capture button, a button to turn cursors on and off, the timebase control, trigger controls and a button to start and stop capture. The timebase can also be adjusted using the mouse wheel when the cursor is in the display area.

Channel settings

Next to the channel selects is a button (cog icon) for channel settings, including the channel offsets, math function, input scaling, full scale lines, persistence and data export. Offsets can be reset to zero using the **Reset** buttons. The zero volts position for each trace is shown as a triangle in the trace colour on the right of the display. Offsets can be adjusted using the channel settings controls or by clicking and dragging on the offset indicators. If **Full scale lines** is selected dotted lines will be drawn at the voltages corresponding to positive and negative full scale for each input channel. The persistence setting emulates analog persistence, with the last 10 sweeps shown progressively more faintly.
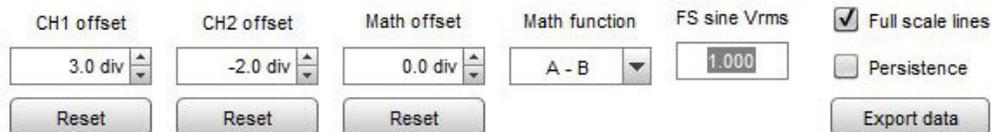
Export data writes the scope data and its settings to a text file using the delimiter set in the File → Export menu. The image below shows a file exported as .csv and loaded into a spreadsheet. Note that data is exported at the captured sample rate as normalised float sample values, where 1.0 is full scale. The sample value corresponding to one volt is included in the header information.
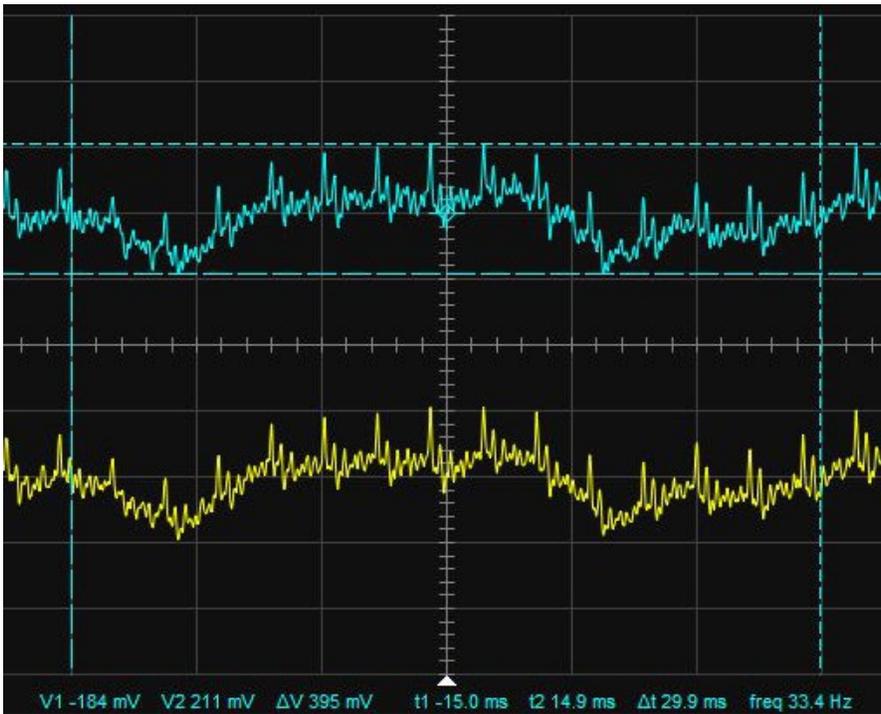
	A	B	C	D
1	Scope data saved by REW V5.20 Beta 12			
2	Trigger channel:	CH1		
3	Trigger edge:	Rising		
4	Trigger position:	2	div	
5	Trigger index:	30767.54		
6	Trigger level:	0.6	div	
7	Trigger hysteresis:	0.1	div	
8	Trigger holdoff:	0	s	
9	Trigger mode:	Manual		
10	Timebase:	0.02	s/div	
11	Sample rate:	48000		
12	One volt:	0.707107		
13	Date:	09-Jun-2019 19:07:13		
14	CH1 scale:	0.2	V/div	
15	CH1 offset:	3	div	
16	CH2 scale:	0.2	V/div	
17	CH2 offset:	0	div	
18	Math scale:	1	V/div	
19	Math offset:	-3	div	
20	Time	CH1	CH2	CH1+CH2
21	0	0	0	0
22	2.08E-05	3.05E-05	0	3.05E-05
23	4.17E-05	0	3.05E-05	3.05E-05
24	6.25E-05	0	-3.05E-05	-3.05E-05

Cursors

The cursor button shows or hides the scope cursors. The settings icon next to the button brings up a menu to control the cursor operation. There are two cursor modes, **Manual** and **Tracking**, In both modes the readings for the cursors are shown below the graph. Times are relative to the trigger position, which is t=0.

Manual mode provides two time cursors and two voltage cursors. Either or both sets can be shown according to the settings selection. The channel selection determines which channel the voltage cursors are referenced to. Cursors are drawn in the colour of the selected channel. Each pair of cursors has a primary cursor (1, long dashes) and a secondary cursor (2, short dashes). When the primary cursor is moved the secondary moves with it, the secondary cursor can be moved independently. Cursors are moved by clicking and dragging them.

Mode	Channel	Cursors
● Manual	● CH1	○ Time
○ Tracking	○ CH2	○ Voltage
	○ Math	● Time & voltage

Tracking mode provides two time cursors which can be adjusted and two voltage cursors which track the selected channel. The cursors can show readings from the same channel or from separate channels according to the settings.

Trigger settings

The scope can be triggered on channel 1 or channel 2, rising or falling edge. The trigger level, hysteresis and holdoff can be adjusted in the trigger settings menu. They can also be adjusted by clicking and dragging on the trigger indicators. When dragging the trigger level the hysteresis can be adjusted using the mouse wheel. In **AUTO** mode the scope will sweep automatically if a trigger has not been detected within the span of the display. In **SINGLE** the scope will wait for a trigger and stop capture after a single sweep.

Position	Level	Hysteresis	Holdoff
20 %	0.0 div	0.1 div	0 ms
	0V	50.0 mV	

[Help Index](#)

Level Meters

The Level meters show the RMS level as a coloured bar and a numeric value at the bottom of the meter and the peak value as a red line and a numeric value at the top of the meter. Levels shown on the bar are in dB below Full Scale.

The numeric value for the output meters is in the units that have been selected on the generator (dBu, dBV, Volts or dBFS).

The numeric value for the input meters is in the units selected by right-clicking on the meter (dB SPL, dBFS, dBu, dBV, dBW, Volts or Watts). Values other than dBFS depend on the input full scale setting on the RTA window or, for SPL, the SPL meter calibration. The SPL figure is unweighted and does not take into account any mic calibration file effects.
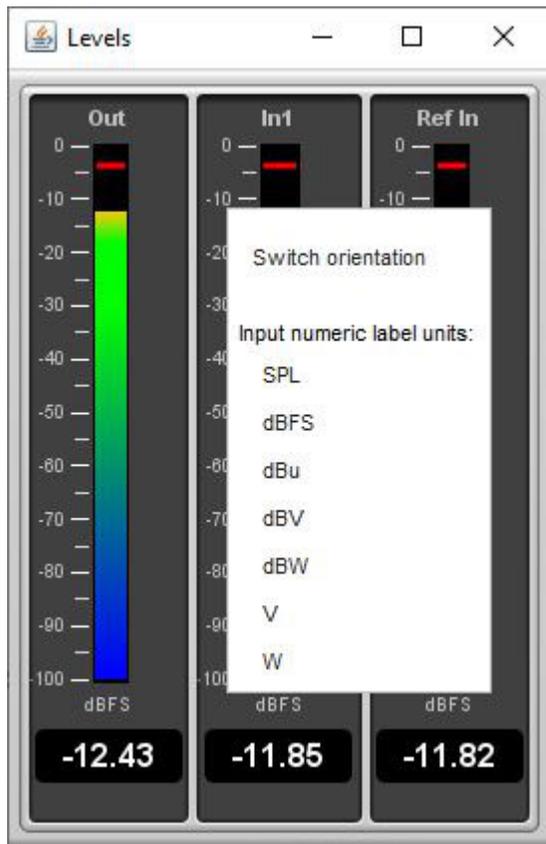
The orientation of the meters can be switched between vertical and horizontal using the right click menu.

When using multi-input capture (Pro upgrade) the meters can show levels for up to 16 inputs, according to the setting in the View preferences.

Help Index

Graph Panel

The graph panel shows plots for the currently selected measurement. The plots are selected via the buttons at the top of the graph area.

| SPL & Phase | All SPL & Impedance | Distortion | Impulse | Filtered IR | GD | RT60 | Clarity | Decay | Waterfall | Spectrogram | Captured |

The various graph types are:
SPL and Phase
All SPL
Distortion
Impulse
Filtered IR
Group Delay
RT60
Clarity
Spectral Decay
Waterfall
Spectrogram
Captured

Options that affect the appearance of the traces can be found in the View Preferences.

Each trace can be turned on or off via the selection buttons to the left of the trace name in the legend panel. Trace names are in the same colour as the trace itself, whilst the line style for the trace is shown between the label and the trace's value at the current cursor position. If a trace has smoothing applied the octave fraction is shown (1/12th octave in the example below). Moving the mouse cursor over a trace name in the legend will highlight that trace on the graph, fading other traces to help it stand out.

☑ Zaphod	1/48	81.1dB	☑ Phase	1/48	73deg
☑ Min phase	1/48	-85deg	☑ Excess phase	1/48	158deg
☑ Mic/Meter Cal	——	-5.8dB	☑ Soundcard Cal	------	-0.1dB

Capture Graph Image Button

Capture

This button at the top left corner of the graph area allows the current graph view to be captured as an image. It can be activated using the Alt+C shortcut. A dialog pops up to set the desired width of the image (click **Default** to set the image to be the same width as the graph). The thickness of the graph traces can be selected, the REW default is 2 pixels. If the **Include title** box is checked the graph type will be shown at the top of the graph. If the **Include legend** box is checked the image includes the graph legend. **Include cursor** shows the cursor lines and values on the captured image. If **Include timestamp** is selected the current date and time will be shown at the bottom right corner of the graph. If **Monochrome image** is selected the traces will be drawn with one of four different line styles instead of in different colours: solid, dotted, dashed or dash-dotted.

Graph aspect ratio is only applicable for graphs that have a horizontal logarithmic frequency axis and a vertical dB axis. It adjusts the vertical size of the captured graph image so that one decade on the frequency axis (e.g. 20 to 200 Hz) has the same length as a 10, 25 or 50 dB span on the dB axis, per the options offered in IEC263. This ensures a consistent visual appearance of frequency response data and helps avoid graphs that may seem unduly flat or have response slopes that appear shallower or steeper than they are.

Text typed into the box that shows **Type any comment text here** will appear on the graph image in the selected **Comment position**, either above, below or on the graph. Text placed above or below the graph can

be styled using HTML 4.01 tags and aligned according to the **Comment alignment** selection. The **Insert notes from current measurement** button inserts the notes from the currently selected measurement in the comments. The **Preview comment** button allows the formatting to be checked before proceeding. Comment text placed on the graph is centred. The graph will be either saved as JPEG or PNG according to the **Image type** selected, or copied to the clipboard.

Scrollbars Button

The Scrollbars button toggles scrollbars for the graph area on/off, hiding the scrollbars provides more room for the graph. The setting is remembered for the next startup. If the scrollbars are off, the graph can still be moved by holding down the right mouse button while in the graph area.

Frequency Axis Button

Freq. Axis

The Freq Axis button toggles the frequency axis between logarithmic and linear modes. This function is also available via a command in the Graph menu and the associated shortcut keys.

Graph Limits Button

Limits

The Graph Limits button allows desired top, left, bottom and right graph limits to be defined. A dialog pops up in which the values are entered, they are applied as they are entered or by clicking the **Apply Settings** button.

Graph Controls Button

Controls

The Graph Controls button brings up a menu of control options for the currently selected graph type, if there are any.

Horizontal Axis Zoom Buttons

The horizontal axis zoom buttons appear when the mouse pointer is inside the graph area, they zoom in or out by a factor of approximately 2 centred around the cursor position. You can also zoom with the keyboard.

Vertical Axis Zoom Buttons

The vertical axis zoom buttons appear when the mouse pointer is inside the graph area, they zoom in and out on the Y axis. You can also zoom with the keyboard.

Variable Zoom

REW provides a variable graphical zoom capability by either pressing and holding the middle mouse button, or pressing the right button then, while the right button is held down, pressing and holding the left button and dragging the pointer.

When variable zoom is active a cross is displayed, split into quadrants allowing horizontal and/or vertical

zooming in our out depending on the mouse position. The amount of zoom is governed by how far the mouse pointer is dragged from the start position.

Zoom to Area

When the Ctrl key is pressed followed by the right mouse button a zoom box can be drawn by dragging the mouse. Note that on a Mac it may be necessary to press both the Ctrl and fn keys, or to hold down Ctrl and drag two fingers on the trackpad. Measurement cursors are shown on the outside of the box, to zoom to the shaded area click within it. If the shaded area is too small to zoom in to a message will indicate which dimension is too small for zooming and what the limit is to allow zoom.
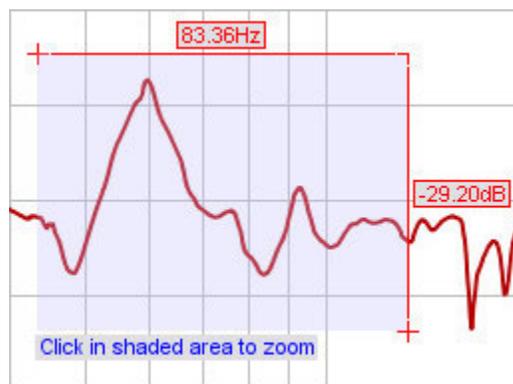
Zoom with Keyboard

To zoom in on the horizontal axis press `Shift+x`, to zoom out press `x`. To zoom in on the vertical axis press `Shift+y`, to zoom out press `y`. The zoom is centred on the cursor position. You may need to click on the graph first.

Undo Zoom

To undo the last Variable Zoom or Zoom to Area, press Ctrl+Z or select the Undo Zoom entry in the Graph menu. This will restore the graph axes to the settings they had when the right or middle mouse button was last pressed. This Undo feature can be used even if you have not zoomed, just press the right mouse button when the axis settings are to your preference then you can return to these settings (undoing any subsequent movements or control changes) by pressing Ctrl+Z.

Arrow Keys

The arrow keys can be used to move the graph cursor in single pixel increments. You may need to click on the graph first. To move the graph instead of the cursor, hold down the `Shift` key while pressing the arrow keys.

Help Index

SPL and Phase Graph

The SPL and Phase plot (or Impedance and Phase for an Impedance measurement) shows the frequency and phase responses of the measurement. The frequency response is labelled with the measurement name, the phase response uses a dotted trace and the right hand plot Y axis. The left hand Y axis can be set to show dB SPL, dBFS, dBr, dB V/V, dBu, dBV, dBW, volts, watts or, for impedance traces, ohms. The dBr and dB V/V values are effectively a transfer function view, showing either the relative input to output dBFS levels (dBr) or voltage levels (dB V/V). Values for dBW and watts are derived from the measured voltages using the reference impedance value in the RTA Appearance settings.

For stepped level measurements the graph will show a plot of the input level versus the generator level, and a linearity plot showing the ratio of input level to generator level.

For impedance measurements when the mouse cursor is within the graph panel the equivalent series resistance + inductance or resistance + capacitance and parallel resistance||inductance or resistance||capacitance of the impedance at the cursor position is shown at the bottom left corner of the graph. This may be useful when making measurements of inductors or capacitors to check their value, but the Component Model provides a much more accurate equivalent circuit.
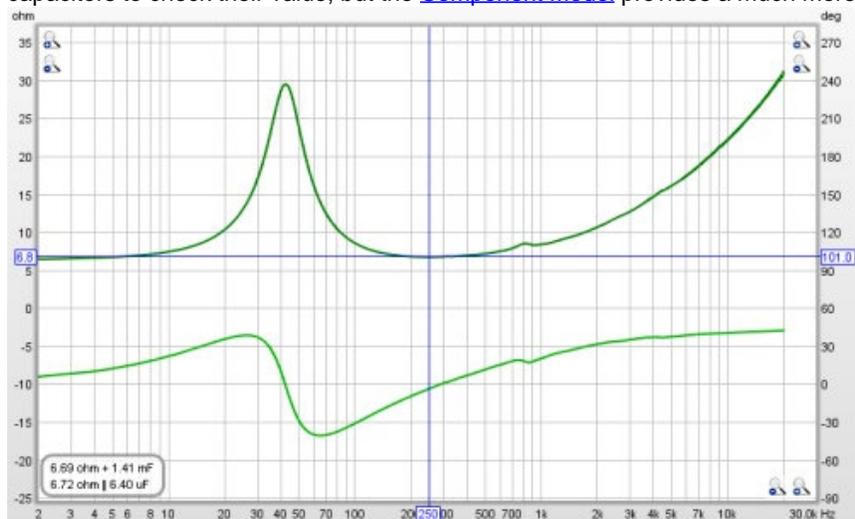
Note that to have valid phase information it is necessary to remove any time delays from the Impulse Response. A time delay causes a phase shift that increases with frequency - for example, a delay of just 1 ms results in a phase shift of 36 degrees at 100Hz but 3,600 degrees at 10kHz, because 1 ms is 1/10th of the 10 ms period of a 100Hz signal but is 10 times the 0.1 ms period of a 10kHz signal, and each period is 360 degrees. The time delay of a measurement can be adjusted by changing the zero position of the time axis using the **Offset t=0** controls, or by using the **Estimate IR delay** control, both described below.

In addition to the measured phase, the plot can show minimum and excess phase plots that result from generating a minimum phase version of the response, described further below. The plot also shows any mic/meter or soundcard calibration data for the measurement. The calibration data can be changed or removed by selecting **Change Cal...** on the measurement panel.

Minimum Phase/Excess Phase

If the **Generate minimum phase** button has been used to produce a minimum phase version of the response the minimum and excess phase traces are activated. The minimum phase trace shows the lowest phase shift a system with the same frequency response as the measurement could have, while the excess phase trace shows the difference between the measured and minimum phase. Note that it is best to make full range measurements if the minimum phase response is to be generated as a good result relies on measuring far beyond the bandwidth of the system being measured. Even then there are limits to the accuracy of minimum phase responses generated from sampled data, particularly at frequencies above about one quarter of the sample rate. There are options to improve the result, for details see Generating minimum phase below. For more about minimum and excess phase and group delay see Minimum Phase.

Mic/Meter Cal

The Mic/Meter Cal trace shows the frequency response of the Mic calibration data for this measurement (the calibration file to use for new measurements is specified in the Cal files Preferences). If **C Weighted SPL Meter** was selected this curve will show the effect of C weighting (outside the range of the calibration data file, if there is one). The trace is not shown if there is no mic/meter calibration data. The trace is drawn relative to the middle of the graph.

Soundcard Cal

The Soundcard Cal trace shows the measured frequency response of the soundcard relative to its level at 1kHz (if a calibration file has been loaded via the Soundcard Preferences). The trace is not shown if cal data has not been loaded. The trace is drawn relative to the middle of the graph. Fractional octave smoothing can be applied or removed via the Graph menu and its shortcut keys. The smoothing is applied to the SPL, phase and Group Delay traces. This is mainly used for full range measurements, as reflections can cause severe comb filtering which makes it difficult to see the underlying trend of the response. Smoothing should rarely be used for low frequency measurements as it obscures the true shape of the response. When smoothing has been applied an indicator appears in the trace legend.

Graph metrics

If a frequency range is selected by holding the shift key then pressing and dragging with the left mouse button a best-fit line for the selected range is displayed, with lines above and below it to show +3 dB and -3 dB. The slope of the line in dB/octave is shown (if the frequency axis is logarithmic) along with the span of the data relative to the best-fit line and the minimum, maximum and average of the range. The range can be adjusted after selection by dragging the start or end of the area, or the whole range can be moved.

SPL and Phase Controls

The control panel for the SPL and Phase graph has these controls, there may be more or fewer controls depending on the measurement type.
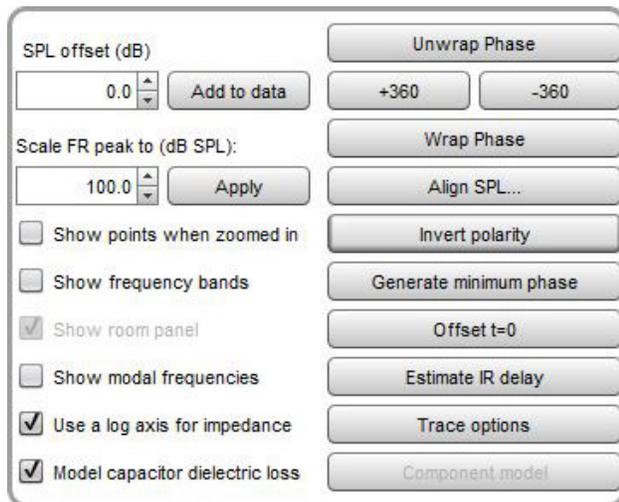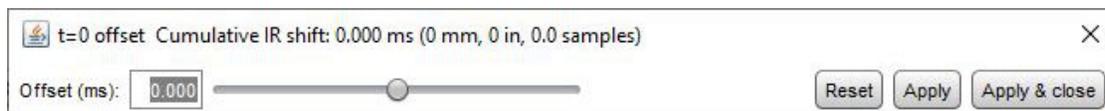
The phase trace normally wraps at +180/-180 degrees. This is because phase is cyclic over a 360 degree range (+90 is the same phase as -270). The trace can, however, be displayed without wrapping which is what the **Unwrap Phase** control does. A difficulty with unwrapped phase is knowing where the correct zero phase is, another is being able to view parts of the trace where the unwrapped value has become very large. The unwrapped phase is offset (by a multiple of 360 degrees) so that it is within the range -180..180 degrees at the cursor frequency. The **+360** and **-360** buttons will also shift the phase trace in 360 degree steps.

Wrap Phase changes the phase trace back to a conventional wrapped view with vertical lines where the trace crosses 180 or -180 degrees.

Align SPL... adjusts the measurement so that the average SPL over a selected span is a specified level.

Invert polarity inverts the polarity of the measurement which shifts the phase values by 180 degrees.

Offset t=0 allows the position of time zero in the impulse response to be altered, with a live preview of the effect the offset will have on phase. The measurement is not changed unless either the **Apply** or **Apply & close** button is pressed. Fine adjustment of the offset can be made with the left and right arrow keys after clicking on the slider knob. If a timing reference was used the System Delay figure (which can be viewed in the measurement Info panel) is shifted by the same amount as the zero time. The cumulative shift that has been applied to the impulse response is shown at the top of the dialog.

Estimate IR delay calculates an estimate of the time delay in the measurement by comparing it with a minimum phase version. The delay it calculates can be removed from the impulse response by pressing the **Shift IR** button on the panel shown after the delay is calculated. Note that shifting the impulse response will clear any spectrogram which had been generated as the plot would no longer be valid.

The trace offset value moves the graph position, but does not alter the data so the legend values do not change. If the **Add to data** button is pressed the current offset value is transferred to the measurement data and the legend readings will update accordingly.

The **Scale FR Peak** control adjusts the SPL offset to achieve a desired maximum SPL figure in the frequency response. This may be useful to adjust levels for an imported impulse response, for example.

If **Show points when zoomed in** is selected the individual points that make up the SPL and phase responses are shown on the graph when the zoom level is high enough for them to be distinguished (which may only be over part of the plot)

If **Show frequency bands** is selected the audio frequency bands are shown in a stripe above the graph. The bands are:

- Infrasound: below 20 Hz
- Sub bass: 20 to 60 Hz
- Bass: 60 to 250 Hz
- Low midrange: 250 to 500 Hz
- Midrange: 500 Hz to 2 kHz

Upper midrange: 2 to 4 kHz
- Presence: 4 to 6 kHz
- Brilliance: 6 to 20 kHz

Show room panel is only applicable for measurements generated from the [Room simulator](#). If selected a plan view of the room configuration used to generate the simulated response is shown.
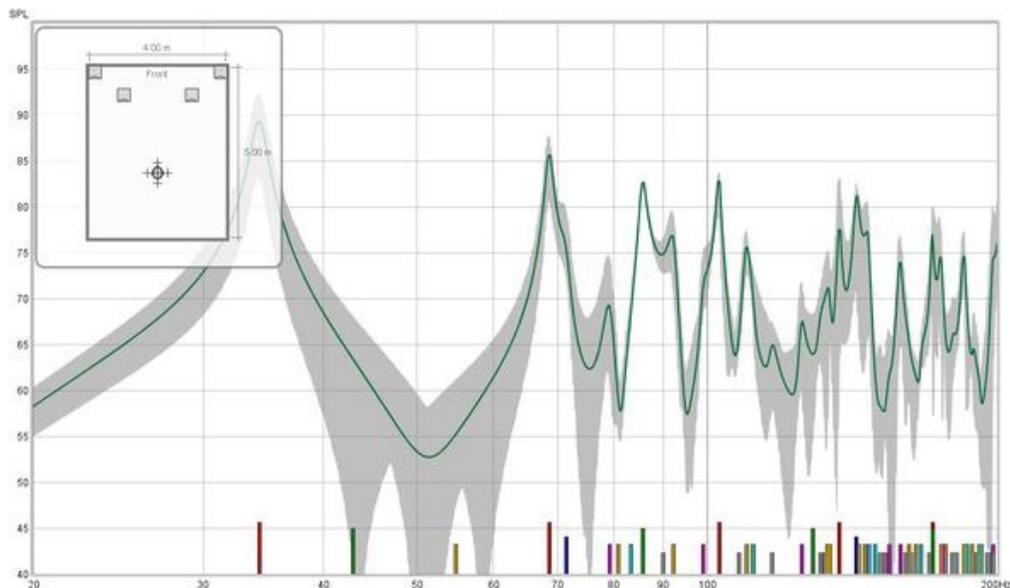
If **Show modal frequencies** is selected the theoretical modal frequencies for the room dimensions entered in the [Modal Analysis](#) section of the EQ Window (or in the Room simulator for responses generated from there) are plotted at the bottom of the graph.

The **Trace options** button brings up a dialog that allows the colour and line type of the graph traces to be changed. If a change is made it will be used for all measurements shown on this graph. Traces can also be hidden, which will remove them from the graph and from the graph legend.

When impedance is being plotted the axis has a span from 0 to 1 kohm. If a larger impedance range is required the axis can be switched to logarithmic with a range up to 1 Mohm using the **Use a log axis for impedance** check box. If selected a log axis will be used wherever impedance is plotted.

Generating minimum phase

Generate minimum phase will produce a minimum phase version of the measurement. Using this control also generates a minimum phase impulse plot and, for measurements that already had an impulse response, minimum and excess phase and group delay plots. If the measurement did not have an impulse response the minimum phase impulse is set as the IR for the measurement, this allows an IR and phase to be produced for measurements which do not have them, such as imported text files.

If the system that was measured was inherently minimum phase (as most crossovers are, for example) the minimum phase response should be the same as removing any time delay from the measurement. Individual driver measurements are usually minimum phase, but multi-driver loudspeaker measurements are usually not. Room measurements are typically not minimum phase except in some regions, mainly at low frequencies. For more about minimum and excess phase and group delay see [Minimum Phase](#).

There are limits to the accuracy of the minimum phase response that can be generated for sampled data systems which mean they may not reflect the true minimum phase response of the system being measured, particularly at frequencies above about one quarter of the sample rate. DAC reconstruction and ADC antialias filters are commonly not minimum phase and have very steep roll-offs, so results for systems whose bandwidth exceeds half the sample rate will be affected by them. REW's responses are generated by using the FFT, which means they will have zero phase at DC and half the sample rate. REW uses the real cepstrum to produce its minimum phase response.

The **Generate minimum phase** button brings up a dialog to configure the settings for producing the minimum phase response. The dialog can also be shown by using the Ctrl+Alt+M shortcut. Minimum phase generation is not possible for some types of measurement (stepped sine, multi-input average, impedance).

If **Include cal file effects** is selected the response used to generate the minimum phase will include any calibration files the measurement has. That means their influence will become part of the minimum phase impulse response. Note that measurement impulse responses themselves do not include cal file effects.

If the measurement does not cover the full frequency span from 0 to half the sample rate, or the device response drops into the noise floor at its extremes, there are options to define what data the minimum phase calculation should use outside the measurement range or the valid range of the measurement. The first option is to **Replicate data outside range**, which repeats the measurement start value for frequencies below the start frequency and the end value for frequencies above the end frequency. This may be useful for measurements that cover a limited frequency span, or have been imported from text files with a limited span.

An alternative is to add 'tails' to the measurement at low frequencies, high frequencies or both. Some systems will have a known low frequency roll-off. For example, sealed box loudspeakers ultimately roll off at 12 dB/octave, while ported loudspeakers typically roll off at 24 dB/octave (in both cases below their respective resonant frequencies). The measured low frequency response can be replaced by a tail with a user-defined slope to facilitate the generation of an accurate low frequency minimum phase response. To help with this a slope overlay is shown at the cursor position, indicating the local slope of the response.
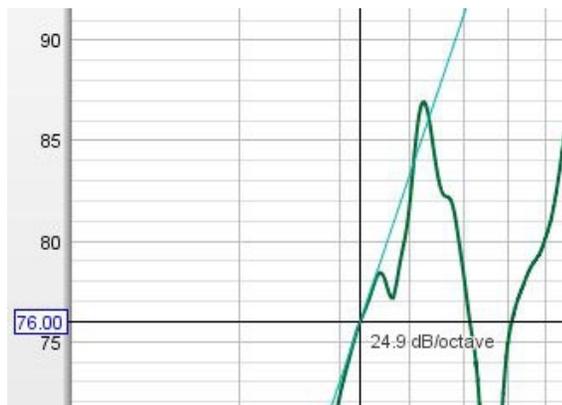

A tail may also be added at high frequencies, though the results may be less successful if the bandwidth of the system being measured is not far below the sample rate. If a high frequency tail is being added REW will oversample the response by a factor of 4 to reduce the influence of the sample rate. There is also an option to **Apply frequency warping to HF tail** which replicates the effect approaching half the sample rate typically has on the slopes of digitally-generated filters.
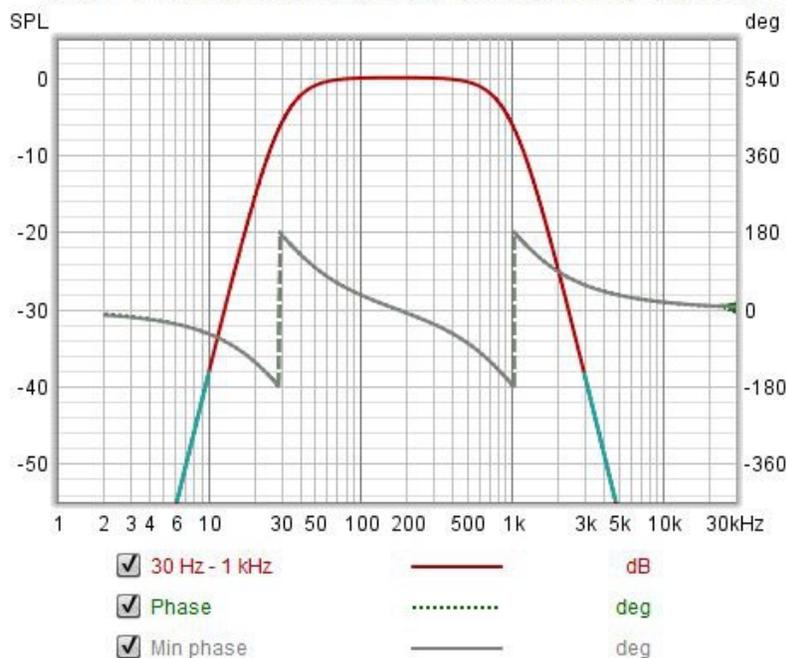
The images below show the effect of adding tails to a crossover response measurement. The tails are overlaid on the

measurement in cyan. The result is very good, but this is an ideal case for the high frequency behaviour since the crossover was measured at 96 kHz and rolls off from 1 kHz, so the measurement bandwidth far, far exceeds the system bandwidth.

No tails

✓ 30 Hz – 1 kHz	———	-0.0 dB
✓ Phase	··········	15 deg
✓ Min phase	———	15 deg
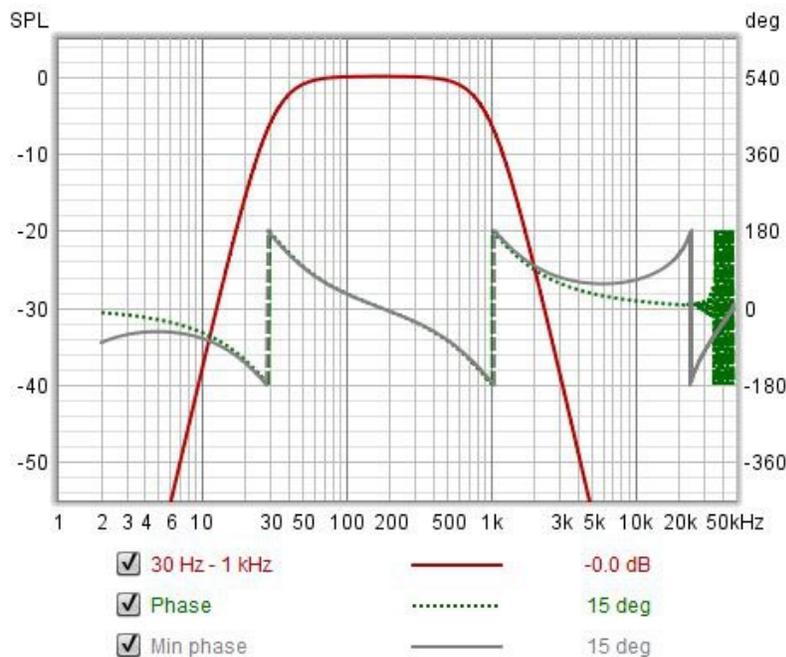
LF tail 10 Hz 24 dB/octave, HF tail 3 kHz -24 dB/octave

✓ 30 Hz – 1 kHz	———	dB
✓ Phase	··········	deg
✓ Min phase	———	deg

Multi-input Adjust average

Multi-input average measurements have an **Adjust average** button which brings up a dialog that allows the SPL alignment and weighting of the individual inputs to be adjusted before the average is recalculated. The **Revert to measured** button removes all SPL alignment offsets and recalculates the average.

Component model

The **Component model** button is only shown for impedance measurements and only enabled for measurements of inductors or capacitors. REW identifies the measurement as being of a capacitor if the phase at 20 Hz (or the start of the measurement, if higher) is below -60 degrees, or an inductor if the phase at the end of the measurement is above 45 degrees. For those measurements pressing the button will carry out a curve fit over the range from 10 Hz to 20 kHz for inductors or 100 Hz to 20 kHz for capacitors (or the measurement range if smaller) to derive equivalent circuit component values. Measurements should span at least those ranges, preferably beyond them so that the model behaviour can be seen beyond the fit range. The equivalent circuit is shown below the button. The equivalent circuit impedance and phase are shown as dashed lines overlaying the measured traces. When measuring components make sure the component lead length is close to the length it will have in circuit, otherwise the measurement will include lead resistance and inductance which will not be present when the part is used.

For capacitors the equivalent circuit is a series combination of a resistor (the ESR), a capacitor and an inductor (likely to be very small at audio frequencies, not shown if less than 1 nH). The ESR value is likely to be over-estimated for very small capacitors (below 100 nF or so) as the impedance at the end of the model fit range (20 kHz) is still quite high, so it is limited to 1 ohm maximum in those cases. If the option to **Model capacitor dielectric loss** is selected an additional R and C may be included in parallel with the capacitor, to capture the effects of dielectric loss. The dielectric loss components are omitted if the loss capacitance is less than 0.1% of the main capacitance. The series R is omitted if it is less than 1 milliohm. Whilst accounting for the dielectric loss usually gives a more accurate fit to the impedance data the equivalent circuit becomes more complex and does not yield a single capacitor value to use. A combined Csum capacitance value is shown, but this may only be valid at low frequencies. The best option if a single figure is required is to use the impedance readout value shown at the bottom left of the graph when the cursor is at the frequency of interest. The images below show results without and with dielectric loss components.

Electrolytic capacitor measurement, nominally 4700 uF
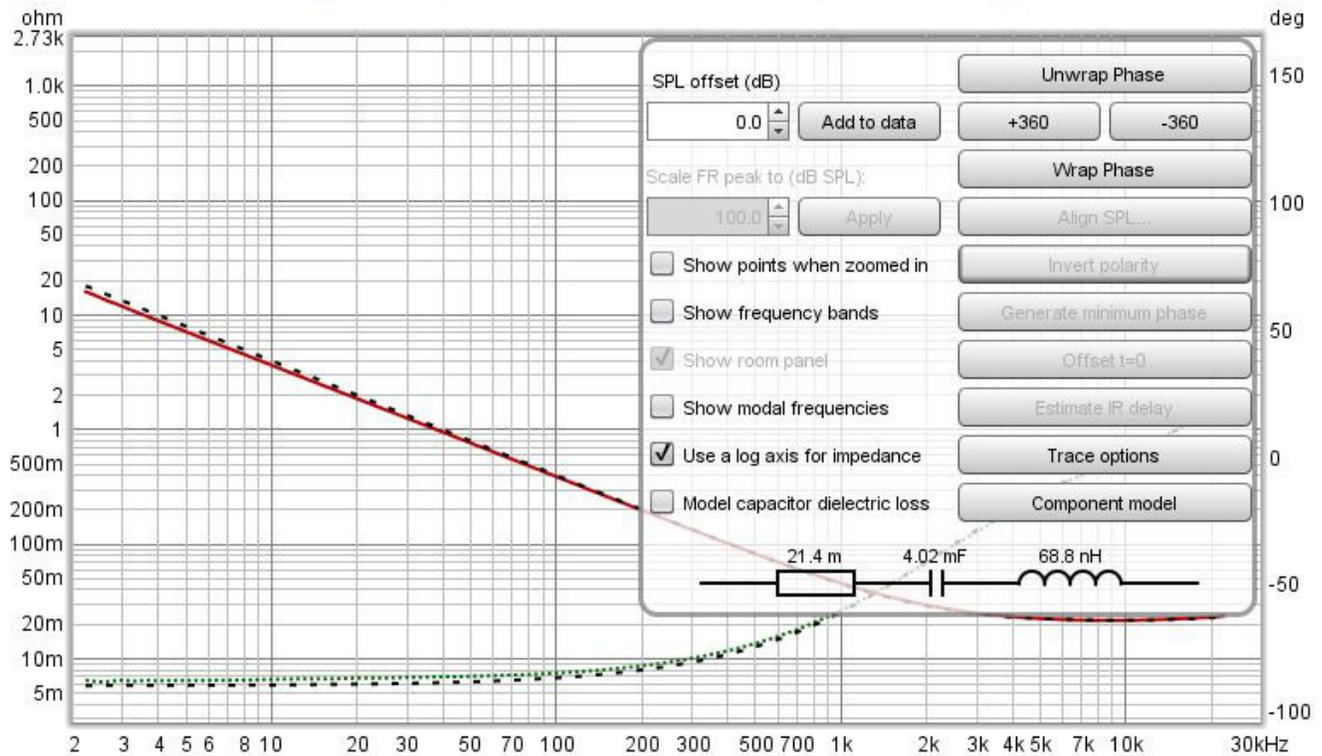
Electrolytic capacitor measurement, nominally 4700 uF
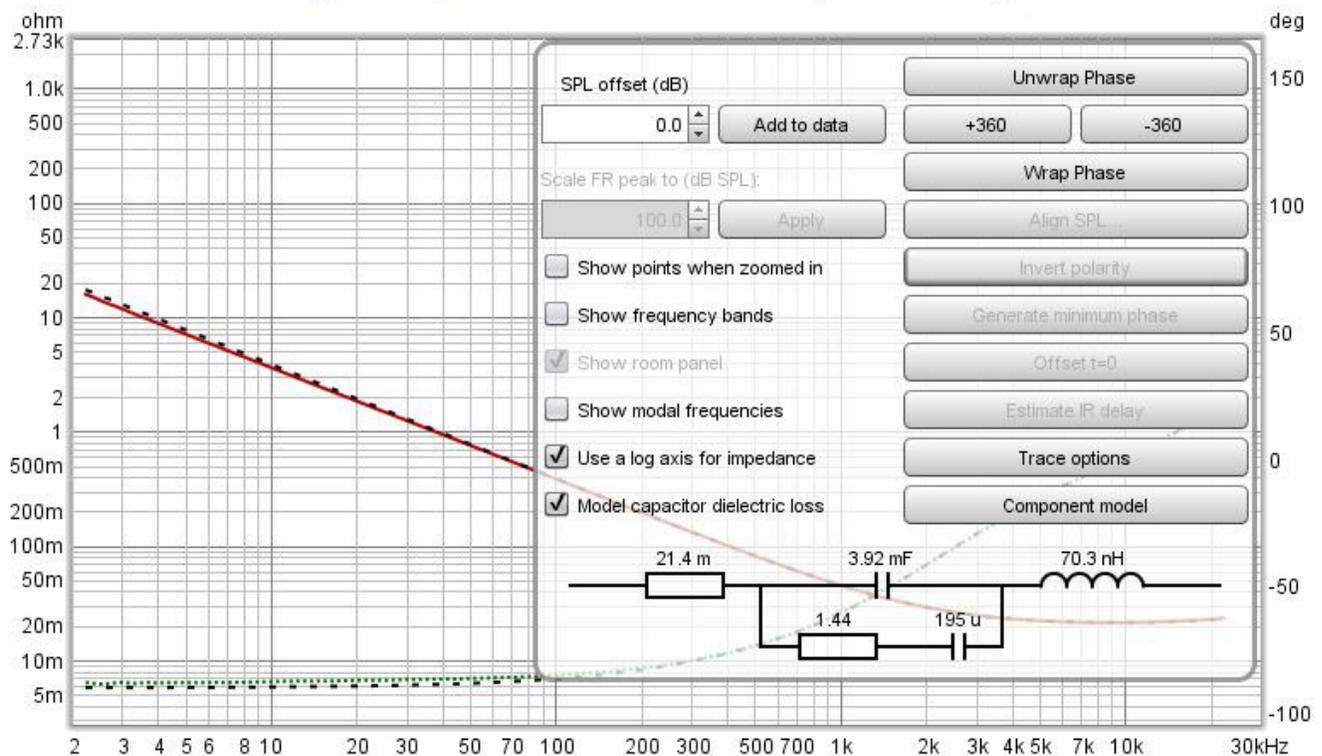
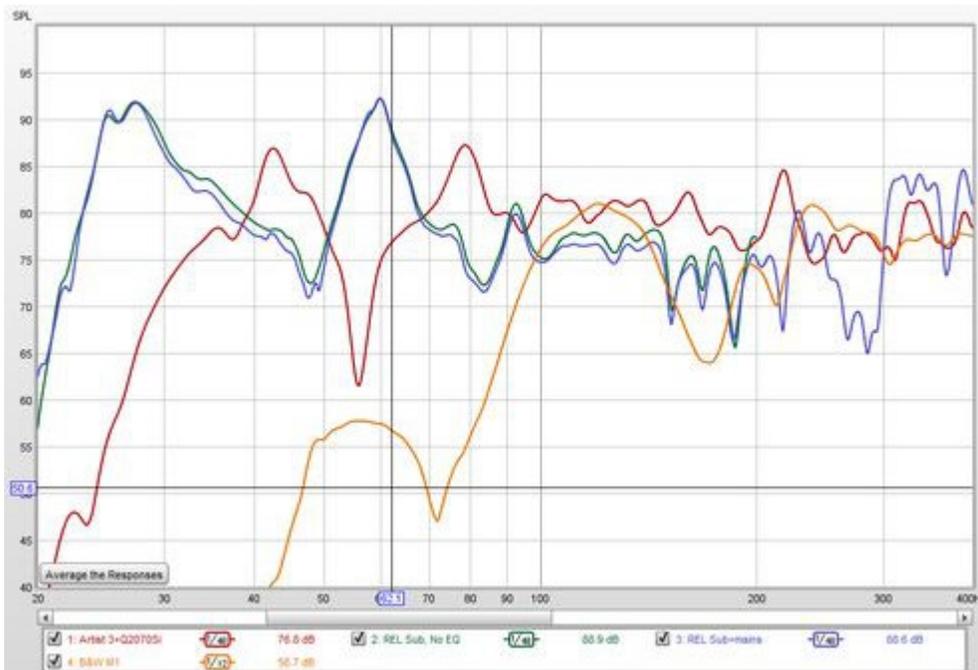
For inductors the equivalent circuit is an LR2 model consisting of a series resistor, a series inductor and a resistor and inductor in parallel.

## Inductor measurement, nominally 0.9 mH
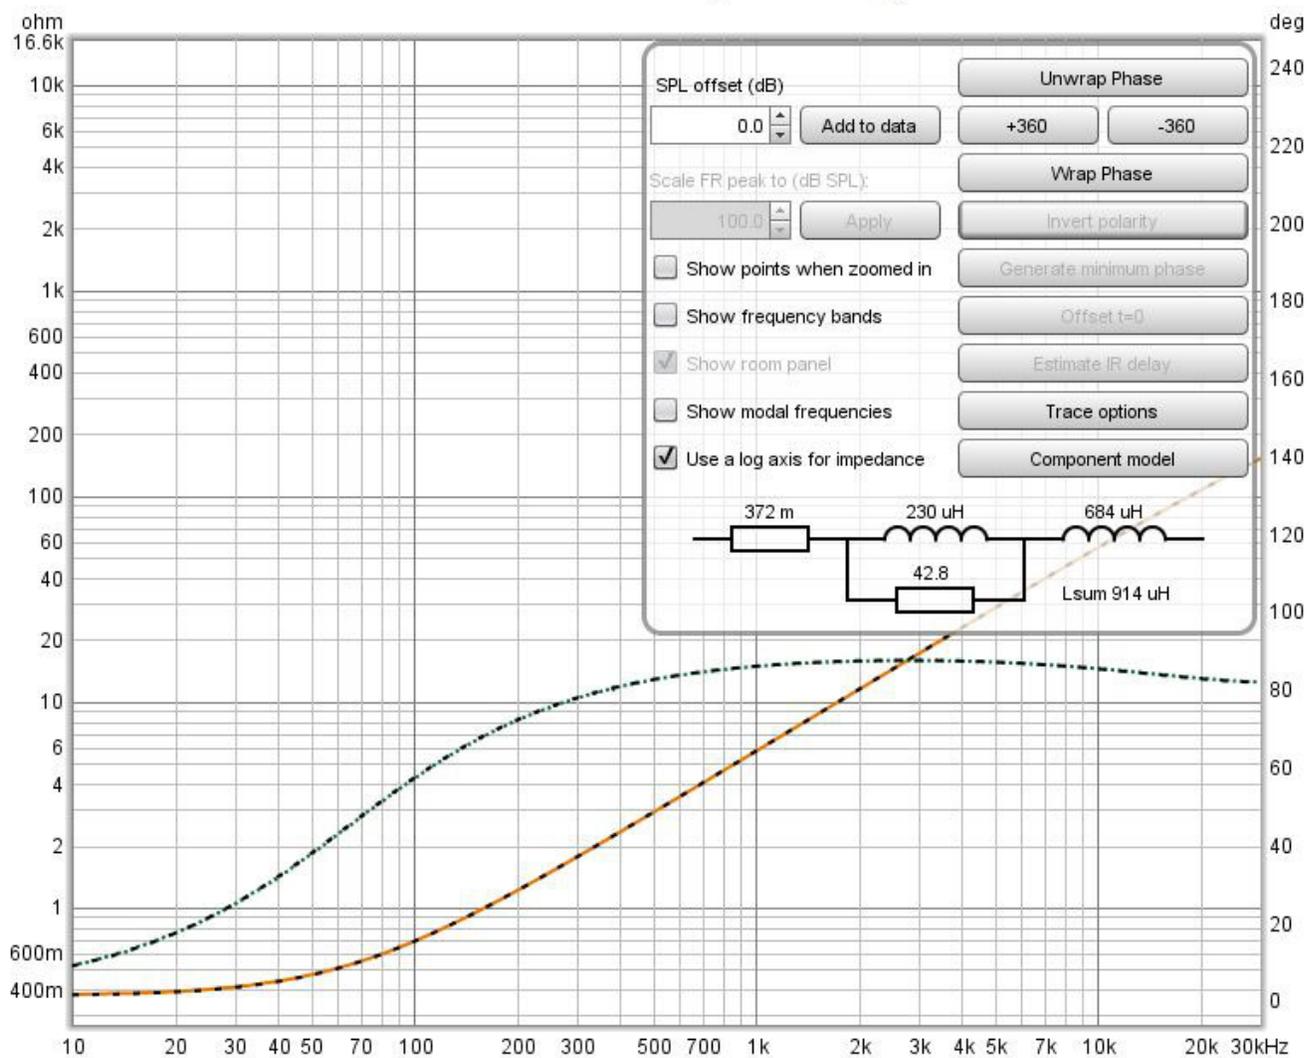
All SPL Graph

The All SPL graph is an overlay graph that shows all measurements (SPL and/or Impedance) that have been made. It allows an average to be generated of all selected traces or arithmetic operations to be carried out on pairs of traces to generate a new trace.

Average The Responses calculates an rms (root mean square) average of the SPL values of those traces which are selected when the button is pressed. That means the dB values are converted to linear magnitudes, those magnitudes are then squared, summed and divided by the number of measurements, the square root of the result is taken, then the value is converted back to dB. Phase is not taken into account, measurements are treated as incoherent. The frequency range of the averaging result covers the region where the traces overlap, for example if one trace was measured to 200Hz, another to 500Hz and a third to 1000Hz the average would range to 200Hz (to the lowest end frequency). If only a single trace is selected the result has the magnitude data from the source measurement and no phase data.

All SPL Controls

The control panel for the All SPL graph has these controls:
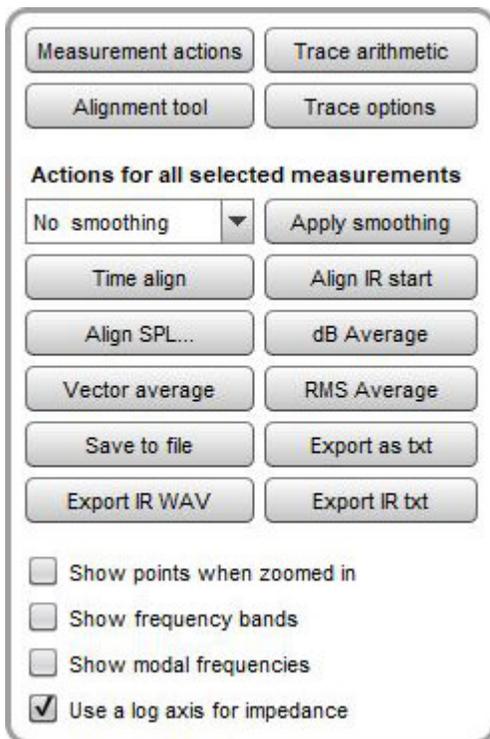
The controls are split into four groups:

- Those which apply to an individual measurement, shown by the **Measurement actions** button
- Alignment tool controls, shown by the **Alignment tool** button
- Trace Arithmetic controls, shown by the **Trace arithmetic** button
- Those which apply to all measurements currently selected in the graph legend

The **Actions for all selected measurements** are:

- A smoothing control that allows the fractional octave smoothing setting for all the currently selected traces to be changed. The chosen setting is applied using the **Apply smoothing** button
- **Time align**, which brings all the currently selected measurements into time alignment. If the measurements have been made with a timing reference (a loopback connection or an acoustic timing reference) the impulse response is shifted according to the measurement delay value, taking into account any IR timing offsets which have been applied since the measurement delay was calculated. Measurements which have been made without a timing reference are shifted according to the estimated IR delay. Time alignment can only be applied to measurements that have an impulse response.
- **Align IR start**, which aligns all the currently selected measurements using their IR start times.
- **Align SPL...**, which adjusts all the currently selected SPL measurements so that they have the same average SPL over a selected span.
- **dB average**, which averages the dB SPL values of the currently selected traces. This produces results that are closer to what one might intuitively expect from looking at the traces, but gives equal weight to peaks and dips which masks the magnitude difference between them. For example, a +20 dB peak has a magnitude that is 100 times larger than a -20 dB dip. dB averaging them produces 0 dB, 10 times smaller than the peak and 10 times larger than the dip. dB averaging may be useful when averaging smoothed traces to derive an EQ target, with unsmoothed data the dips would have a disproportionate effect on the result.
- **Vector average**, which averages the currently selected traces taking into account both magnitude and phase. It can only be applied to measurements that have an impulse response and is most appropriate for multiple measurements taken from the same position, or measurements which have been time and level aligned before averaging them.

- **RMS average**, which calculates an rms average of the SPL values of those traces which are selected when the button is pressed. This does the same as the **Average The Responses** button, the dB values are converted to linear magnitudes, those magnitudes are then squared, summed and divided by the number of measurements, the square root of the result is taken, then the value is converted back to dB. Phase is not taken into account, measurements are treated as incoherent. If only a single trace is selected the result has the magnitude data from the source measurement and no phase data. If the measurements were made at different positions (spatial averaging) it is usually best to first use the **Align SPL...** feature to remove overall level differences due to different source distances.
- **Save to file**, which saves the currently selected measurements to an mdat file.
- **Export as txt**, which exports the currently selected measurements to text files.
- **Export IR WAV**, which saves the impulse responses of the currently selected measurements to WAV files.
- **Export IR txt**, which saves the impulse responses of the currently selected measurements to text files.

Measurement Actions

The **Measurement actions** dialog has the following controls:
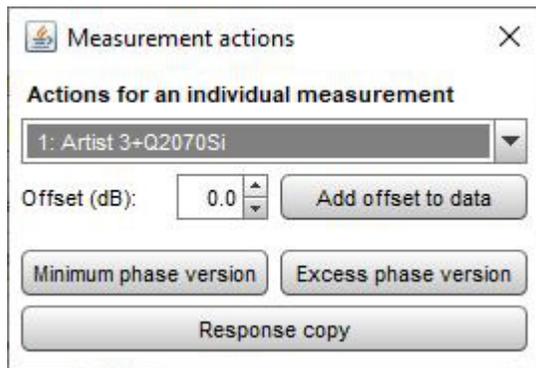

The options for each selected measurement are to:

- Offset the measurement either temporarily or (by using **Add offset to data**) permanently
- Make a **Minimum phase version** of the selected measurement to use in calculations - the copy has a minimum phase impulse response and the same magnitude response and calibration data as the original
- Make an **Excess phase version** of the measurement to use in calculations - the copy is the original measurement divided by a minimum phase version of the measurement
- Make a new measurement with **Response copy** that has the same response as the measurement - the copy has the same impulse response and calibration data as the original but does not include distortion data or equaliser settings, for example

Trace Arithmetic

The **Trace Arithmetic** controls allow the chosen pair of traces to be added, subtracted, multiplied, divided, coherently (aka vector) averaged or merged. If both the chosen traces have impulse responses, the result will also have an impulse response, however the sample rates must be either the same or related by an integer. For example, traces at 44.1kHz and 11.025kHz can be combined via an arithmetic operation, the result will have the higher of the two rates. This allows operations on band limited measurements which may have been decimated to a lower sample rate.
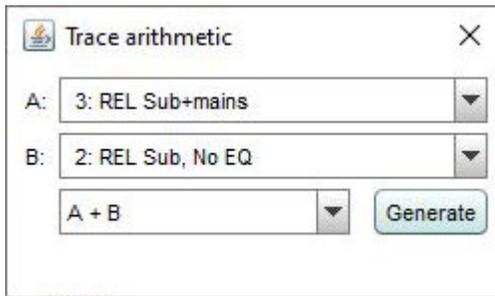
If the traces have incompatible sample rates, or either does not have an impulse response, the result will not have an impulse response, but it may have both magnitude and phase data if both the traces it was applied to had magnitude and phase data, otherwise the result will only have magnitude data and the traces will be treated as incoherent.

The frequency span of the result of an arithmetic operation will be from the lowest start frequency to the highest end frequency of the traces operated on. Outside their frequency range traces are treated as being zero valued, with the exception of the divisor in a division operation which is treated as being unity outside its range. If the measurements actually have significant levels outside the measurement range the zero setting will generate oscillations in frequency and time domains, for best results use traces that span the full frequency range.

Trace arithmetic notes

- For meaningful results measurements that have impulse responses or phase data should be properly time aligned before they are combined. An exception is the Merge operation, for which REW will automatically align both magnitude and phase at the merge frequency, adjusting the trace B time delay as required for the phase match. The amounts of the adjustments are shown in the notes of the newly generated measurement.

- The currently applied impulse response window settings are used for each trace. The result uses the same window settings as trace A unless the operation was Merge B to A, in which case the window settings for trace B (the low frequency portion) are used. The window reference time for the result will be at the peak of the result IR. The left window width will be increased if necessary so that it is at least twice the span from the start of the result IR to its peak to avoid truncation artefacts in the result. Any frequency-dependent settings are excluded, applying an FDW to the result would amount to applying the window twice, as it is already applied to the data used to produce the result.

- The result of arithmetic on measurements that have compatible impulse responses is smoothed using the measurement A smoothing, unsmoothed data is used during the calculations. Other measurements use whatever smoothing they already had applied during the calculations and the result is treated as unsmoothed (or 1/48 octave smoothed if data is 96 PPO).

Alignment Tool

The **Alignment tool** allows gain, polarity and delay adjustments to be made to a pair of measurements with a live preview of the summed result updated as the parameters are changed (the preview uses log spaced data at 96 PPO so may differ slightly from the final result).

When **Show phase traces** is selected the All SPL graph is split with the SPL traces at the top and phase traces for the two selected measurements and the aligned sum at the bottom.
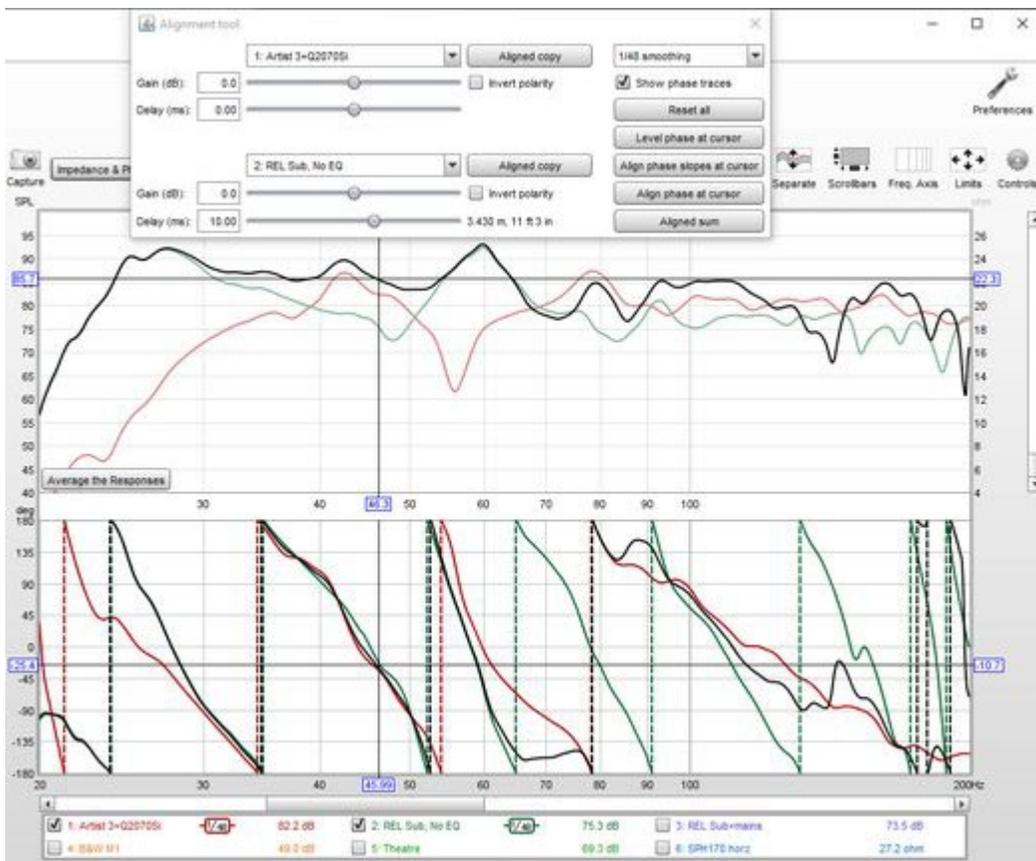
The **Aligned copy** buttons generate a copy of the selected measurement with the alignment gain, delay and polarity selections applied. **Aligned sum** generates a new measurement with the summed results of the alignment settings.

The **Level phase at cursor** button calculates the delay required to bring the first measurement phase trace to be approximately horizontal around the frequency of the cursor in the phase trace graph. Both measurements have the same delay applied so the relative delays required to align them are not affected. Shifting the phase traces in that way makes it easier to align them. The delay is only used to update the phase traces, it is **not** applied when using **Aligned copy** or **Aligned sum**. **Undo levelling** removes the levelling delay from both measurements.

The **Align phase slopes at cursor** button calculates the measurement 2 delay required to give both phase traces approximately the same slope around the frequency of the cursor in the phase trace graph, so they have approximately the same group delay. Note that the phase traces may still be far apart, particularly if one measurement is out of phase with the other. In that case inverting one of the measurements would bring the phase traces closer together.

The **Align phase at cursor** button calculates the measurement 2 delay required to bring both phase traces as close as possible to each other around the frequency of the cursor in the phase trace graph. Note that a better overall alignment might be achieved by inverting one of the measurements and then aligning phase, and/or by manually adjusting the delay.

Other Controls

The **Trace options** button brings up a dialog that allows the line type of the graph traces to be changed.

If **Show points when zoomed in** is selected the individual points that make up the SPL and phase responses are shown on the graph when the zoom level is high enough for them to be distinguished (which may only be over part of the plot)

If **Show frequency bands** is selected the audio frequency bands are shown in a stripe above the graph. The bands are:

- Infrasound: below 20 Hz
- Sub bass: 20 to 60 Hz
- Bass: 60 to 250 Hz
- Low midrange: 250 to 500 Hz
- Midrange: 500 Hz to 2 kHz
- Upper midrange: 2 to 4 kHz
- Presence: 4 to 6 kHz
- Brilliance: 6 to 20 kHz

If **Show modal frequencies** is selected the theoretical modal frequencies for the room dimensions entered in the Modal Analysis section of the EQ Window for the currently selected measurement are plotted at the bottom of the graph.

When impedance is being plotted the axis has a span from 0 to 1 kohm. If a larger impedance range is required the axis can be switched to logarithmic with a range up to 1 Mohm using the **Use a log axis for impedance** check box. If selected a log axis will be used wherever impedance is plotted.

Help Index

Distortion Graph
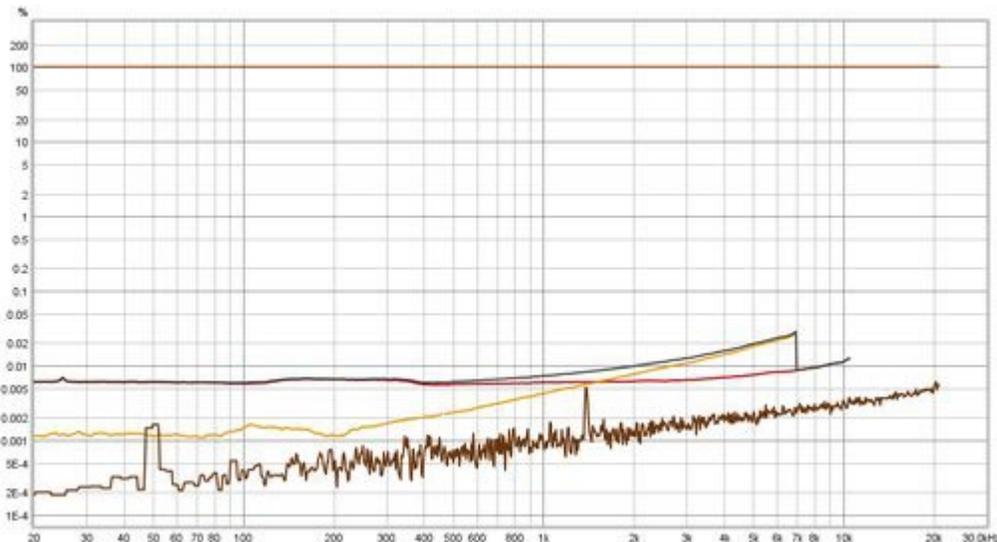
The Distortion graph shows the measurement's fundamental (the linear part of its response), its harmonic distortion components up to the ninth harmonic, Total Harmonic Distortion (THD) and the level of the noise floor, which is captured before the measurement starts.

The plots are derived either from analysis of the impulse response or from stepped sine measurements. Impulse responses measured using logarithmic sweeps separate distortion from the linear part of the system response. The distortion components appear at negative times, behind the main impulse. Analysing the frequency content of these components allows plots of distortion harmonics to be generated. The longer the sweep, the better the distortion components are separated from each other. When measuring a system with high distortion levels use a long sweep setting (e.g. 1M or higher), at shorter sweep lengths the harmonics may affect each other giving misleading results. A spot check can be made at frequencies of interest using the RTA and the signal generator. If discrepancies are observed consider making a stepped sine measurement instead. The noise floor of log sweep distortion measurements rises with frequency. For the lowest noise floor with sweep measurements use multiple sweeps, but note that requires the input and output to be on the same device for reliable results.

Although much, much slower than a log sweep the stepped sine measurement can measure low distortion levels much more accurately than a sweep, particularly at high frequencies and for higher harmonics. Stepped sine distortion measurements show distortion components up to the ninth harmonic, THD and the noise floor, in the same way as the sweep-derived results, but also include THD+N (total harmonic distortion plus noise and non-harmonic distortion) and N (noise and non-harmonic distortion) alone. Note that the noise floor plot shows the spectral content of the noise measured with no signal playing. The 'Noise' in the N and THD+N shows the **summed** level of all non-harmonic distortions and noise across the frequency span for each test frequency. It consequently sits much higher than the noise floor plot. For stepped level measurements the X axis can be dB SPL, dBFS, dBu, dBV, dBW, V or W showing either the generator or input level.

The harmonic plots can only be generated for frequencies within the bandwidth of the measurement. For example, if a measurement is made to 20 kHz, the second harmonic plot can only be generated to 10 kHz, as the 2nd harmonic of 10k Hz is 20 kHz. Similarly the third harmonic plot can only be generated to 6.67 kHz (20/3). The upper limit for distortion plots is 95% of the Nyquist frequency (which is half the sample rate). For example, at 44.1 kHz sampling the upper limit is 0.95*44.1/2 = 20.95 kHz.

The lower frequency limit for distortion plots is 10 Hz or the measurement start frequency, whichever is higher. 10 Hz is the lower limit of the logarithmic part of the measurement sweep. Start frequencies lower

than that use an initial linear swept portion (to avoid spending an excessive proportion of the sweep duration time at very low frequencies) which means that region cannot be used to generate distortion data. Note that sweeps with start frequencies of 20 Hz or higher may show rising distortion at low frequencies as a result of the time domain ringing inherent in band limited impulse responses overlapping the distortion harmonics. To avoid that use start frequencies below 20 Hz where possible and longer sweeps.

Total Harmonic Distortion is generated from the available harmonics. There is a control to select the highest harmonic to use in the THD calculation, the trace name in the graph legend shows which harmonics are included. At higher frequencies the THD plot will incorporate fewer harmonics, according to which are available.

If the Analysis Preference Apply cal files to distortion is selected the results will include corrections for the cal file responses (as is the case for the RTA distortion figures). Applying the cal files provides more accurate results in regions where the fundamental or harmonics are affected by interface roll-offs but boosts the noise floor in those regions. This should be borne in mind when viewing the results. If large cal file corrections are required make sure the Analysis Preference Limit cal data boost to 20 dB is not selected. Note that any subsequent changes to the cl files will **NOT** update the distortion results, they are generated from the cal files that were in use at the time the measurement was made.

The fundamental and harmonic plots derived from sweep measurements are smoothed to 1/24 octave. This cannot be adjusted. The distortion data can be exported to a text file using File → Export → Distortion data as text.

Distortion Controls

The control panel for the graph has these controls:

By default the plot shows the absolute levels of the fundamental and the harmonics. The Y axis can be set to show dB SPL, dBFS, dBu, dBV, dBW, volts, watts, dBr or percent. Values for dBW and watts are derived from the measured voltages using the reference impedance value in the RTA Appearance settings. If the Y axis is set to percent or dBr the harmonics are divided by the fundamental to show their relative level and the fundamental appears as a flat line at 0 dBr or 100 %. The legend value for the fundamental will continue to show its absolute level, the readings for the harmonics and THD will depend on the **Distortion Figures** setting. Normalising the plot will cause the distortion traces to rise at high frequencies if the response of the system being measured rolls off (as is usually the case). This is exaggerated if **Use harmonic frequency as ref** is selected (see next section). The boosting due to low fundamental level can be controlled by selecting **Limit norm. to 30 dB below peak**, this sets a lower limit on the fundamental that is 30 dB below the peak level of the fundamental - for example, if the peak of the fundamental were 95 dB the minimum level used for normalising would be 65 dB.

The harmonic and THD plots in normalised mode use the level at the fundamental for each frequency as

their reference by default - for example, the distortion figures for each harmonic at 1 kHz will depend on the level of the fundamental at 1 kHz. If **Use harmonic frequency as ref** is selected the reference will be the frequency of the harmonic - for example, at 1 kHz the 2nd harmonic figure will depend on the level of the fundamental at 2 kHz, the 3rd harmonic will depend on the level of the fundamental at 3 kHz and so on. This follows a recommendation made by Steve F. Temme in "How to graph distortion measurements" presented at the 94th AES convention in March 1993. If the response of the system being measured is flat this makes no difference to the results, but when the response is not flat (as for most acoustic measurements) it can remove the influence of the loudspeaker's fundamental response from the distortion figures. As an example, suppose the loudspeaker response was flat apart from a 6 dB peak at 2 kHz. 2 kHz is the 2nd harmonic of 1 kHz, so the 2nd harmonic level shown at 1 kHz will be increased by 6 dB due to the boost in the fundamental when using the excitation frequency as the reference. Similarly the 3rd harmonic level at 667 Hz (2/3 kHz) will be boosted by 6 dB. If the harmonic frequency were used as the reference the distortion figures would not show this boost. Using the harmonic frequency as the reference also provides a more meaningful view of distortion at frequencies below the LF roll-off of the system as otherwise the distortion levels are boosted as the level of the fundamental drops. Note that this option will not affect the traces when the plot is not normalised, but will still affect the values in the legend if the distortion figures are set to read in percent or in dB relative to the fundamental.

The **Mask harmonics below noise floor** control greys out harmonics that lie below the noise floor (at the frequency of the harmonic). The THD trace is also greyed out if all the harmonics that contribute to it are below the noise floor. In applying the masking for sweep measurements REW takes into account the (limited) ability of the sweep to discriminate harmonics that lie a below the noise floor at the harmonic frequency, which varies from 3 dB for the second harmonic to almost 10 dB for the ninth.

Plot harmonics at the harmonic frequency changes the way the harmonic orders for sweep measurement distortion are plotted - rather than plotting the distortion at the frequency of the fundamental it is plotted at the frequency of the harmonic, so (for example) the 2nd order harmonic distortion level for 1 kHz would be plotted at 2 kHz, where the distortion occurs. This makes it easier to correlate harmonic levels with the level of the fundamental at the harmonic frequency and also helps distinguish the effects of external noise, since noise affects all harmonics at the frequency at which it occurs. The THD trace is not affected by this control, it will continue to show the THD value corresponding to the harmonic levels at their fundamental frequency. When plotting harmonics at the harmonic frequency the harmonic frequency is used as the reference, regardless of the state of **Use harmonic frequency as ref**.

The **Distortion figures** control selects the units that are used for the harmonic distortion levels displayed on the graph legend. The choices are **As Y axis**, which shows the level using the selected Y axis unit; **dB relative**, which shows how many dB the harmonic is below the fundamental; and **Percent**, which shows the harmonic level as a percentage of the fundamental. The frequency at which the fundamental level is taken depends on the setting of **Use harmonic frequency as ref** (see above).

The **Highest harmonic to show** control allows the higher harmonics to be hidden if they are not of interest. For example, if **Highest Harmonic to show** were set to 3 only the second and third harmonic traces would appear on the graph and in the graph legend. The measured harmonics used to calculate THD are controlled by **Highest harmonic in THD**.

The **Highest harmonic in THD** control allows the higher harmonics to be excluded when calculating THD. This may be desired if they are in the noise floor and so adding noise to the THD calculation.

The **Trace options** button brings up a dialog that allows the colour and line type of the graph traces to be changed. If a change is made it will be used for all measurements shown on this graph. Traces can also be hidden, which will remove them from the graph and from the graph legend.

Distortion Examples - soundcard measurements

Here is a distortion plot generated from a loopback measurement of a soundcard, produced at a high sweep level (-4 dBFS, which resulted in 2 dB of headroom at the gain settings used). The readings in the legend are with the cursor at 1 kHz.
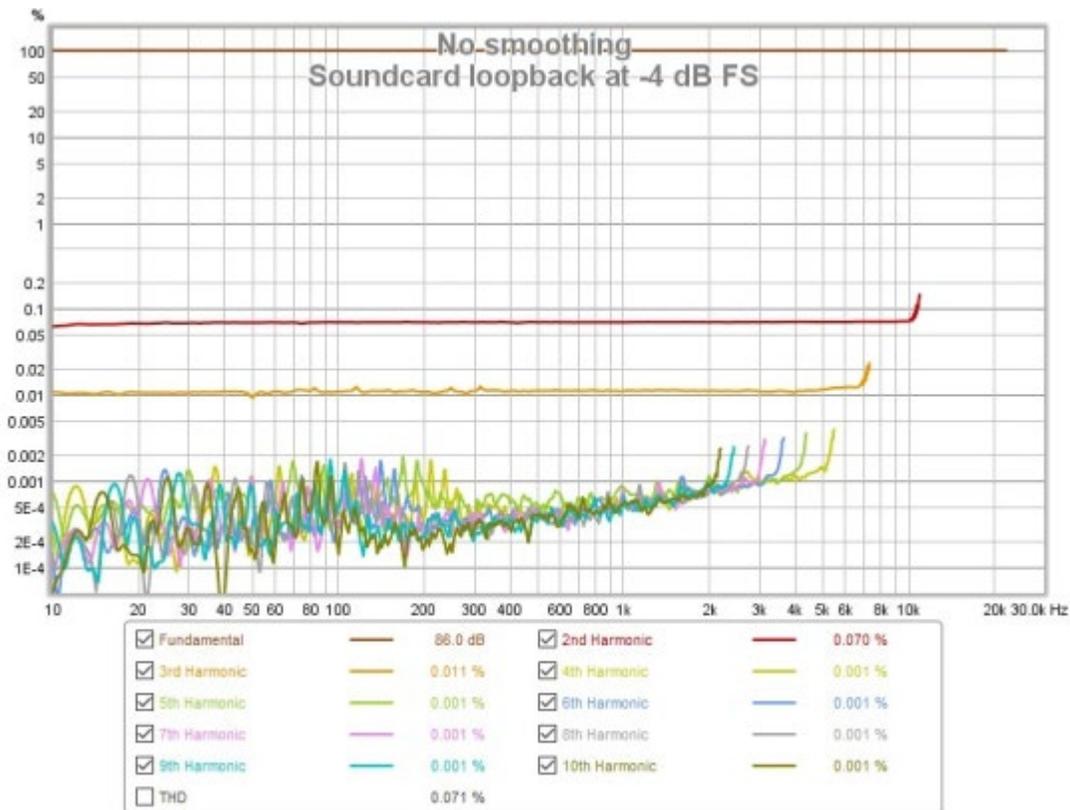
The THD trace has been omitted, as it overlays the 2nd harmonic trace (in red) which is the dominant component, 0.07%. The 3rd harmonic (in orange) is much lower at around 0.01%, whilst the higher harmonics are largely within the noise floor.

This is the impulse response for that measurement, the distortion peaks are to the left of the main peak. The first peak to the left is the 2nd harmonic, the next is the third harmonic and so on.
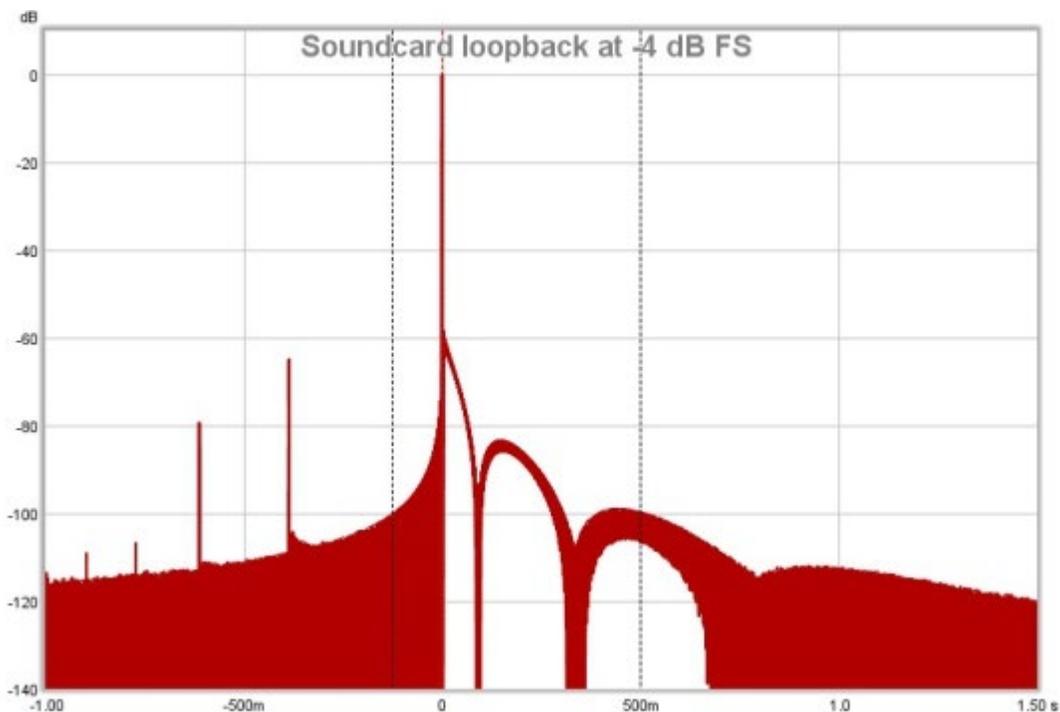
The next plot is from a room measurement. The 2nd (red), 3rd (orange) and 4th (yellow) harmonic traces are shown, along with the THD (black). Higher harmonics were within the noise floor. The measurement shows a sharp rise in 3rd harmonic distortion at 94 Hz, and a dramatic rise in all distortion components from about 2 kHz upwards. Further measurements at differing signal levels established that this distortion was being introduce by the SPL meter used for the measurement.
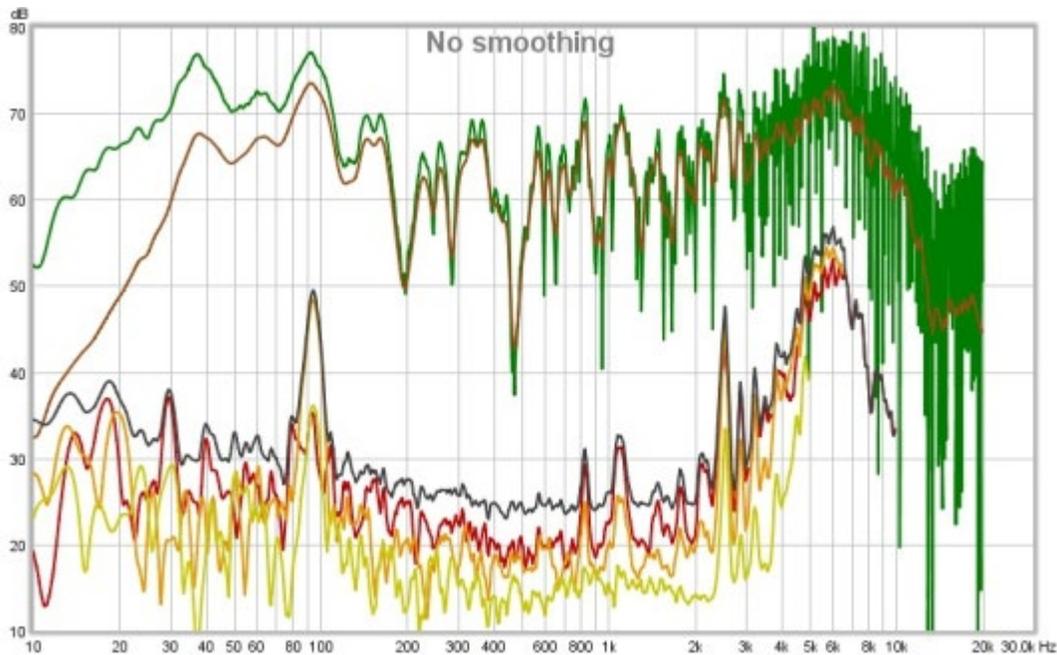

This is the impulse response for the in-room measurement, the distortion peaks are clearly visible to the left of the main peak.
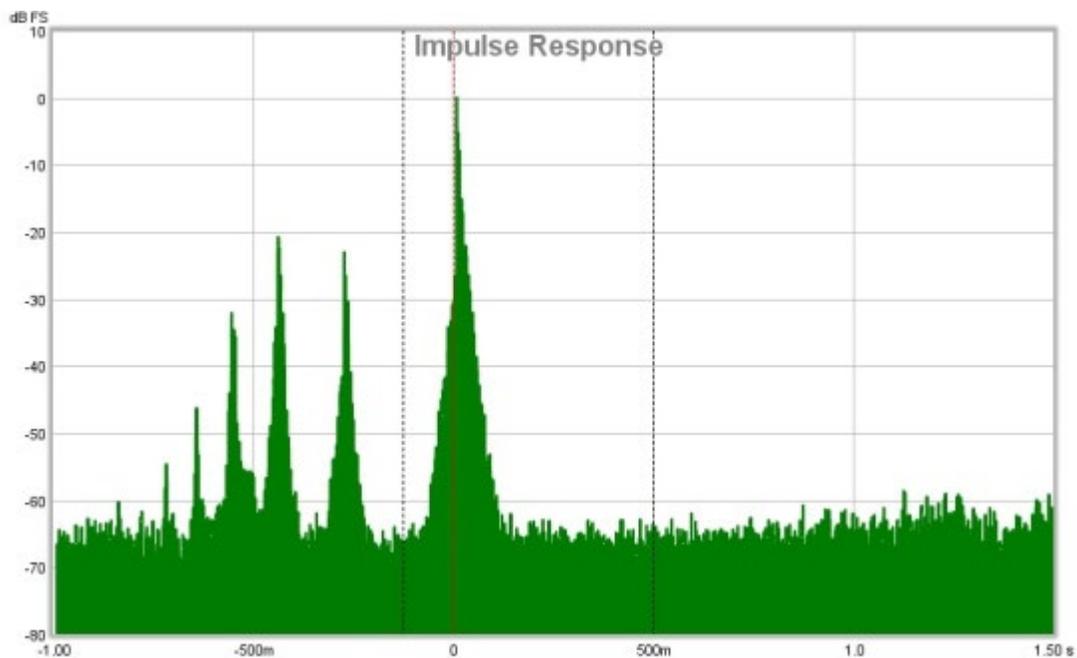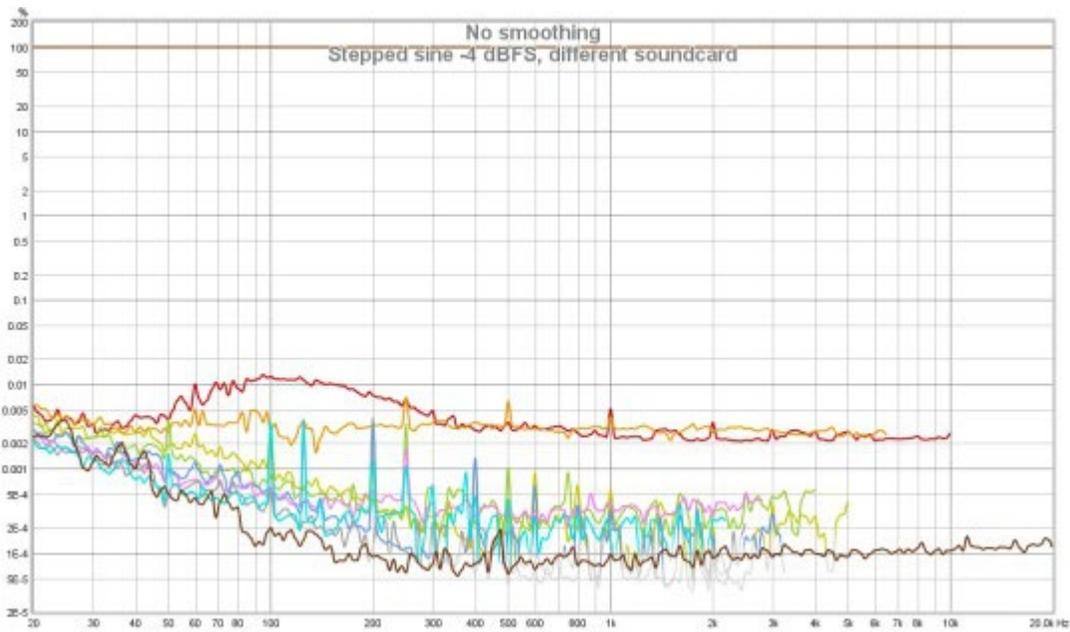

The plots below show a different soundcard loopback measurement at -4 dBFS measured with stepped sine (64k FFT, 24 ppo) and with a 1M log sweep. Note the rise in the noise floor with frequency when measuring with the sweep, all but the 2nd and 3rd harmonic lie below the noise floor. In addition, the 3rd harmonic measures higher with the sweep than with stepped sine at frequencies above 100 Hz or so. With the stepped sine measurement the contributions of the 4th, 5th and 7th harmonics are also visible above the noise floor

(dark brown trace). Stepped sine is a much more accurate method when measuring low distortion levels.
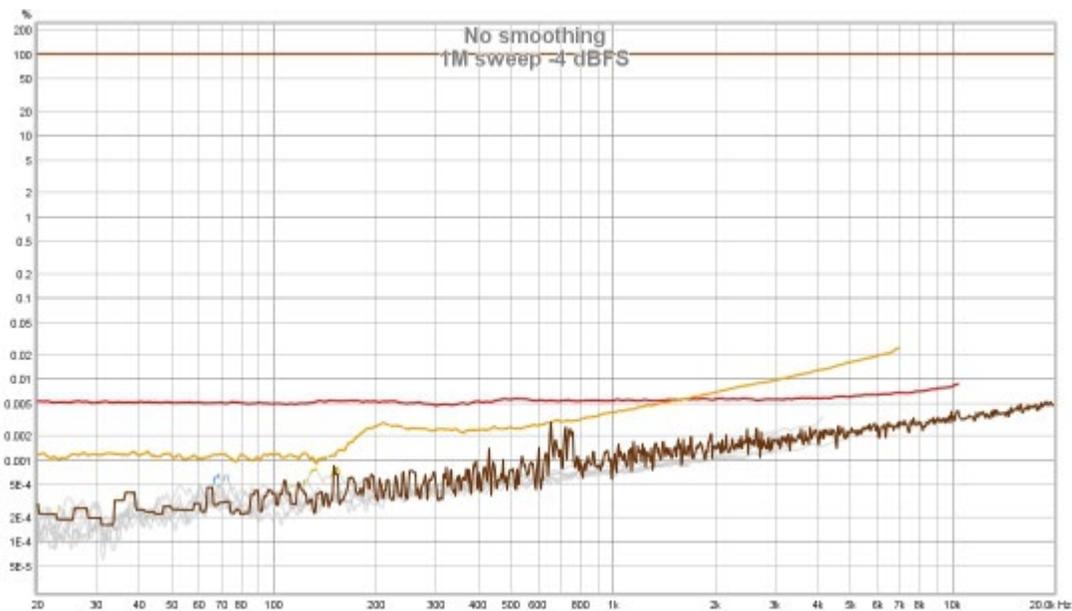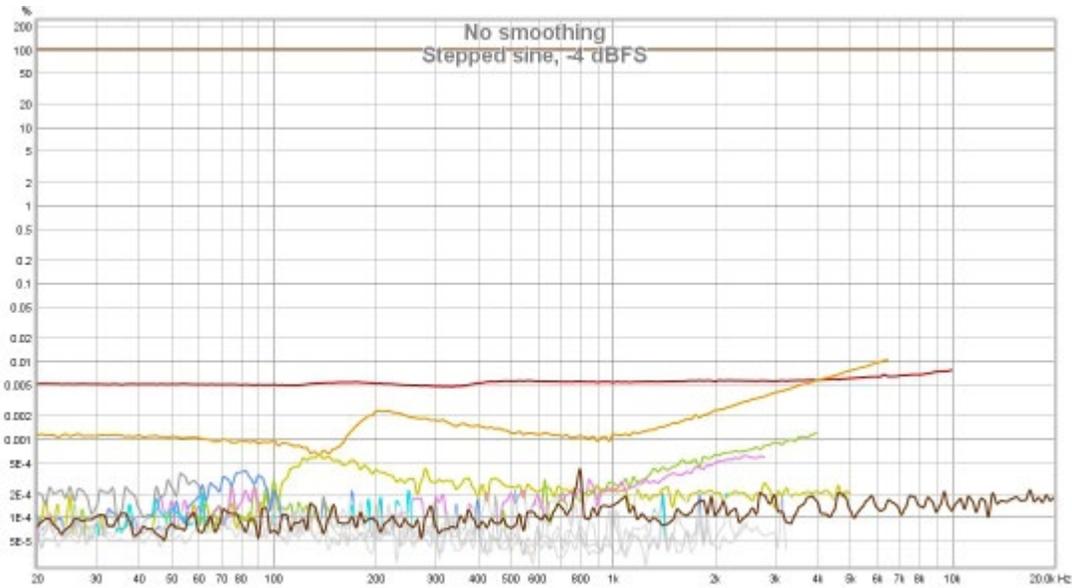
Further insight into the distortion behaviour can be obtained by looking at waterfall or spectrogram views of the spectrum data captured at each test frequency. Note that these can only be generated for stepped sine measurements that had the option to **Capture spectrum data at each frequency** selected. Here is a spectrogram of the stepped sine measurement above.

Here is its waterfall plot.

Here is a stepped sine measurement of another soundcard, superficially looking to have similar performance to the one above.

However, its spectrogram makes clear that it has far worse behaviour with significant intermodulation products with the 1 kHz harmonics of the USB frame rate.
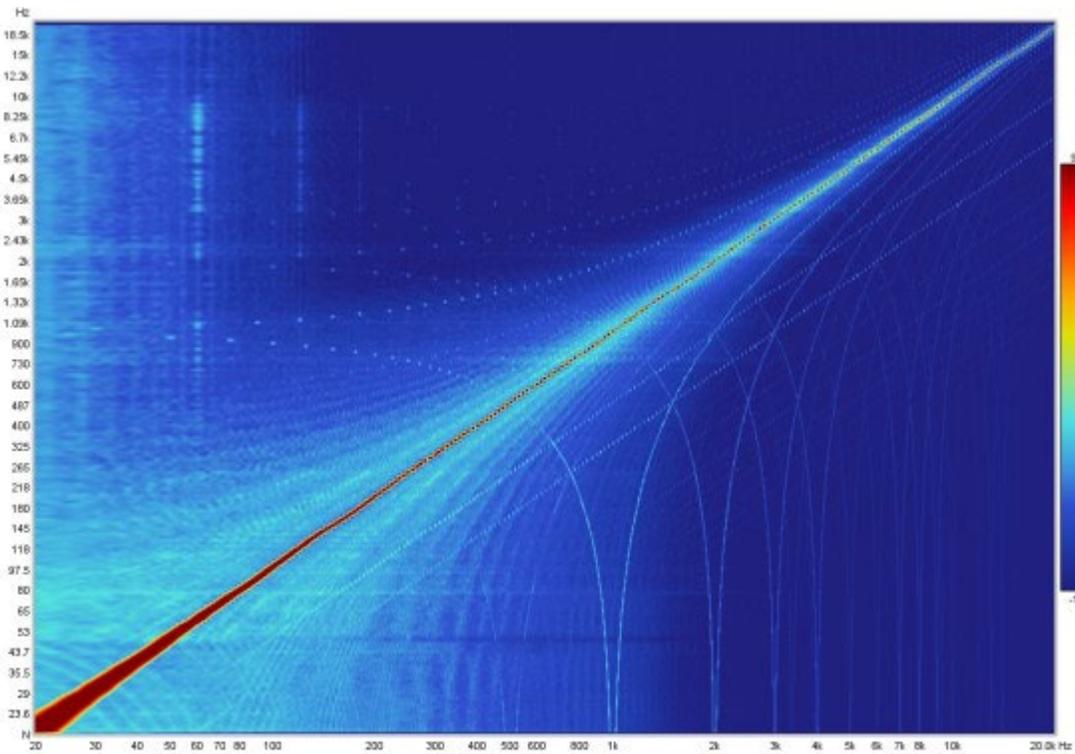

Its waterfall plot is similarly revealing.
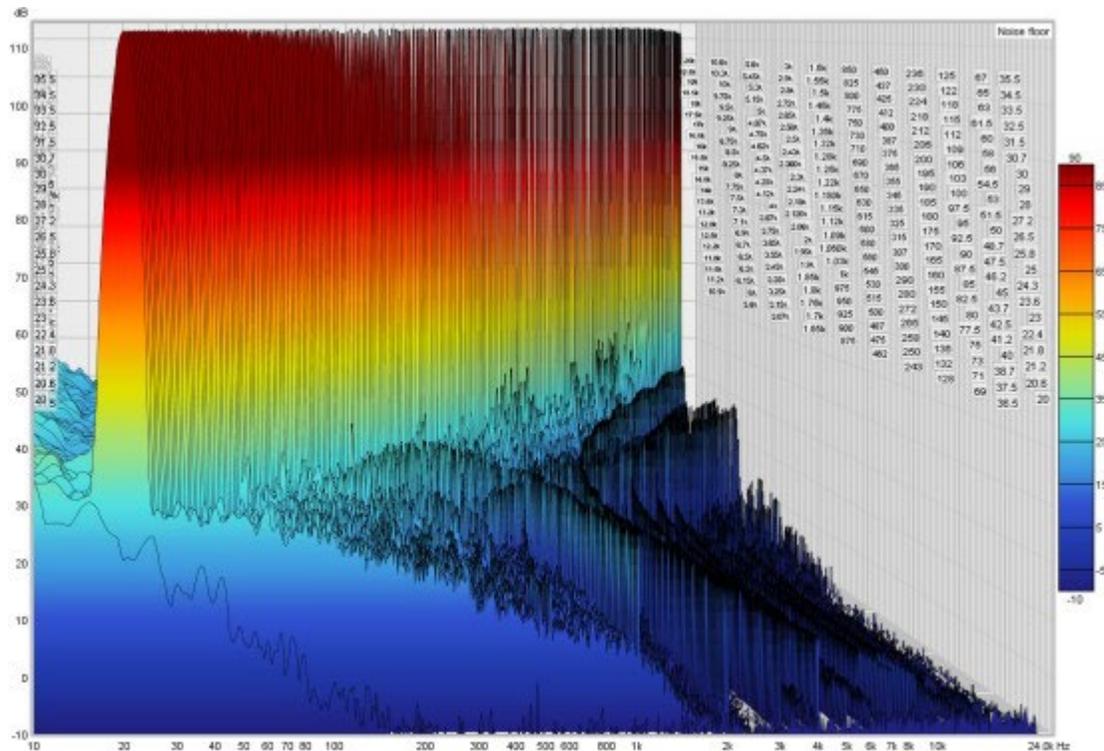
Distortion Examples - speaker measurements

Speaker measurements, particularly those made in-room rather than with the benefit of an anechoic chamber, present additional challenges for interpretation. Distortion results for electronic components are typically shown as percentages, values relative to the fundamental. Plotting these percentages for speaker measurements can make the results challenging to interpret due to the often highly irregular frequency response of the speaker. Using the harmonic frequency as reference (see distortion controls above) helps somewhat, but is by no means a complete solution. My personal preference is to view the distortion results without normalisation, which helps make the effect of the response irregularities apparent and also emphasises the acoustic level of the distortion components, which can be helpful in deciding where they are likely to be audible and where not. When viewing the distortion graph as SPL levels the trace legend can still be configured to read out distortion percentages using the **Distortion figures** control.

Here is a measurement of a small desktop speaker (Adam Artist 3). The mic (UMIK-1) was just 15 cm from the speaker, hence the relatively high SPL measured - the levels correspond to about 85 dB at the listening position, about 70 cm from the speakers.
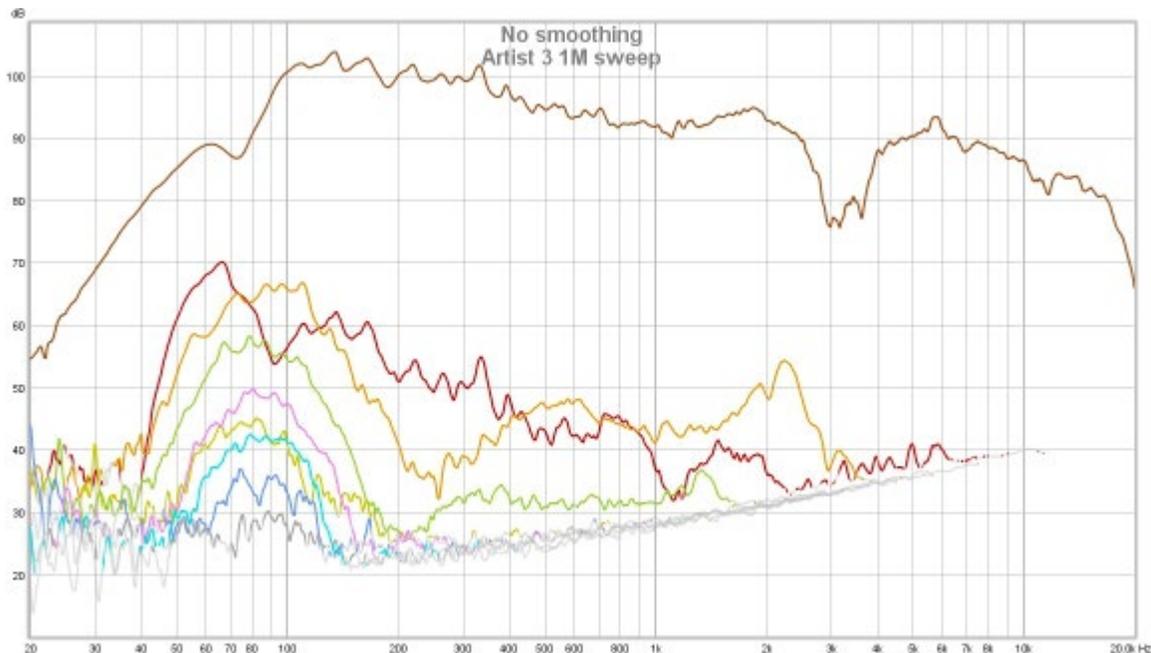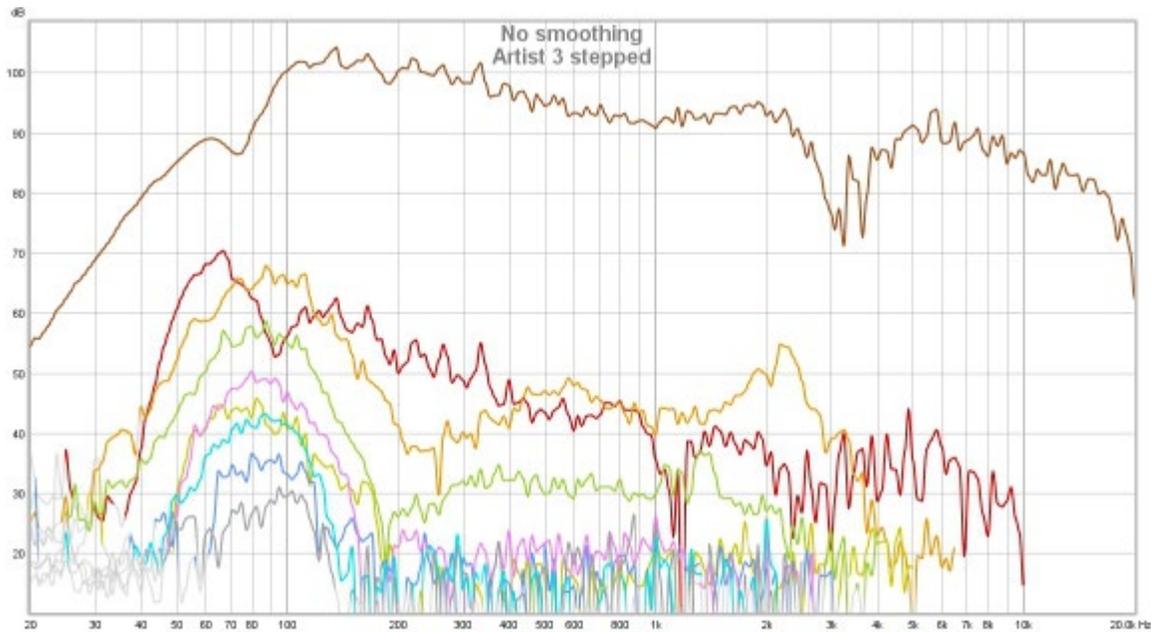
The hump of odd harmonics centred around 80 Hz or so looks to be mild clipping. Above 3 kHz, where the speaker's ribbon tweeter comes into play, distortion levels drop below the noise floor for the sweep though the stepped sine tracks them comfortably. As this is a near field measurement room effects are reduced, but the frequency response variation can still make a normalised view more difficult to fathom. Below is the response normalised, without using harmonic frequency as ref. The noise floor (dark brown) is included to help illustrate the effect normalisation has on boosting the noise floor where the response dips.
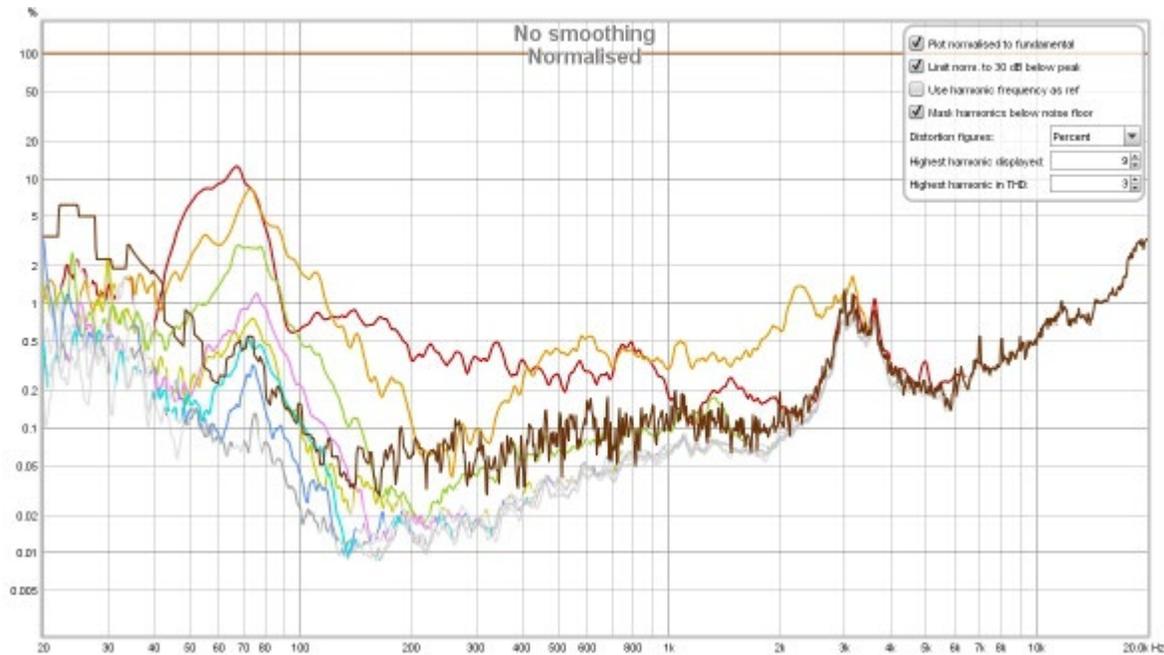
It is apparent that the peak in second harmonic distortion around 3 kHz is simply the noise floor being boosted and so can be ignored - though that would not be so easily identified without the noise floor trace to guide us. However, distortion at low frequencies is now shown much higher than it should be due to the low frequency roll off of the speaker's response meaning it is being normalised against a much lower level. That 10% 2nd harmonic might give cause for concern until it is recognised as an artefact of the normalisation.

The normalisation artefacts at low frequencies can be ameliorated by using the level of the fundamental at the harmonic frequency as the normalisation reference. That gives us the plot below, the low frequency behaviour is now more accurately portrayed. However, new peaks have appeared in the 2nd and 3rd harmonics above 1kHz, consequences of the dip in the response centred around 3 kHz. That dip is probably due to the crossover between mid/bass and tweeter integrating poorly at such a near field measurement position. It is unlikely that the distortion peaks are real. The view without normalisation is least susceptible to misinterpretation.
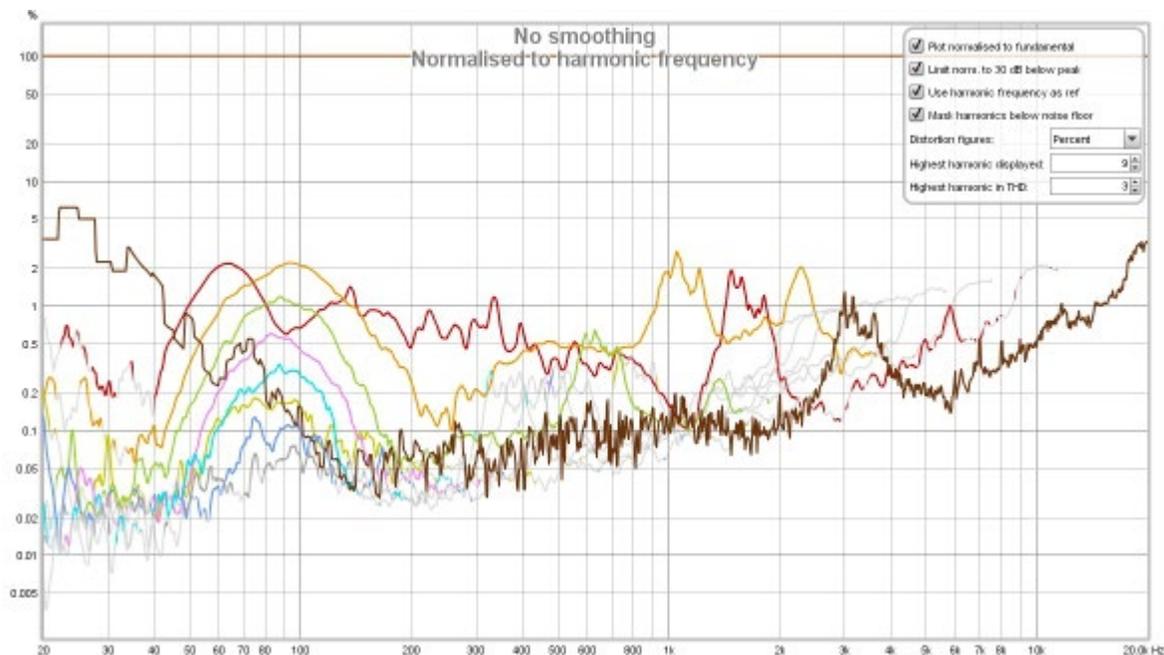
[Help Index](#)

Impulse Graph

The Impulse graph shows the impulse response for the current measurement. It can also show the left and right windows and the effect of the windows on the data that is used to calculate the frequency response; a minimum phase impulse; the impulse response envelope (ETC) and the step response.
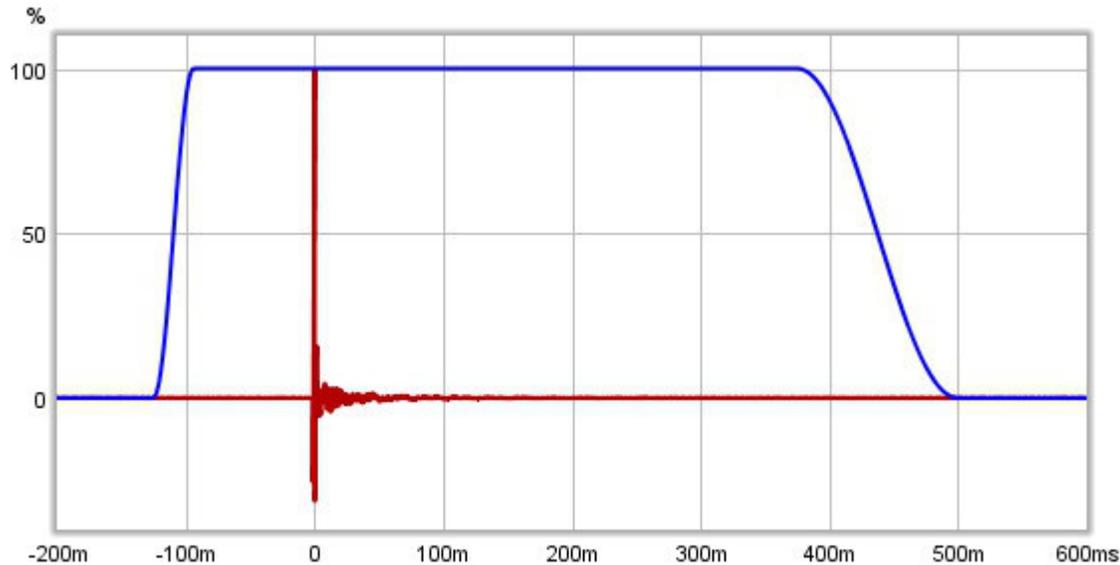

The Y axis used for the impulse response can be selected as % FS or dBFS (FS = Full Scale) via a control in the top left corner which appears when the mouse cursor is inside the graph area. The dBFS scale is equivalent to a "log squared" view of the impulse.
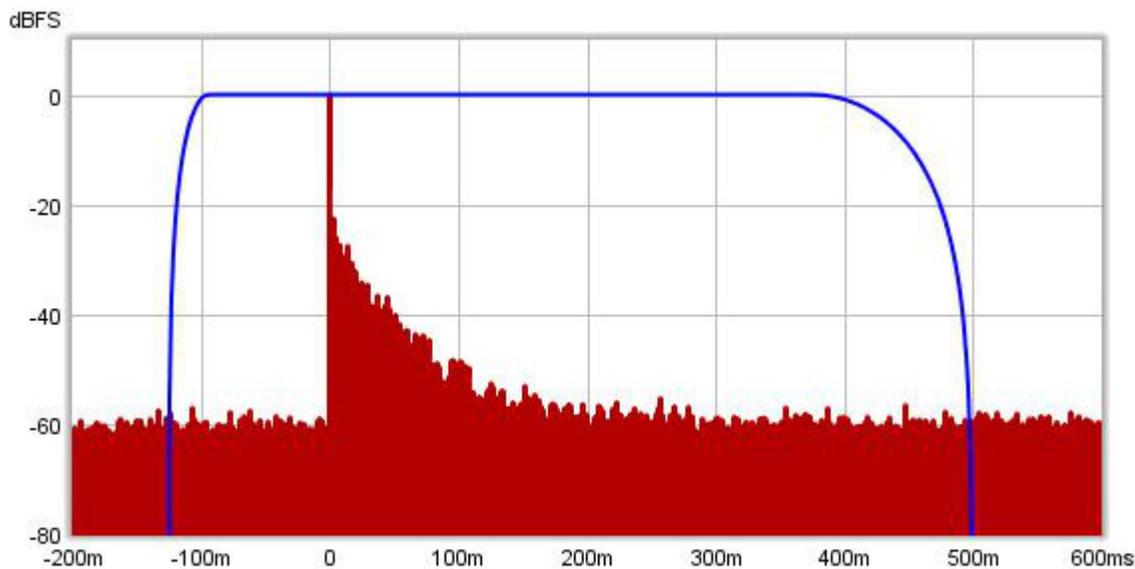
Indicators above the graph show the extents of the impulse response windows and the window reference position. The window settings may be changed by clicking and dragging on the indicators, the new values are applied when the mouse button is released. While the window settings are being changed the region

outside the new area is shown shaded until the settings are applied. When the window indicators are dragged a preview of the effect the new setting will have on the SPL response is shown below the impulse response. It is best to set the Y axis to dB to adjust the windows as it is then much easier to see where the response has decayed into the noise.
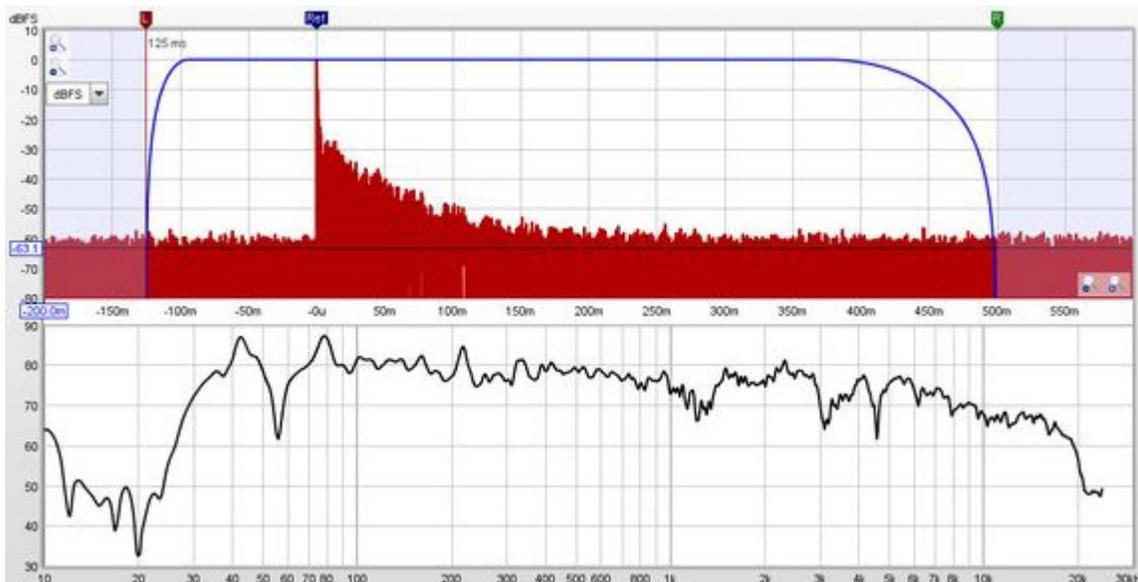
After each measurement the left window width is automatically set up. For full range measurements (and down to end frequencies of 1kHz) the width is 125ms, below that it increases to allow for pre-ringing effects of using a limited sweep range. To change the window settings for a measurement click the **IR Windows** button:


IR Windows

The impulse response is that of the whole system, including the mic/meter and the soundcard. **The mic/meter and soundcard calibration files are only applied when calculating the frequency response, their effect is not included in the impulse response.**

Minimum Phase Impulse

If the **Generate minimum phase** control has been used to produce a minimum phase version of the current measurement's magnitude response a minimum phase impulse trace is activated, showing the impulse response the minimum phase system would have. Note that it is necessary to make full range measurements if the minimum phase response is to be generated as a good result relies on measuring beyond the bandwidth of the system being measured.

Impulse Response Envelope

The envelope of the impulse, also called the energy-time curve or ETC, is useful to identify reflections and see the overall shape of the impulse response. The plot below shows the envelope, the spikes after the initial peak are due to reflections from room surfaces, the first spike occurs 3.25ms after the initial peak indicating that the sound travelled an additional 1.11m or 3.7 feet to reach the microphone.
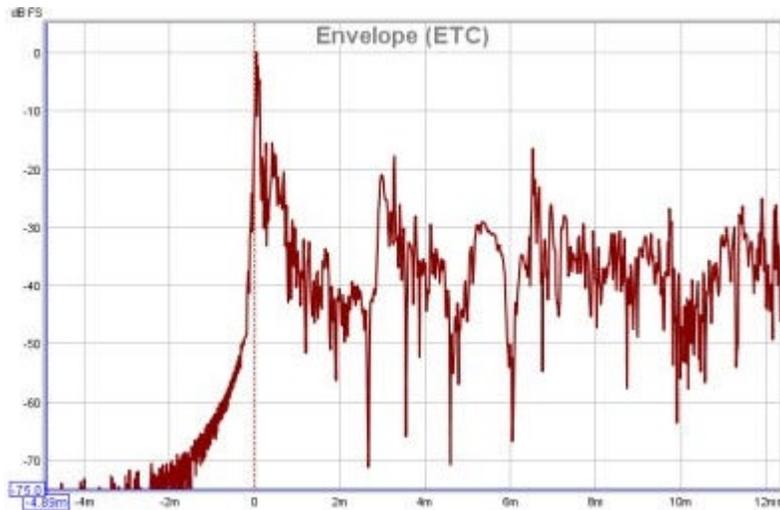
Step Response

The step response shows the output which would result if the input signal jumped to a fixed level and stayed there. It is the integral of the windowed impulse response. If there is an offset in the measurement input chain the step response will show an overall rise or fall as time progresses, rather than tending back to zero.
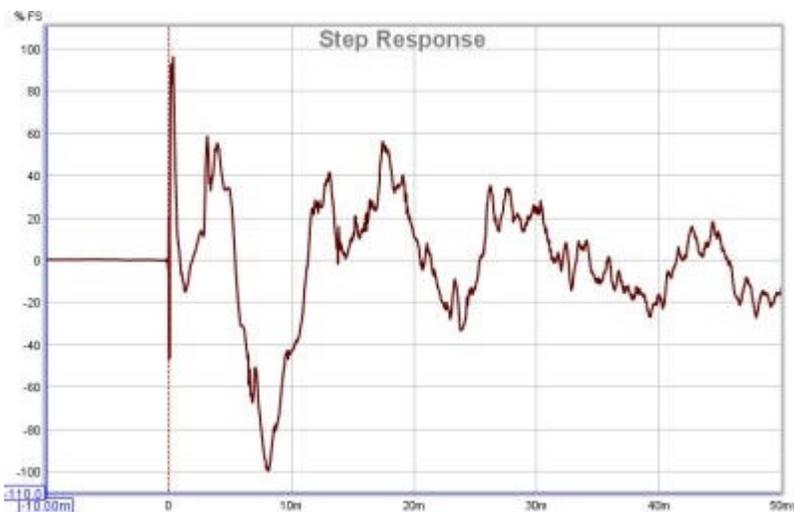
Distortion Components

A property of the log sweep analysis method is that the various harmonic distortion components appear as additional impulses at negative time, with decreasing spacing as the distortion order increases. For example, this plot shows spikes from distortion components up to the 8th harmonic on a laptop soundcard loopback measurement:
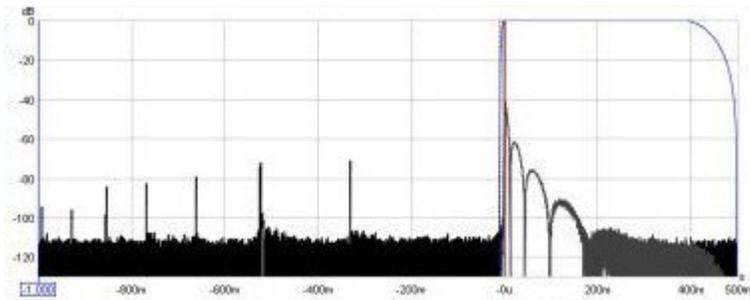
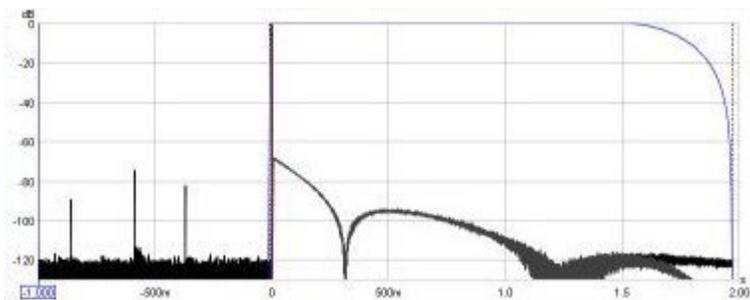Here is a similar measurement for an external USB soundcard, it is a 44.1k card rather than 48k, which limits us to the 6th harmonic in the 1s pre-impulse period - however, only the 2nd, 3rd and 5th harmonic peaks are evident, the 4th harmonic peak is barely visible above the noise floor (which is about 10dB lower than the laptop card). The extended lobes after the impulse are due to the card's much lower -3dB frequency, 1.0Hz versus 22.1Hz (note that the right side of the time axis is 2.0s in this plot compared to 0.5s in the previous plot):

Impulse Controls

The control panel for the Impulse graph has these controls:

The **Set t=0 at IR start** button will align the zero position to the time at which the impulse starts (emerges from the noise floor). It can be activated using the Alt+y keyboard shortcut when the controls panel is visible.

The **Set t=0 at cursor** button will align the zero position to the current cursor position. It can be activated using the Alt+z keyboard shortcut when the controls panel is visible.

Offset t=0 allows the position of time zero in the impulse response to be altered, with a live preview of the effect the offset will have on phase. The measurement is not changed unless either the **Apply** or **Apply & close** button is pressed. Adjustment of the offset can be made by entering a value or by using the up and down arrow keys after clicking on the offset spinner. The cumulative shift that has been applied to the

impulse response is shown at the top of the dialog. That includes any shifts applied to the original measured IR to align it to the peak or to the timing reference, if used. The offset can be set to the value required to set the cumulative shift to zero by pressing the **Zero cumulative shift** button. If a timing reference was used the System Delay figure (which can be viewed in the measurement Info panel) is shifted by the same amount as the zero time.
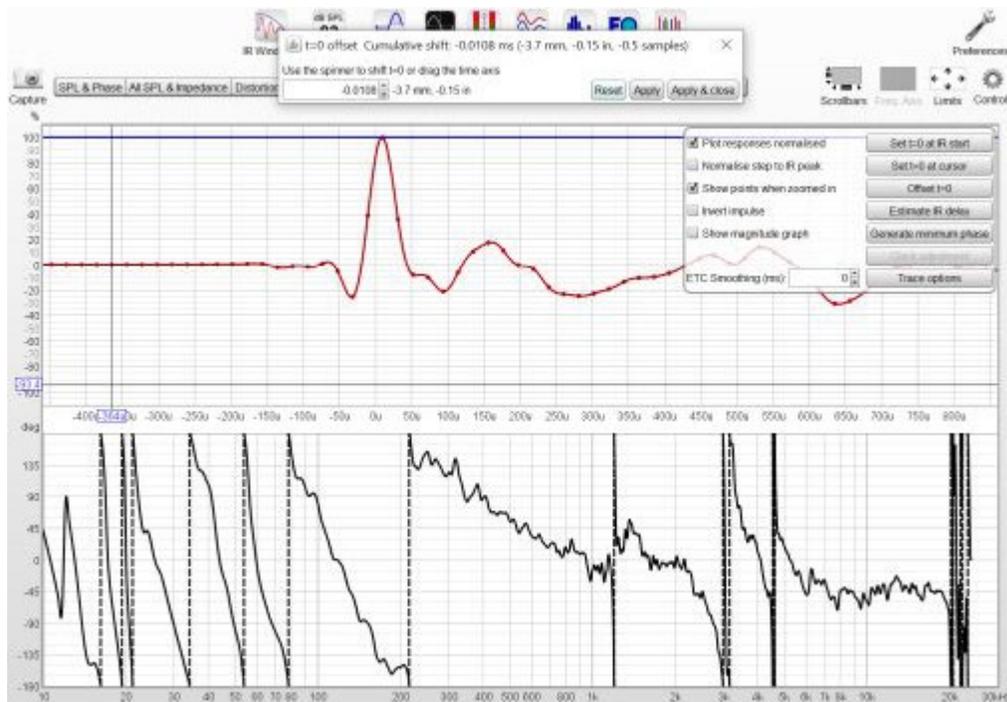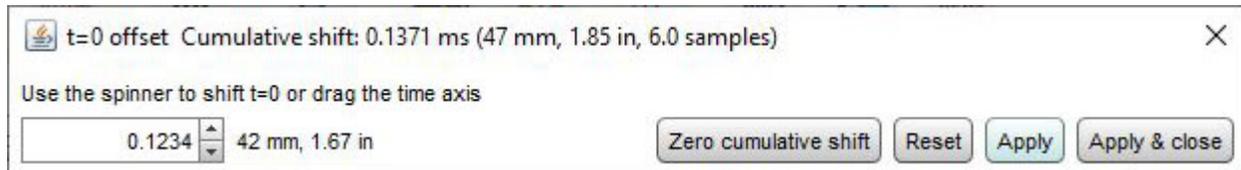
Estimate IR delay calculates an estimate of the time delay in the measurement by comparing it with a minimum phase version. The delay it calculates can be removed from the impulse response by pressing the **Shift IR** button on the panel shown after the delay is calculated and can additionally be applied as a timing offset for subsequent measurements.
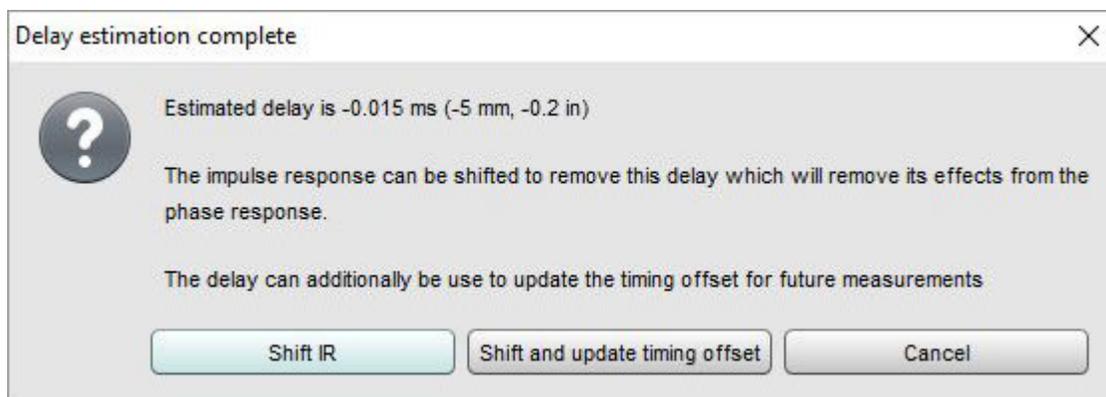
Generate minimum phase will produce a minimum phase version of the measurement using the current IR window settings. The minimum phase impulse then shows the response of a system having the same frequency response as the measurement but with the lowest phase shift such a system could have. Note

that it is best to make full range measurements if the minimum phase response is to be generated as a good result relies on measuring beyond the bandwidth of the system being measured. This control also activates minimum and excess phase and group delay traces on the SPL & Phase and GD graphs respectively.

Note that the IR window settings are important as the minimum phase response is derived from the frequency (magnitude) response of the measurement, which in turn is affected by the IR window settings. If the window settings are subsequently changed **Generate minimum phase** should be used again to reflect the new settings. Note also that the shape of the left side window (the window applied before the peak) affects the minimum phase result, a rectangular window will produce a response with lower phase shift than, for example, a Hann window.

If the system being measured was inherently minimum phase (as most crossovers are, for example) the minimum phase response is the same as removing any time delay from the measurement. Room measurements are typically not minimum phase except in some regions, mainly at low frequencies. For more about minimum and excess phase and group delay see [Minimum Phase](#).

The **Clock adjustment** control is enabled for log sweep measurements made with REW V5.20 beta 2 or later. It is used to compensate for differences between the clock rates of the output device and input device. When the output and input device are different their clocks may differ sufficiently to distort the shape of the impulse response, particularly when using longer sweeps. The clock rate difference affects the phase response (magnitude is not affected). The adjustment is in units of parts per million, ppm. While the clock rate is being adjusted a live preview shows the effect the adjustment will have on the impulse response and phase. The measurement is not changed unless either the **Apply** or **Apply & close** button is pressed. Fine adjustment can be made with the left and right arrow keys after clicking on the slider knob. The control has a range of -500 to +500 ppm, but greater adjustments can be made by using the **Apply** button then applying further changes as required. The cumulative clock adjustment is shown at the top of the dialog and in the measurement Info dialog. Note that clock adjustment is a computationally intensive task so the display updates are not immediate.
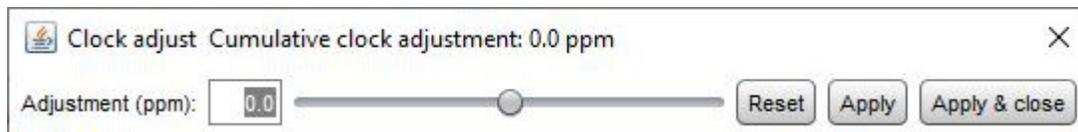
The **Trace options** button brings up a dialog that allows the colour and line type of the graph traces to be changed. If a change is made it will be used for all measurements shown on this graph. Traces can also be hidden, which will remove them from the graph and from the graph legend.

The impulse response may be plotted with or without normalisation to its peak value according to the setting of the **Plot responses normalised** control. When normalised plotting is selected the peak will be at 100% or 0 dBFS.

The step response may be normalised to its own peak value or to the peak value of the impulse response according to the setting of the **Normalise step to IR peak** control.

If **Show points when zoomed in** is selected the individual points that make up the response are shown on the graph when the zoom level is high enough for them to be distinguished.
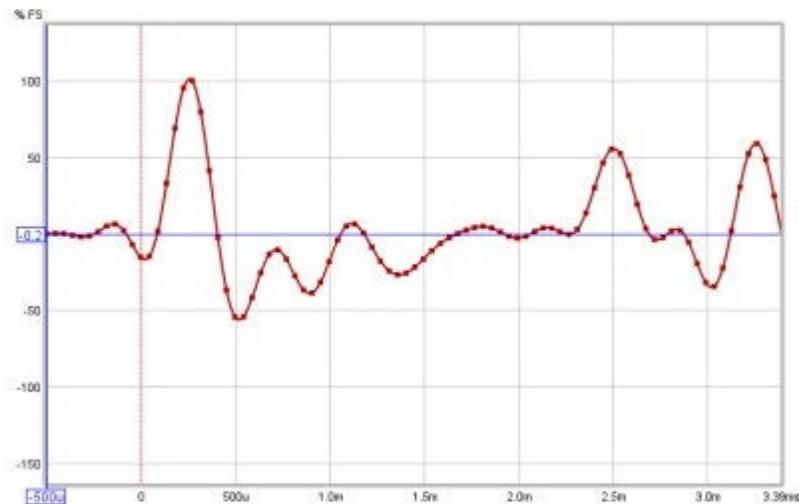
The impulse response may be inverted according to the setting of the **Invert impulse** control. Note that this has no effect on the displayed impulse response when the Y axis is set to dBFS. If the soundcard you are using inverts its inputs that can be corrected using the **Invert** checkbox in the Soundcard Preferences Input Channel controls.

If **Show magnitude graph** is selected the response SPL graph is shown below the impulse graph. The view is the same as that shown when adjusting the IR windows using the indicators at the top of the graph.

ETC Smoothing is used to smooth the envelope (ETC) trace using a moving average filter of the duration specified in the spinner.

Help Index

Filtered IR Graph

The Filtered IR graph allow octave and one-third octave filters to be applied to the measurement. The octave band filters are Butterworth bandpass, with the order selected in the graph controls. The graph is primarily aimed at examining decay behaviour in different frequency bands and analysing the results per ISO 3382. In addition to the filtered impulse response itself this graph includes traces of the impulse response envelope (ETC) and the Schroeder integral, from which RT60 estimates are derived.
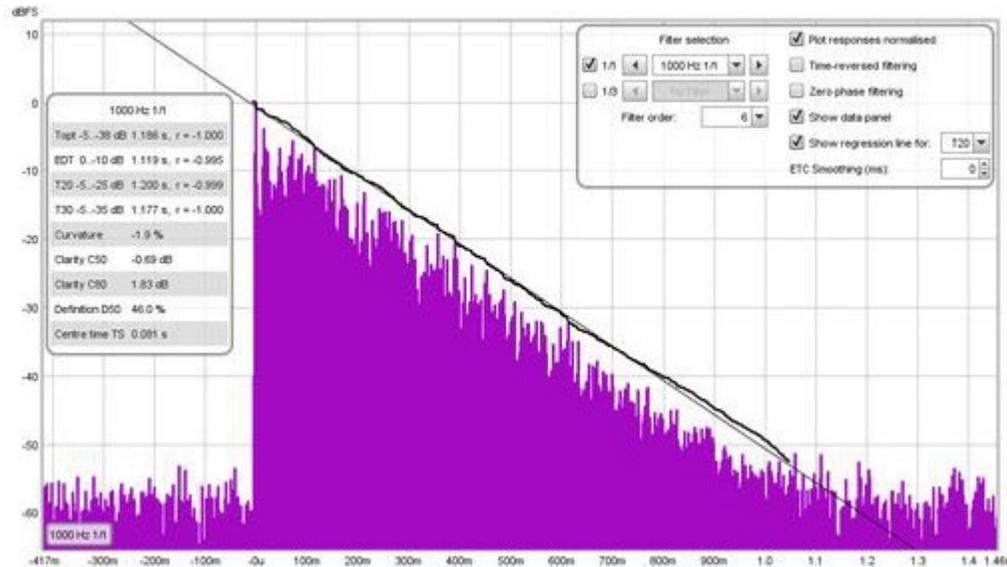
Schroeder Integral

The Schroeder Integral is a curve obtained by backwards integration of the squared impulse response, ideally starting from a point where the response falls into the noise and applying a correction (a starting value for the integral) which assumes the rate at which the Schroeder curve is falling continues for the whole response. REW uses an iterative procedure to estimate the best starting point for the integration, based on "Lundeby's Method" (from the paper by A. Lundeby, T. E. Vigran, H. Bietz, and M. Vorländer, "Uncertainties of Measurements in Room Acoustics," Acustica, vol. 81, pp. 344–355 (1995)). The slope of this curve is used to measure how fast the impulse response is decaying, deriving a figure for "RT60" which is the time it would take sound to decay by 60dB. The curve shown on the graph is for the currently applied filter, if any. When calculating decay data for the octave and one-third octave RT60 results the impulse is first filtered to the corresponding bandwidth and centre frequency before the Schroeder Integral for that band is determined and the various RT60 measures calculated.
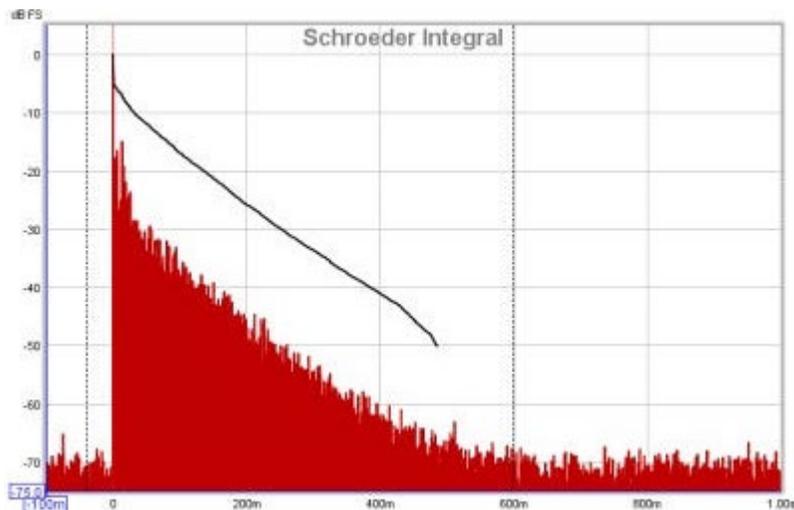
Graph metrics

If a Schroeder Integral range is selected by holding the shift key then pressing and dragging with the left mouse button a regression line for the selected range is displayed, with the corresponding T60 value. The range can be adjusted after selection by dragging the start or end of the area, or the whole range can be moved. The Y axis must be set to dBFS.
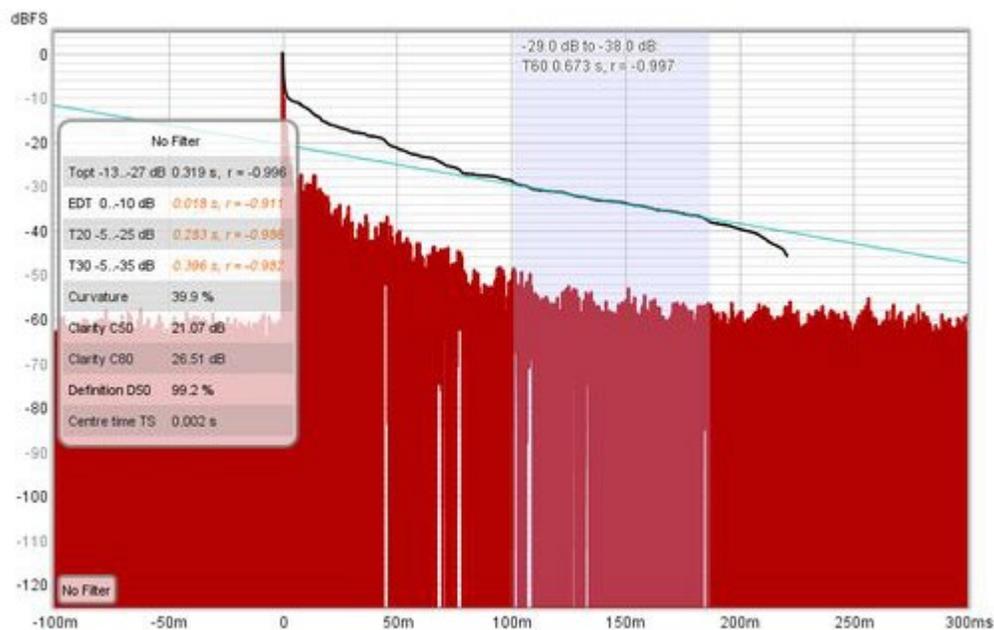
Filtered IR Controls

The control panel for the Filtered IR graph has these controls:

Octave and 1/3 octave filters can be selected using the **1/1** and **1/3** boxes, which enable the corresponding set of controls. The required filter frequency can be chosen directly from the drop-down list or stepped through using the buttons either side. The filter is applied to the Impulse Response upon selection, a label in the bottom left corner of the graph shows the current filter setting. **The filter remains active until "No Filter" is selected or the boxes are unchecked. The measurement name on all graphs is shown with an indication of the applied filter, for example "Auditorium [250Hz 1/3]".** The filters are Butterworth bandpass, whose order is determined by the selection in the **Filter order** control. Orders from 6 to 48 may be selected, the default is 6. Note that higher filter orders have correspondingly higher group delays, which may affect the RT60 estimates if the RT60 values are low (fast decays in a heavily treated space, for example). Using **Time reversed filtering** can avoid that potential issue.

The impulse response may be plotted with or without normalisation to its peak value according to the setting of the **Plot responses normalised** control. When normalised plotting is selected the peak will be at 100% or 0 dBFS.

The **Time reversed filtering** control applies the octave band filters backwards in time, this reduces the filter's own contribution to the measured decay. When using 1/3 octave filters at low frequencies the filter decay time can be significant, over 200 ms for a 100Hz 1/3 octave 6th order filter, for example. Applying the filter in reverse reduces this decay to less than 50 ms, but it does affect the response somewhat, such that Early Decay Time (EDT) figures using Time-Reversed filters may not be valid. Note that this control is independent of and does not affect the time-reversed filtering control on the RT60 graph, but it is initialised to the setting last used for RT60.

Zero phase filtering applies the filters in two passes through the data, one forwards and one reversed, to give a response with overall zero phase shift. This may be useful when comparing the locations of reflection peaks in filtered impulse responses or the ETCs derived from them as the peaks will not be shifted in time by the effect of the filter's phase response. The filter order used for the backward and forward passes is two thirds of the Filter order value, which together with a bandwidth adjustment gives a response that is close to that of a single pass of a filter of the specified order. Note that this control is independent of and does not affect the zero phase filtering control on the RT60 graph, but it is initialised to the setting last used for RT60.

The **Show data panel** control shows a panel on the graph containing the results for the decay values. The RT60 figures include the decay range over which they have been calculated and an "r" value, the regression coefficient, which measures how well the data corresponds to a straight line. A value of -1 would indicate a perfect fit, values lower in magnitude than -0.98 indicate the corresponding decay figure may not be reliable. Unreliable figures are italicised and shown orange.

Topt -5..-49 dB 0.466 s, r = -0.998

EDT 0..-10 dB 0.343 s, r = -0.997

T20 -5..-25 dB 0.405 s, r = -0.997

T30 -5..-35 dB 0.443 s, r = -0.997

Curvature   9.3 %

Clarity C50   9.11 dB

Clarity C80   13.97 dB

Definition D50   89.1 %

Centre time TS 0.020 s

The parameters available are:

EDT

Early Decay Time, based on the slope of the Schroeder curve between 0 dB and -10 dB.

T20

Decay time based on the slope of the Schroeder curve between -5 dB and -25 dB.

T30

Decay time based on the slope of the Schroeder curve between -5 dB and -35 dB.

Topt

An "optimal" decay time based on the slope of the Schroeder curve over a variable range chosen to yield the best linear fit. If the early decay time is much shorter than T30 the Topt measure uses a start point based on the intersection of the EDT and T30 lines, otherwise it uses -5 dB. REW tests every end point in 1dB steps to the end of the Schroeder curve and chooses the one which gives the best linear fit.

Curvature

(T30/T20 - 1) expressed as a percentage, providing an indication of how the slope of the decay curve is changing. Values from 0 to 5% are typical, higher than 10% is suspicious and may indicate that the room has a two-stage decay curve. If curvature is negative the results should be treated with caution as they may be in error.

Clarity C50

The early to late energy ratio in dB, using sound energy in the first 50 ms as the 'early' part. C50 is most often used as an indicator of speech clarity.

Clarity C80

The early to late energy ratio in dB, using sound energy in the first 80 ms as the 'early' part. C80 is most often used as an indicator of music clarity.

Definition D50

The early to total energy ratio as a percentage, using sound energy in the first 50 ms as the 'early' part

Centre Time TS

The time of the 'centre of gravity' of the squared impulse response

The graph can also show the "Regression Line", which is a line obtained by carrying out least squares linear regression on the Schroeder curve over the range applicable to any particular decay parameter. The selector for which regression line is to be shown is next to the **Show Regression Line** check box.

**ETC Smoothing** is used to smooth the envelope (ETC) trace using a moving average filter of the duration specified in the spinner.

The **Trace options** button brings up a dialog that allows the colour and line type of the graph traces to be changed. If a change is made it will be used for all measurements shown on this graph. Traces can also be hidden, which will remove them from the graph and from the graph legend.

Help Index

Group Delay Graph

The group delay for the measurement is calculated from the slope of the phase trace. Note that if smoothing has been applied to the measurement that will also smooth the phase and group delay traces. Smoothing can be applied or removed via the Graph menu and its shortcut keys. Peaks and dips in the frequency response will usually be accompanied by corresponding peaks and dips in the group delay. The group delay will include any delay in the measurement due to time delays in the PC or soundcard, processing delays in the equipment and delays due to the time sound takes to travel from source to microphone. Delays in the PC or soundcard can be eliminated by using the Use loopback as timing reference or Use acoustic timing reference options in the Analysis Preferences. If the group delay is tending towards a level at the upper end of the measurement that level typically corresponds to the overall measurement delay.
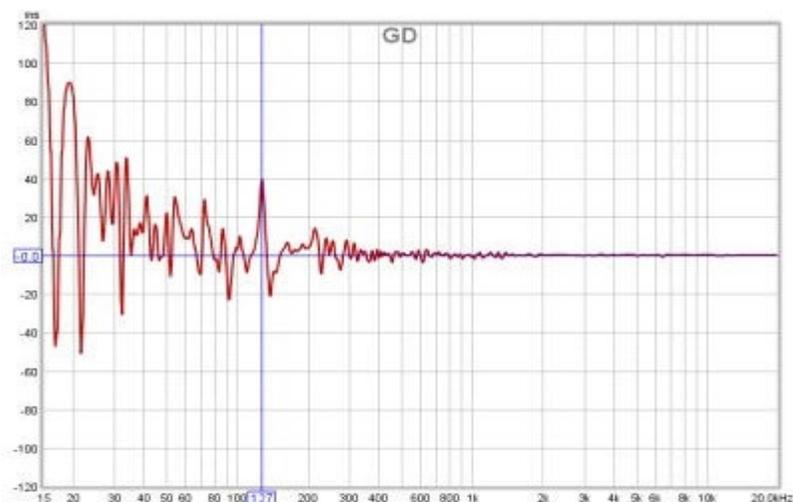
Group Delay Controls

The control panel for the Group Delay graph has these controls:
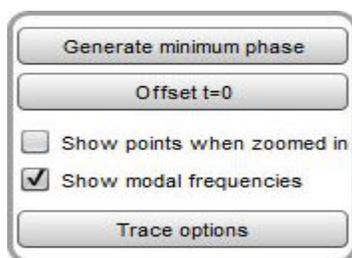
Generate Minimum Phase will produce a minimum phase version of the measurement using the current IR window settings. This activates minimum and excess group delay traces that show how the measurement's group delay compares with the response of a system having the same frequency response but with the lowest phase shift such a system could have. Note that it is best to make full range measurements if the minimum phase response is to be generated as a good result relies on measuring beyond the bandwidth of the system being measured. This control also activates minimum and excess phase and minimum phase impulse traces on the SPL & Phase and Impulse graphs respectively.

Note that the IR window settings are important as the minimum phase response is derived from the frequency (magnitude) response of the measurement, which in turn is affected by the IR window settings. If the window settings are subsequently changed **Generate Minimum Phase** should be used again to reflect the new settings. Note also that the shape of the left side window (the window applied before the peak)

affects the minimum phase result, a rectangular window will produce a response with lower phase shift than, for example, a Hann window.

If the system being measured was inherently minimum phase (as most crossovers are, for example) the minimum phase response is the same as removing any time delay from the measurement. Room measurements are typically not minimum phase except in some regions, mainly at low frequencies. For more about minimum and excess phase and group delay see Minimum Phase.

Offset t=0 allows the position of time zero in the impulse response to be altered, with a live preview of the effect the offset will have on group delay and phase. The measurement is not changed unless either the **Apply** or **Apply & close** button is pressed. Fine adjustment of the offset can be made with the left and right arrow keys after clicking on the slider knob. If a timing reference was used the System Delay figure (which can be viewed in the measurement Info panel) is shifted by the same amount as the zero time. The cumulative shift that has been applied to the impulse response is shown at the top of the dialog.
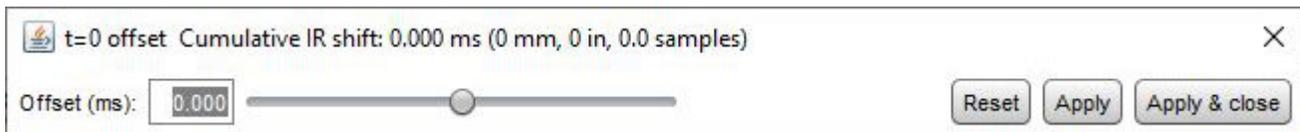
If **Show points when zoomed in** is selected the individual points that make up the measured and minimum phase responses are shown on the graph when the zoom level is high enough for them to be distinguished (which may only be over part of the plot)

If **Show modal frequencies** is selected the theoretical modal frequencies for the room dimensions entered in the Modal Analysis section of the EQ Window are plotted at the bottom of the graph.

The **Trace options** button brings up a dialog that allows the colour and line type of the graph traces to be changed. If a change is made it will be used for all measurements shown on this graph. Traces can also be hidden, which will remove them from the graph and from the graph legend.

Help Index

RT60 Graph

RT60 Reverberation Time values at each octave or one-third octave centre frequency are displayed on this graph, with separate traces for the Early Decay time (EDT), Centre Time (TS) and the 60dB decay times T20, T30 and REW's Topt and T60M. See below for descriptions of each of those parameters and the related clarity, definition and centre time parameters.

RT60 Explanation

RT60 is a measure of how long sound takes to decay by 60 dB in a space that has a diffuse soundfield, meaning a room large enough that reflections from the source reach the mic from all directions at the same level. Domestic rooms are usually too small to have anything approaching a diffuse field at low frequencies as their behaviour in that region is dominated by modal resonances. As a result RT60 is typically not meaningful in such rooms below a few hundred Hz. Use the RT60 Decay, waterfall, spectrogram and Decay plots to examine the decay of low frequencies in domestically-sized rooms.
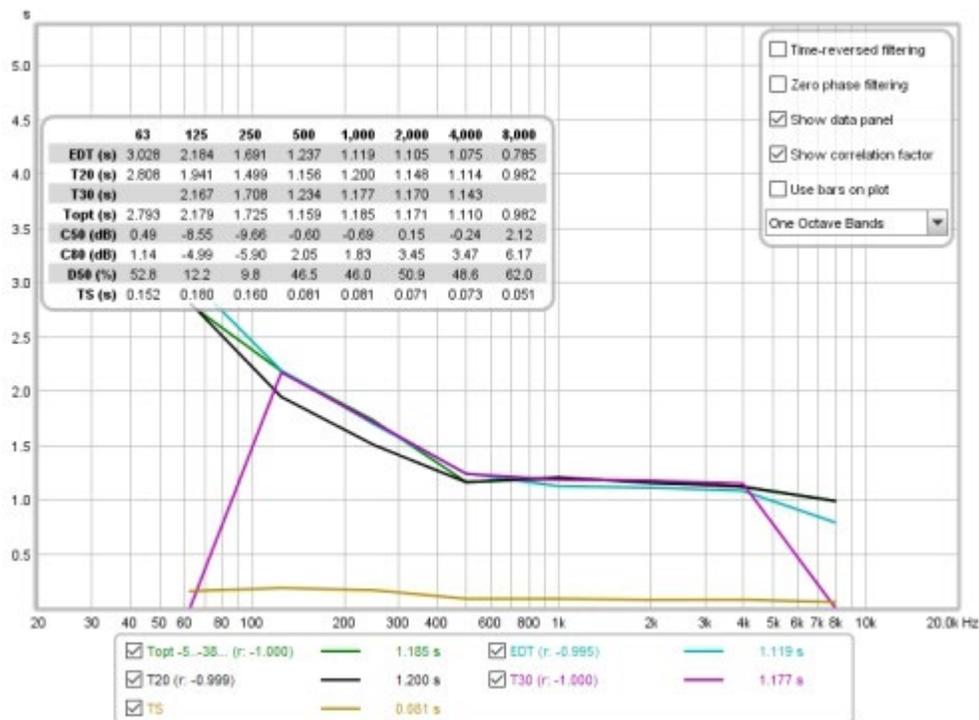
RT60 Calculation

The RT60 values are estimated by calculating the slope of the Schroeder curve, which is a plot of the energy (squared values) of the impulse response that is backwards integrated (summed starting from the end and moving backwards). The vertical axis of the plot is in dB. The different RT60 measures (e.g. T20, T30, REW's Topt) are derived by calculating the slope of a best fit line to the Schroeder curve over different ranges (detailed below). In a diffuse field the curve with the dB vertical scale is quite linear until it reaches the noise floor.

The start point for the classical T20 and T30 measures of RT60 is where the Schroeder curve has dropped 5 dB below its peak. That works well in the large spaces for which RT60 is most applicable, particularly if the source used for the measurement is omnidirectional. In domestically sized rooms using normal, directional loudspeakers as sources the initial drop of the Schroeder curve is quite sharp (the Early Decay Time is quite short) meaning the -5 dB point lies within the early decay region rather than the diffuse field region. That in turn means the T20 and T30 figures underestimate the RT60 time. Where the EDT is much shorter than the T30 RT60 figure REW's Topt RT60 calculation uses a start point based on the intersection of the EDT and T30 regression lines, to determine a point that lies within the diffuse field region. It then tests each possible end point in 1 dB steps and picks the one that gives a regression line with the best linear fit. That produces a more reliable RT60 figure.

The results are presented within octave or 1/3 octave bands, providing a view of how the rate at which sound decays changes with frequency. The lowest centre frequency for the RT60 plot is 50 Hz at 1/3 octave and 63 Hz at 1 octave, per ISO3382. For domestic listening rooms and recording studios with volumes of less than 50 cubic metres (1,800 cubic feet) the recommended RT60 value is 0.3 s. For larger rooms, up to 200 cubic metres (7,000 cubic feet) the recommendation is 0.4 to 0.6 s. In both cases the value should be fairly uniform across the frequency range, though it will typically tend to increase at lower frequencies.

Important: If you are making RT60 measurements in a large space (bigger than a domestic room) change the IR truncation setting in the Analysis preferences to make sure enough of the IR is retained for the decay to reach the noise floor. It may also be necessary to use a longer sweep (with a 256k sweep there are about 6 seconds of IR data after the peak before any truncation is applied).

	63	125	250	500	1,000	2,000	4,000	8,000
EDT (s)	3.028	2.184	1.691	1.237	1.119	1.105	1.075	0.785
T20 (s)	2.808	1.941	1.499	1.156	1.200	1.148	1.114	0.982
T30 (s)		2.167	1.708	1.234	1.177	1.170	1.143	
Topt (s)	2.793	2.179	1.725	1.159	1.185	1.171	1.110	0.982
C50 (dB)	0.49	-8.55	-9.66	-0.60	-0.69	0.15	-0.24	2.12
C80 (dB)	1.14	-4.99	-5.90	2.05	1.83	3.45	3.47	6.17
D50 (%)	52.8	12.2	9.8	46.5	46.0	50.9	48.6	62.0
TS (s)	0.152	0.180	0.160	0.081	0.081	0.071	0.073	0.051

RT60 and related parameters

The parameters which are calculated are:

EDT

Early Decay Time, based on the slope of the Schroeder curve between 0 dB and -10 dB. This is not an RT60 figure, but rather an indication of how quickly the initial sound at the measurement position decays - it is much more location dependent than RT60. A fast early decay (low EDT figure) indicates better clarity than at positions where the EDT is higher.

T20

The RT60 decay time based on the slope of the Schroeder curve between -5 dB and -25 dB.

T30

The RT60 decay time based on the slope of the Schroeder curve between -5 dB and -35 dB.

Topt

An "optimal" RT60 decay time based on the slope of the Schroeder curve over a variable range chosen to yield the best linear fit. If the early decay time is much shorter than T30 the Topt measure uses a start point based on the intersection of the EDT and T30 regression lines, otherwise it uses -5 dB. REW tests every end point in 1dB steps to the end of the Schroeder curve and chooses the one which gives the best linear fit.

Clarity C50

The early to late energy ratio in dB, using sound energy in the first 50 ms as the 'early' part. C50 is

most often used as an indicator of speech clarity.

Clarity C80

The early to late energy ratio in dB, using sound energy in the first 80 ms as the 'early' part. C80 is most often used as an indicator of music clarity.
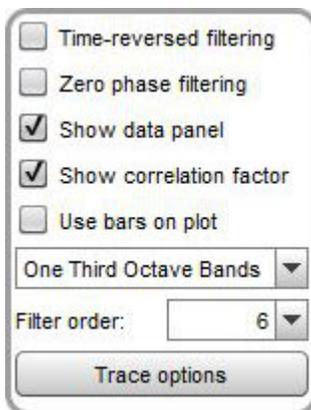
Definition D50

The early to total energy ratio as a percentage, using sound energy in the first 50 ms as the 'early' part

Centre Time TS

The time of the 'centre of gravity' of the squared impulse response

RT60 Controls

The control panel for the RT60 graph has these controls:

The **Time reversed filtering** control applies the octave band filters backwards in time, this greatly reduces the filter's own contribution to the measured decay. When using 1/3 octave filters at low frequencies the filter decay time can be significant, over 200 ms for a 100Hz 1/3 filter, for example. Applying the filter in reverse reduces this decay to less than 50 ms, but it does affect the response somewhat, such that Early Decay Time (EDT) figures using Time-Reversed filters may not be valid.

Zero phase filtering applies bandpass filters in two passes through the data, one forwards and one reversed, to give a response with overall zero phase shift. This reduces the contribution of the filter's decay time similar to (though not as much as) time reversed filtering but without significantly affecting early decay time. The filter order used for the backward and forward passes is two thirds of the Filter order value, which together with a bandwidth adjustment gives a response that is close to that of a single pass of a filter of the specified order.

The **Show data panel** control shows a panel on the graph containing the results for the decay values. Unreliable figures are italicised and shown orange.

	63	125	250	500	1,000	2,000	4,000	8,000
EDT (s)	3.028	2.184	1.691	1.237	1.119	1.105	1.075	0.785
T20 (s)	2.808	1.941	1.499	1.156	1.200	1.148	1.114	0.982
T30 (s)		2.167	1.708	1.234	1.177	1.170	1.143	
Topt (s)	2.793	2.179	1.725	1.159	1.185	1.171	1.110	0.982
C50 (dB)	0.49	-8.55	-9.66	-0.60	-0.69	0.15	-0.24	2.12
C80 (dB)	1.14	-4.99	-5.90	2.05	1.83	3.45	3.47	6.17
D50 (%)	52.8	12.2	9.8	46.5	46.0	50.9	48.6	62.0
TS (s)	0.152	0.180	0.160	0.081	0.081	0.071	0.073	0.051

If the **Show correlation factor** box is checked the graph legend names shows the quality of the line fit for the various decay measures. The "r" value shown after each decay measure is the regression coefficient, which measures how well the data corresponds to a straight line. A value of -1 would indicate a perfect fit, values lower in magnitude than -0.99 indicate the corresponding decay figure may not be reliable. Potentially unreliable values are italicised.

The RT60 plot can show horizontal bars centred on each filter frequency and spanning the filter's bandwidth, or lines joining the filter centre frequencies, according to the **Use bars on plot** control setting.

RT60 data can be shown in octave or one-third octave bands. The **Filter order** selection controls how steeply the fractional octave band filters roll off. The default order is 6 and that gives results comparable with most other software. Higher filter order makes the bands more selective, which can help isolate longer decays to a particular band, but increasing filter order also increases the delay of the filter itself. That can be mitigated by using **Time reversed filtering** or **Zero phase filtering** as described above.

The **Trace options** button brings up a dialog that allows the colour and line type of the graph traces to be changed. If a change is made it will be used for all measurements shown on this graph. Traces can also be hidden, which will remove them from the graph and from the graph legend.

The parameter values (RT60 and clarity) for the current measurement can be written to a text file using the File → Export → **RT60 data as text** menu entry.

Help Index

RT60 Decay Graph

REW's RT60 Decay graph provides a way to examine reverberation time behaviour at much higher frequency resolutions and with much narrower octave fractions than is usually possible, even at low frequencies, thanks to a frequency domain approach.
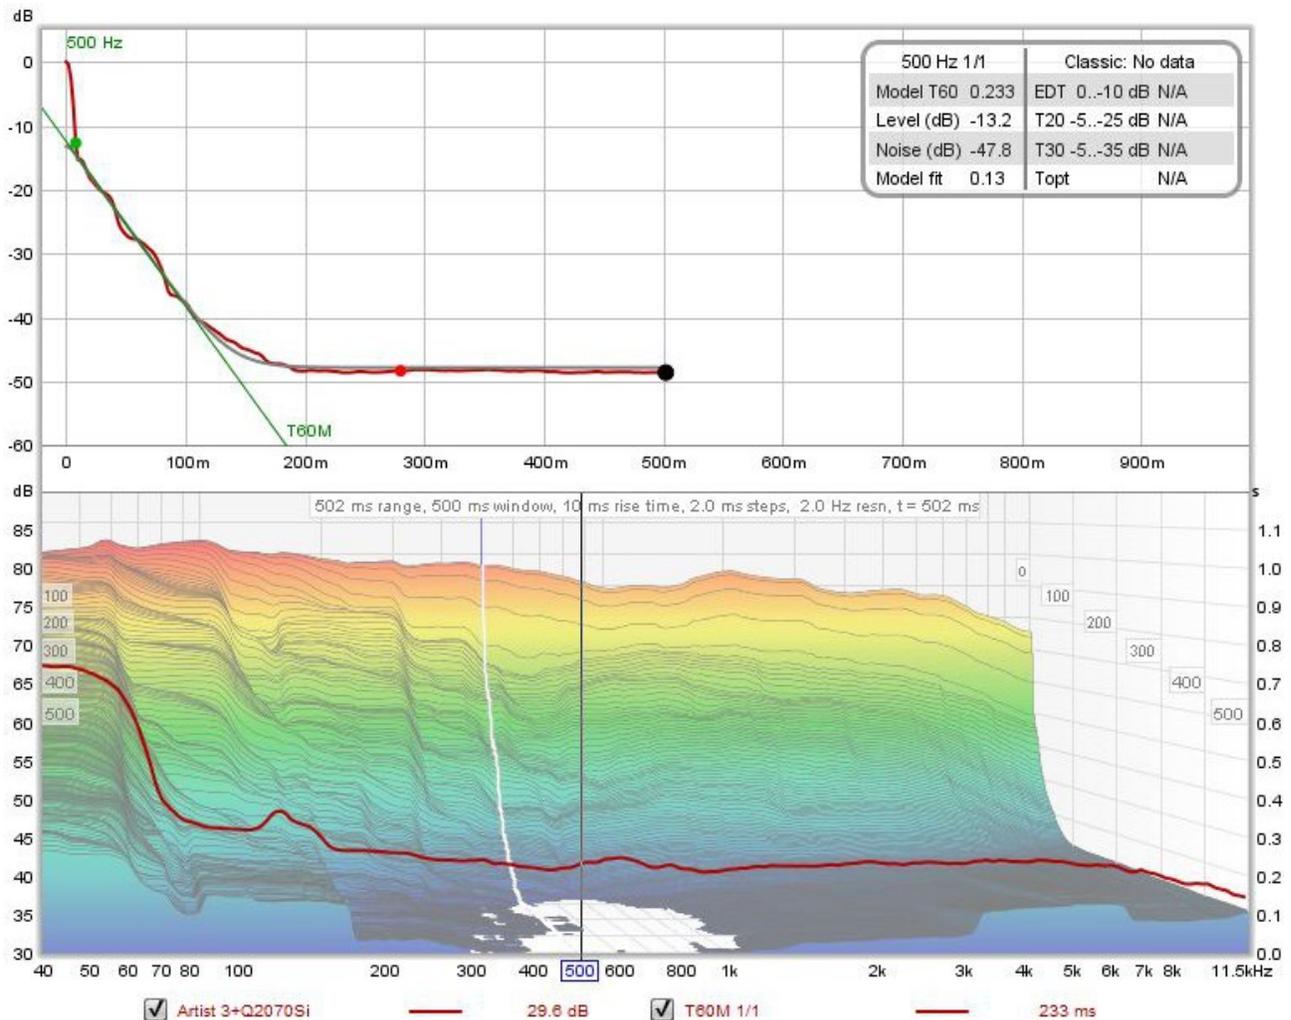
Frequency Domain RT60 Estimation

REW's RT60 Decay Graph takes a frequency domain approach to estimating RT60, rather than the more usual time domain approach.

- Classical RT60 Estimation
- A Frequency Domain Approach to RT60 Estimation
- Selectivity comparison
- Accuracy comparison
- Implementation

Classical RT60 Estimation

Reverberation times are usually determined through time domain processing. The various RT60 values are estimated by calculating the slope of the Schroeder integral, which is a plot of the energy (squared values) of the impulse response that is backwards integrated (summed starting from the end and moving backwards). The different RT60 measures (e.g. T20, T30, REW's Topt) are then derived by calculating the best fit line to the Schroeder integral over different ranges. There are some limitations to this approach. The measurement's noise floor affects the shape of the Schroeder integral, causing it to curve.

Various measures are applied to reduce that affect by estimating the level of the noise floor and altering where the integration starts and the initial values it uses. In addition, the octave band filters have group delays, which may affect the calculated reverberation times, particularly at low frequencies. The filters also usually have poor frequency discrimination, since their order is generally low (6th order typically) to avoid exacerbating the group delay problem. Strong resonances consequently have a masking effect on nearby regions of the response.

A Frequency Domain Approach to RT60 Estimation

Reverberation times may also be estimated through frequency domain processing, examining the decays of the slices of a set of Short-Time Fourier Transform (STFT) plots. These are commonly used to produce waterfall or spectrogram graphs. The slices of an STFT series can be octave band filtered in the frequency domain using brick wall filters with no associated time domain group delays. Frequency discrimination is then determined by the characteristics of the window used for the STFT, the deciding factor generally being the width and shape of the left side of the window (aka the window rise time). Processing in the frequency domain makes it straightforward to present results using much narrower octave band filters, allowing individual resonances to be distinguished with no negative impact on the reverberation time estimate.

The Schroeder integral is not applicable to frequency domain processing. The STFT slices capture the reverberant decays and noise, reverberation times may be estimated by fitting an exponential decay plus noise function to the data series formed by the slice values at each time step. The fitting process must take into account the effect of the left window on the data series, where its width is more than the slice interval. It must also be tolerant of the non-monotonic nature of the decay data.

Here is an example of an RT60 decay plot after generating the STFT data (by pressing the **Generate** button in the bottom left corner of the lower graph) but before calculating the RT60 model. The upper panel shows the decay curve for the cursor position in the lower graph.
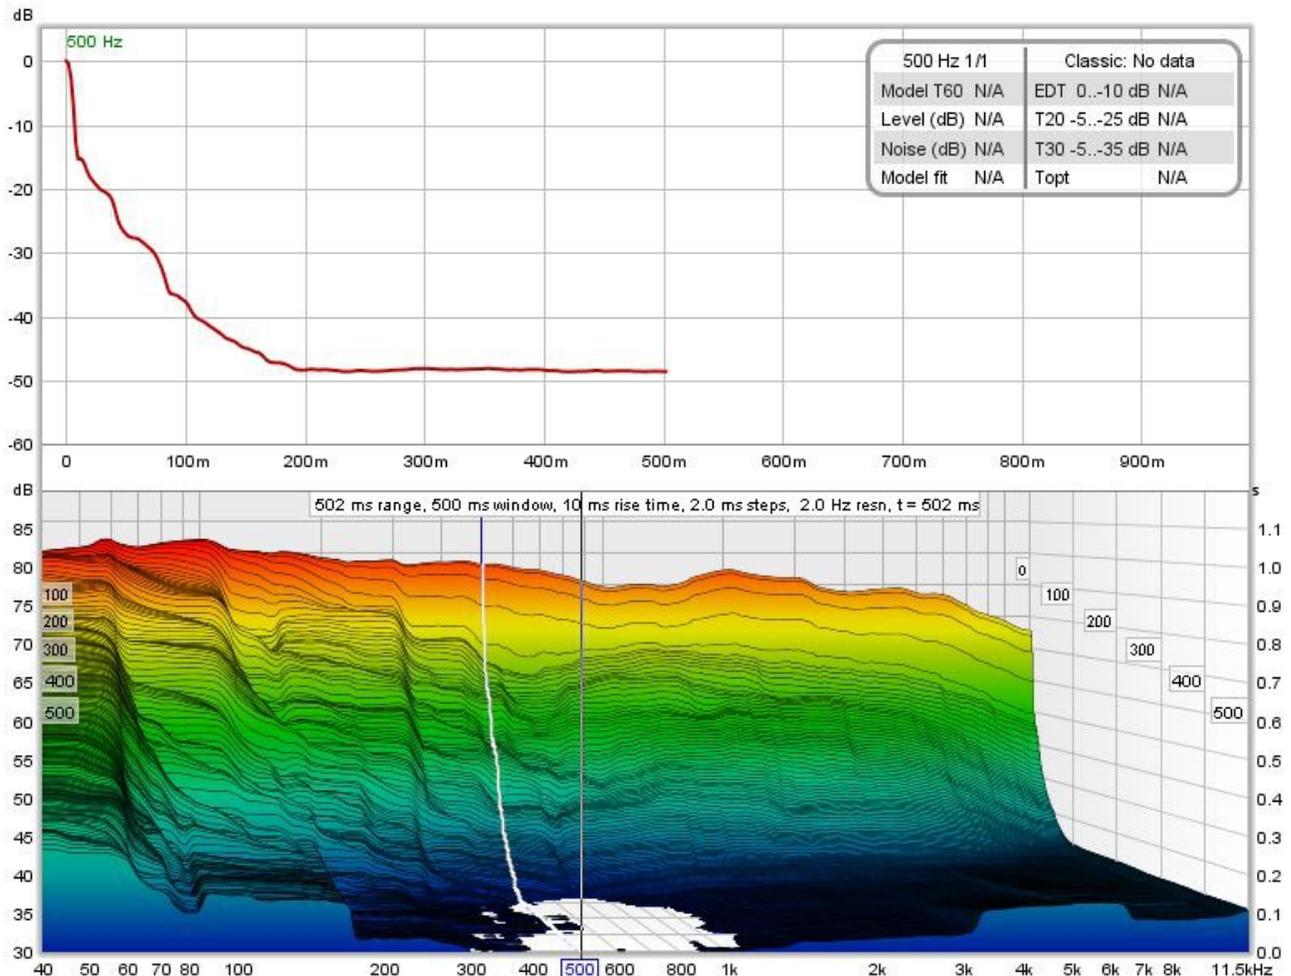

The control panel in the lower graph has a button to calculate the RT60 model.

On the left side of the panel are controls for the STFT data. When **Automatic** is selected REW will determine the rise time, window width and time span to use for the STFT plot. **Manual** allows those settings to be altered, using the button to the right of the Manual selection. **Octave band** is the width of the octave band filter applied to the STFT data, while **Results PPO** determines the frequency resolution at which REW will calculate the model (up to 96 PPO). **Show Schroeder integral** allows the Schroeder integral for the nearest classical RT60 calculation result to be shown alongside the decay profile. The **Prev. 1/1** and **Next 1/1** buttons move the cursor to the previous or next 1 octave centre frequency.

On the right side of the panel are the frequency range over which the model fit will be calculated and the **Calculate RT60 model** button which runs the calculation. **Perspective settings** brings up a dialog to control the presentation of the STFT data while **Colour scheme settings** brings up a dialog to control the colour scheme.

After the model has been calculated for each frequency the model result (T60M) will appear as an overlay on the STFT data, using a time axis on the right hand side of the plot. It can be shown or hidden by selecting or deselecting the T60M trace in the graph legend. In the upper graph the model result is shown as a grey curve, with a regression line corresponding to the calculated RT60 value shown in blue. The data panel shows the T60 time, the starting level of the model curve, the noise level and the model fit error (the lower the value the better the fit). Three dots are shown on the decay curve. Green and red indicate the start and end of the range used for the model fit. Black indicates the point past which REW considers the data not to be valid.
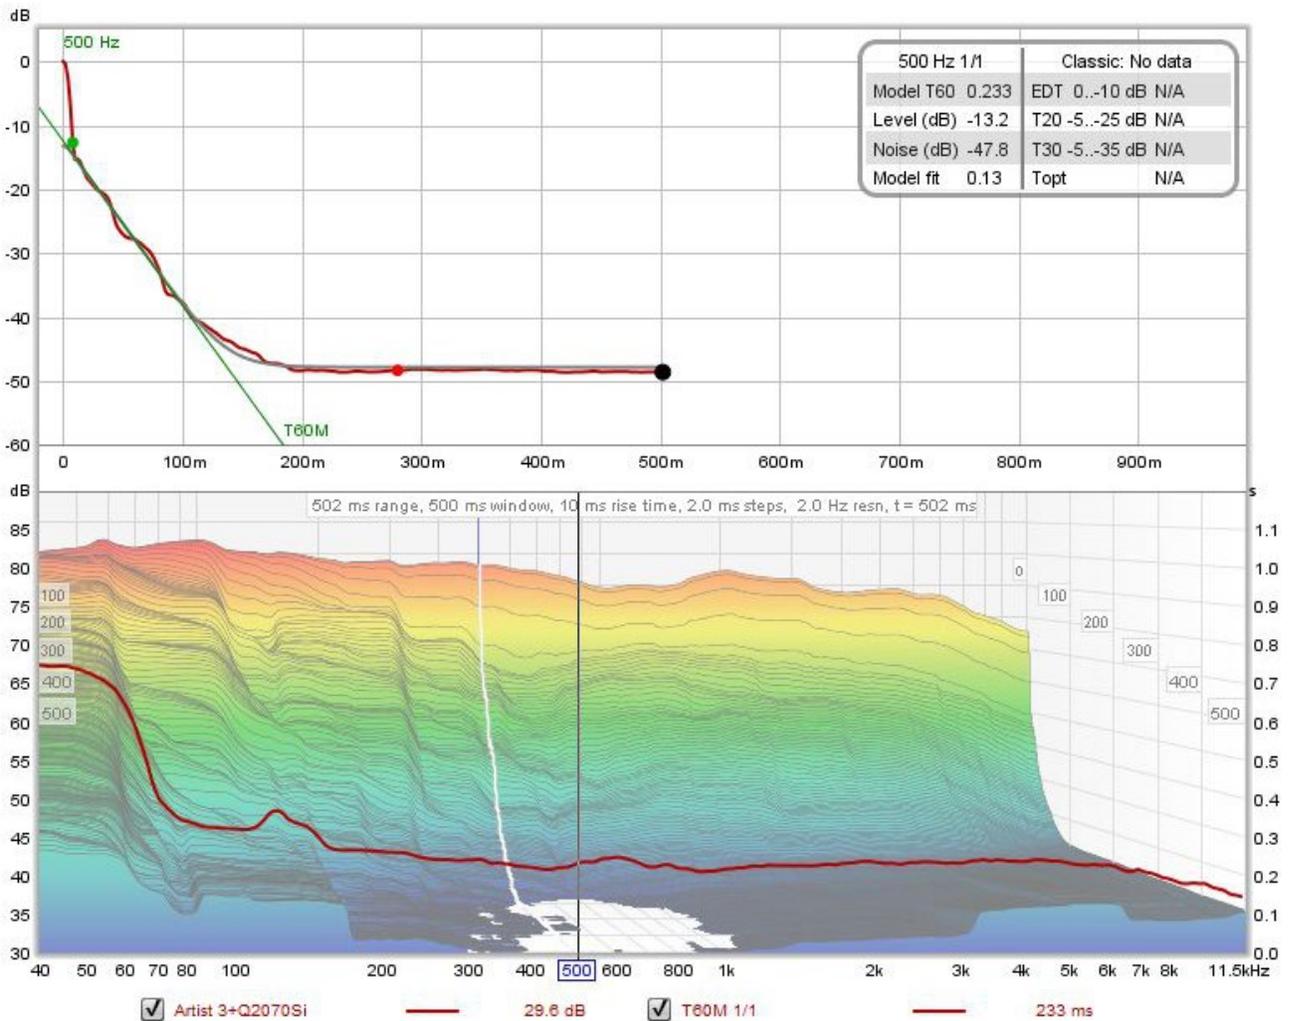
If classical RT60 results are generated for the measurement (from the RT60 graph) those figures will appear, along with the Schroeder integral and its regression line if they have been selected for display.
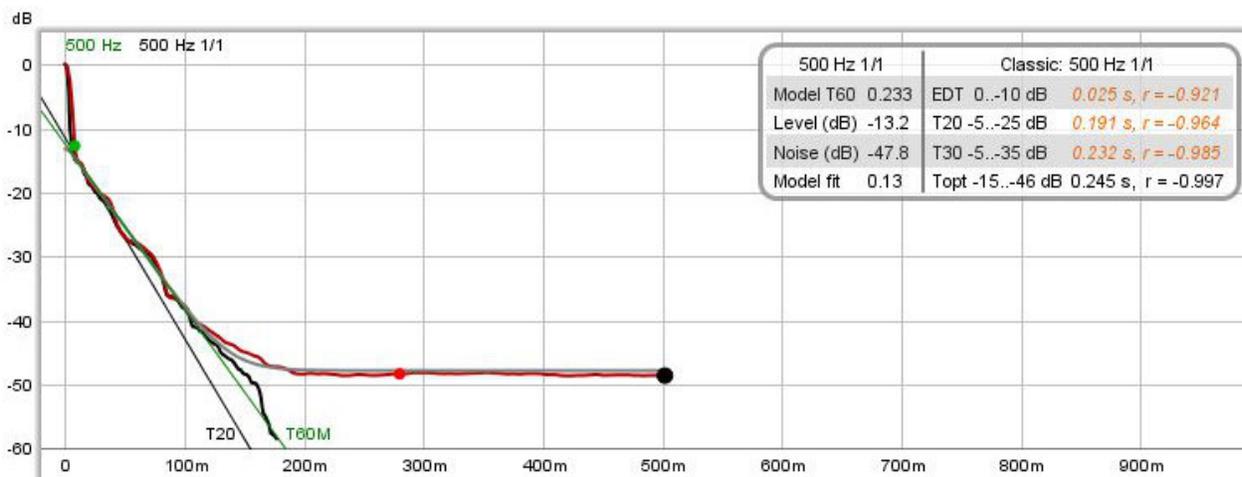
In this example the classical T20 RT60 value is lower than the model value. That is common in responses that have sharp early decays, as may be the case in heavily treated spaces. The -5 dB point on the Schroeder integral which is the starting point for T20 and T30 values is still within the early decay, so biases the results.

After the model has been calculated the result can also be viewed on the RT60 graph, allowing easy comparison with the classical measures.
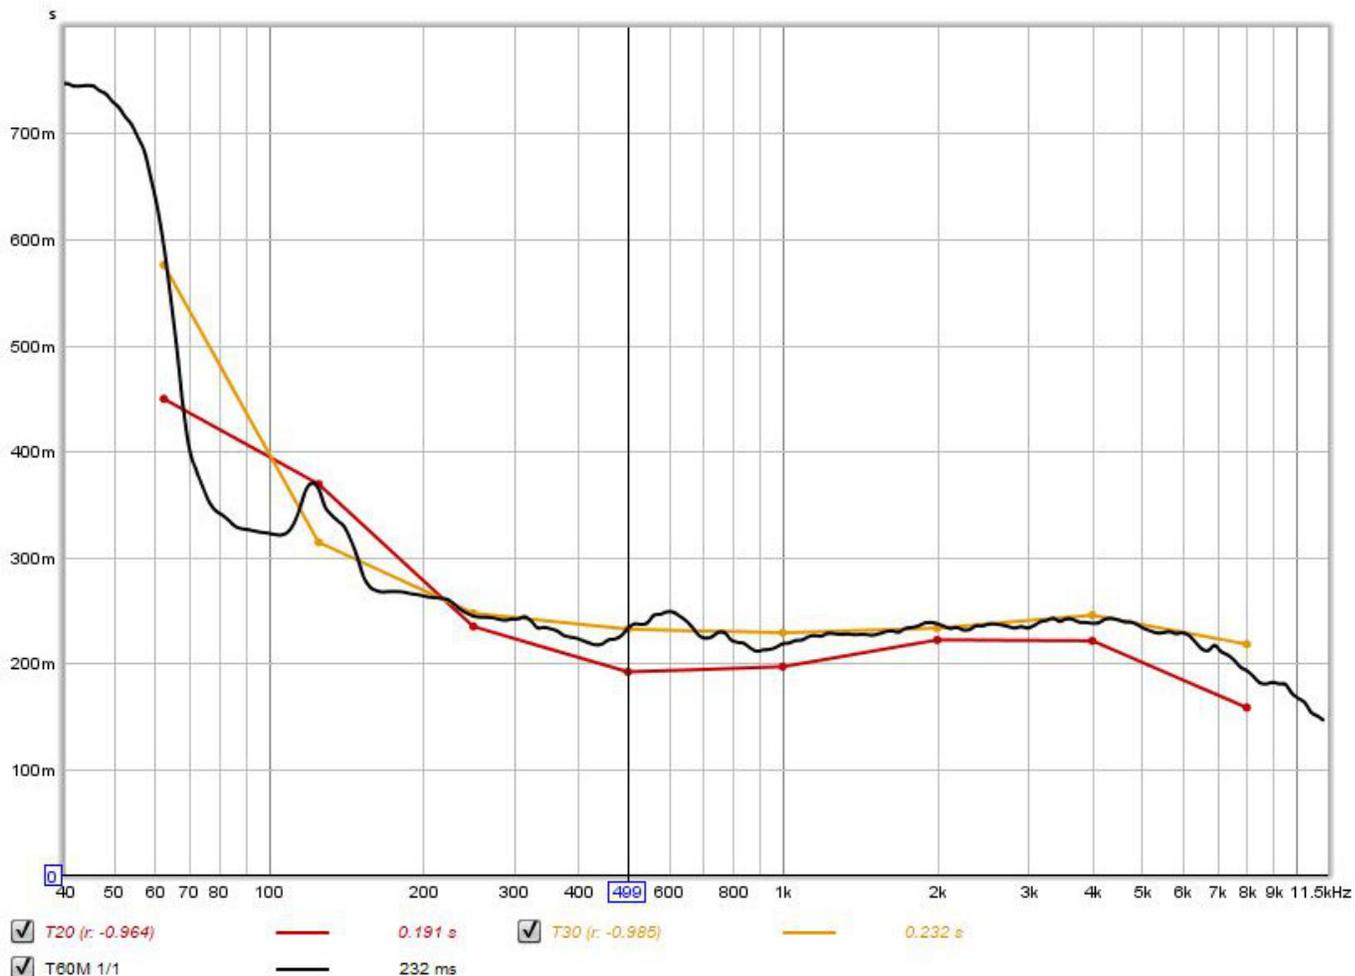
Selectivity comparison

To aid in testing the selectivity of classical and frequency domain approaches a synthetic test signal was generated, consisting of equal amplitude exponentially decaying cosines at the one octave centre frequencies, with noise added at -60 dB relative to the cosine amplitudes. The cosine 60 dB decay rates alternate between 1.000 s and 0.100 s. Here is the spectrum of the test signal.

Selectivity signal spectrum
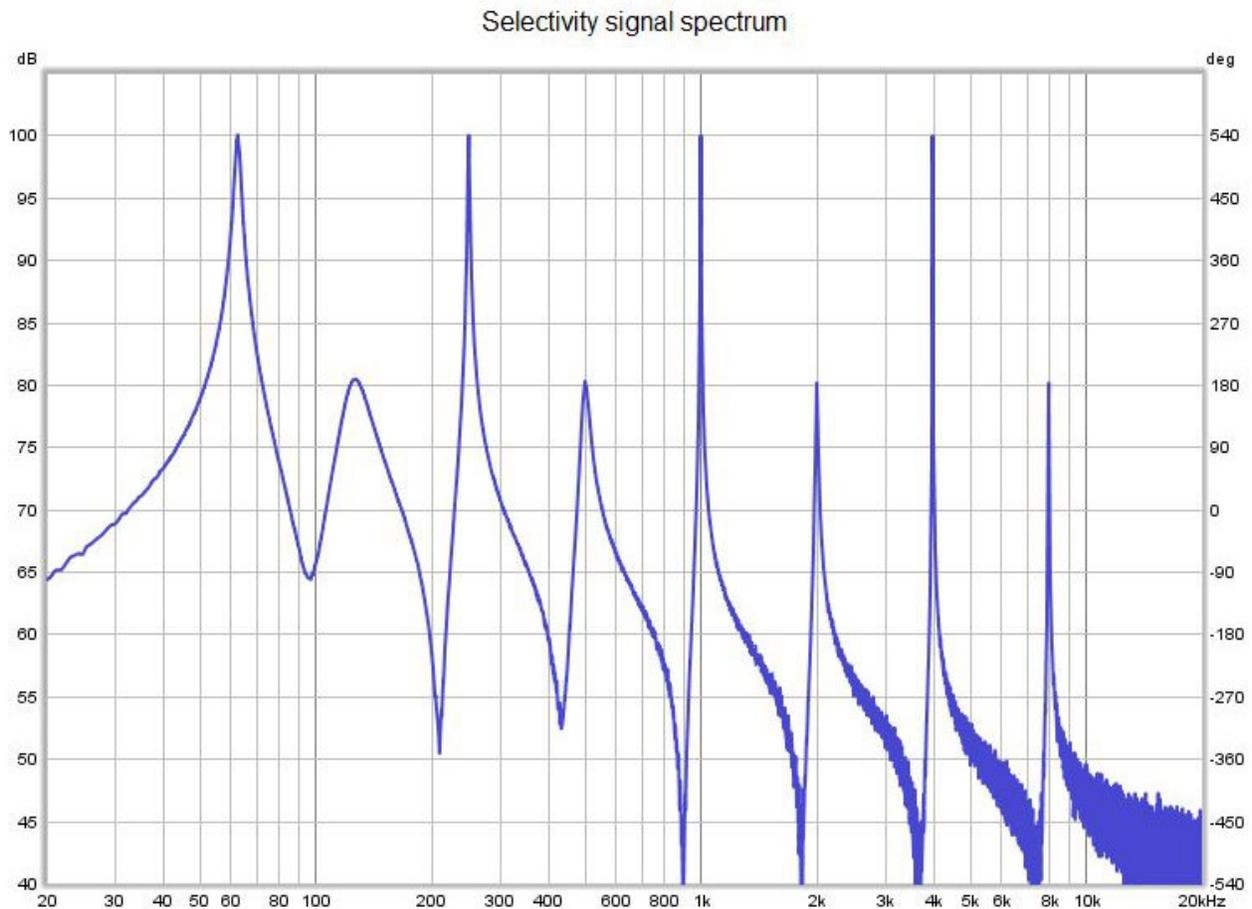
The signal will be octave band filtered in both approaches. We will first look at the classical approach, starting with the 6th order filters typically used.

T20 and T30 with 6th order filters
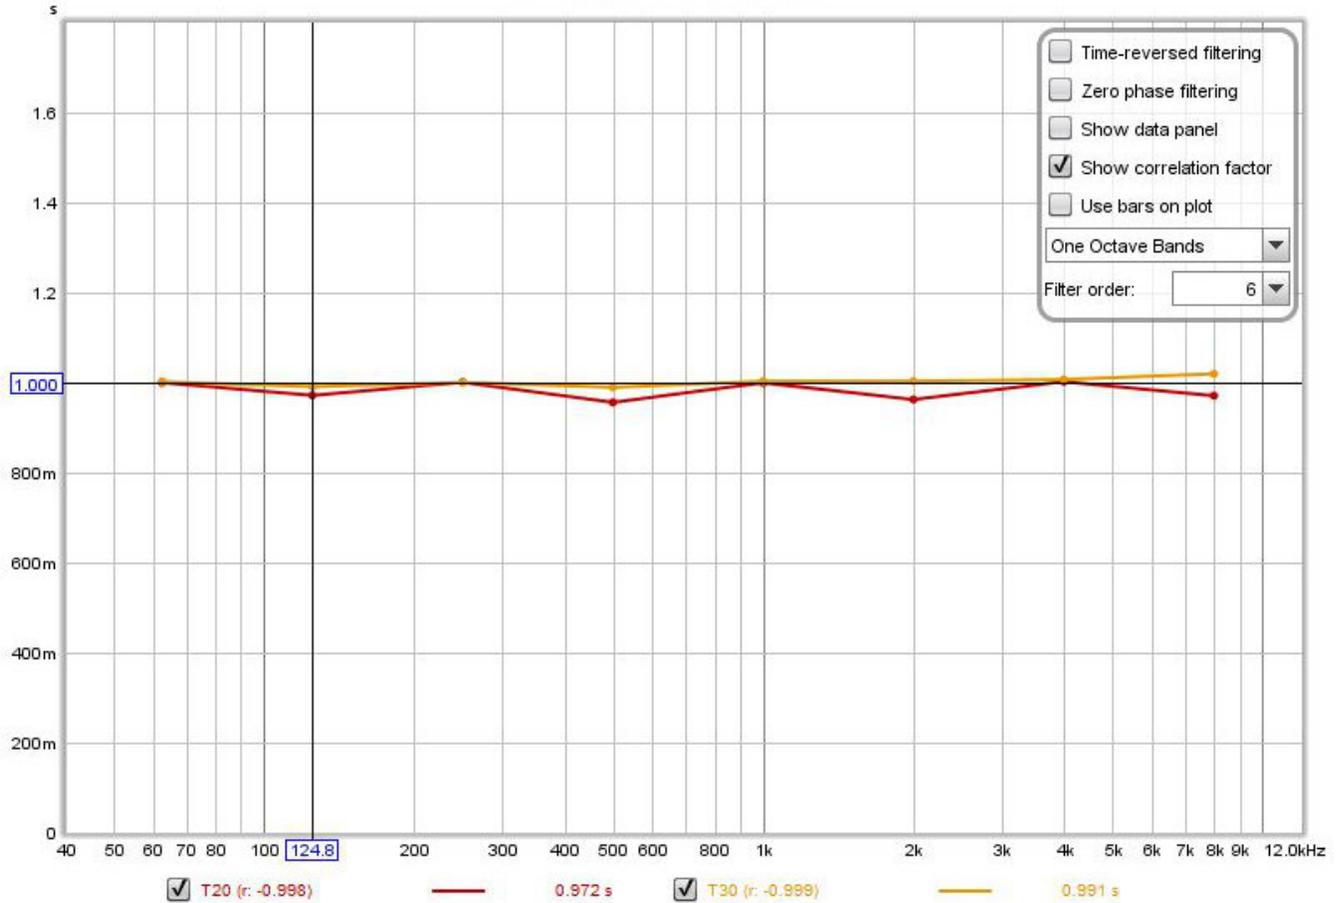
With 6th order filters the 0.100 s decays are not resolved at all. Next we will try 12th order filters.

T20 and T30 with 12th order filters
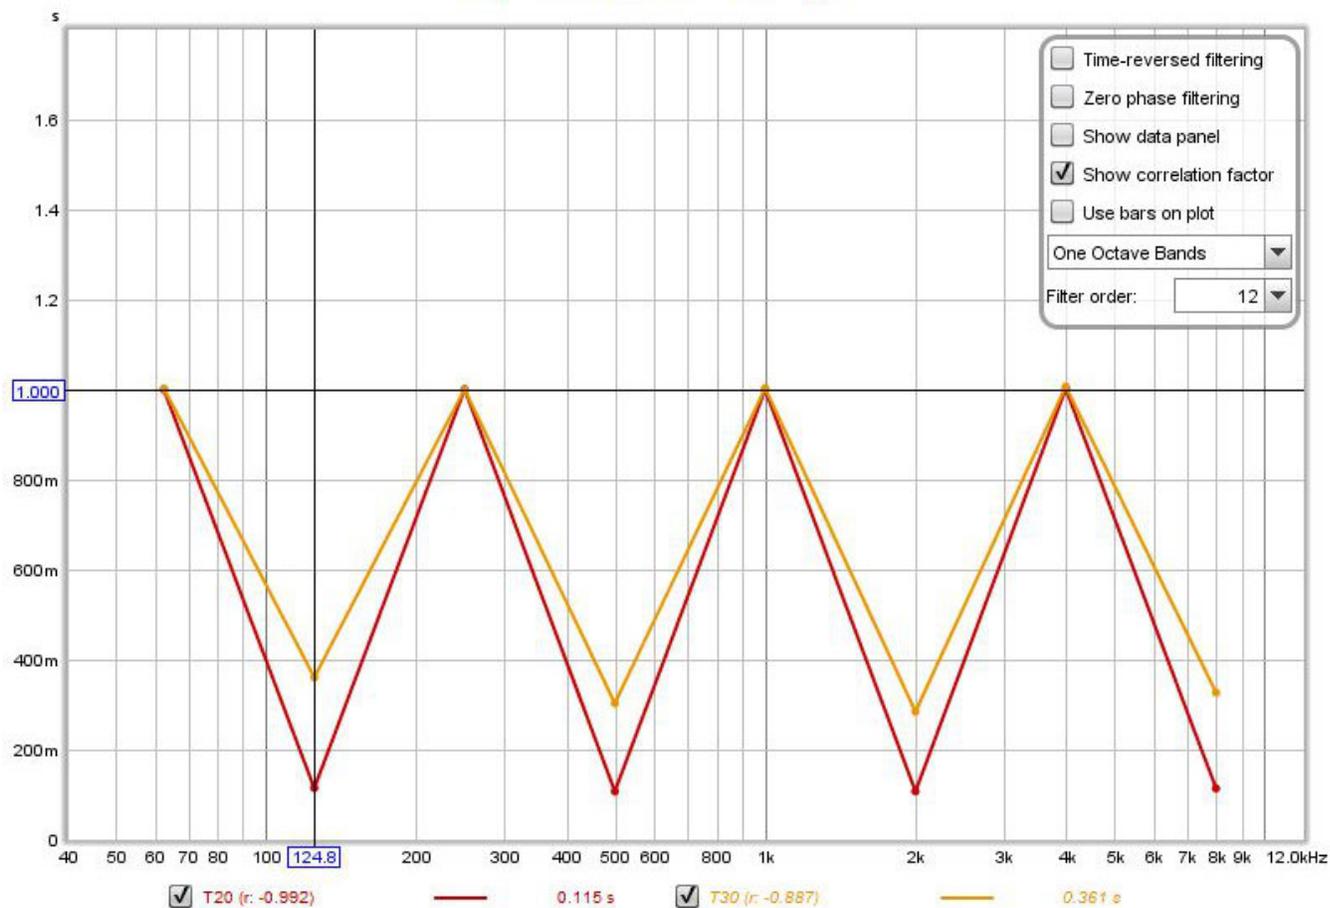
The 12th order filters are beginning to resolve the 0.100 s decays, but the results are still some way off. Next up is 18th order.

T20 and T30 with 18th order filters
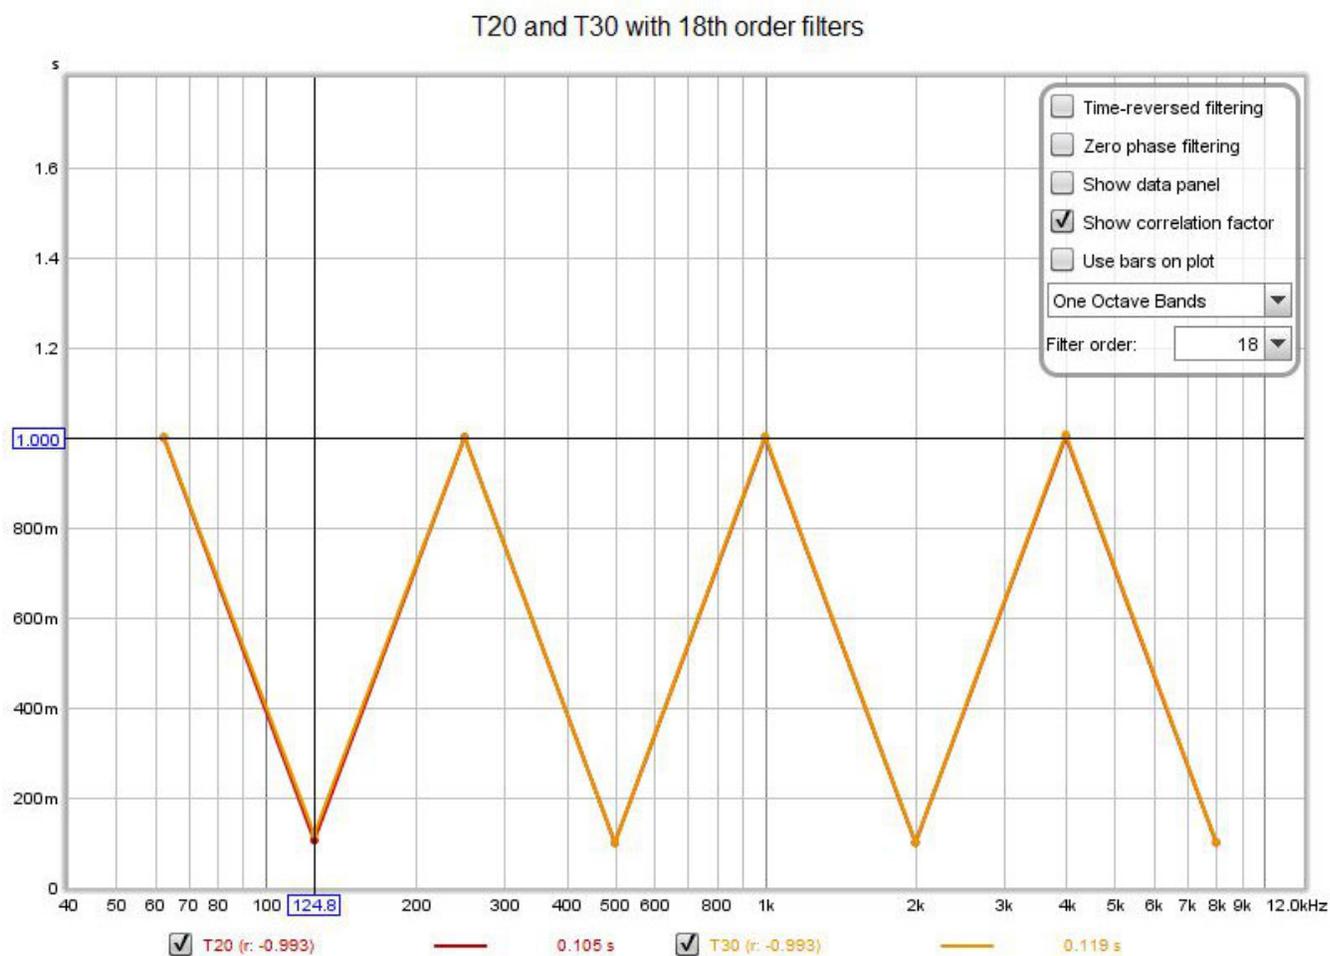
The 18th order filters have done the trick except for the lowest 125 Hz 0.100 s decay, with T30 still slightly worse than T20 there. Increasing filter order further makes the results worse rather than better, as the filter group delays begin to bias the results.

T20 and T30 with 24th order filters
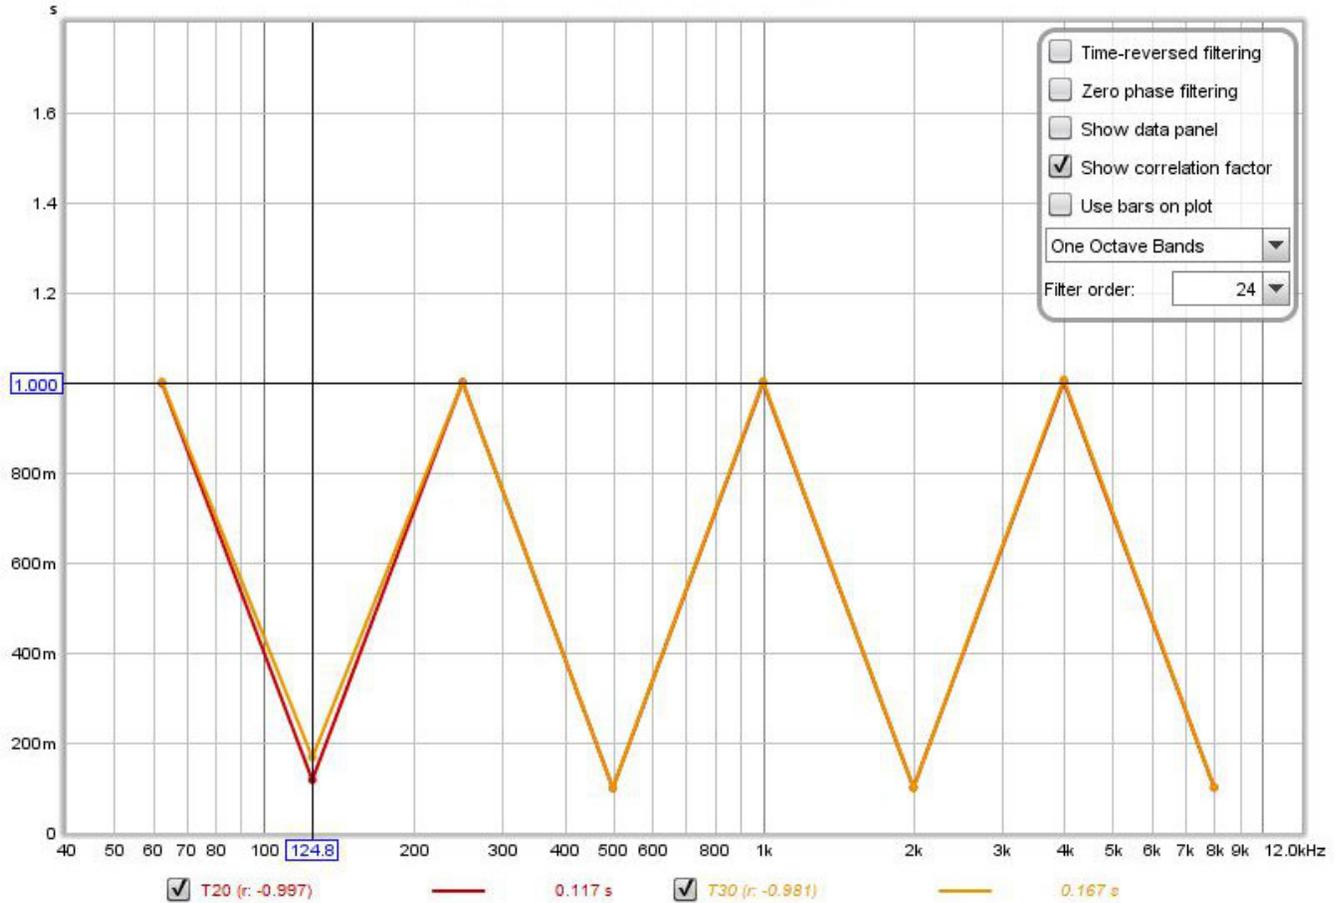
Perfect results can be obtained with 18th order filters and time reversed filtering.

T20 and T30 with 18th order filters, time reversed
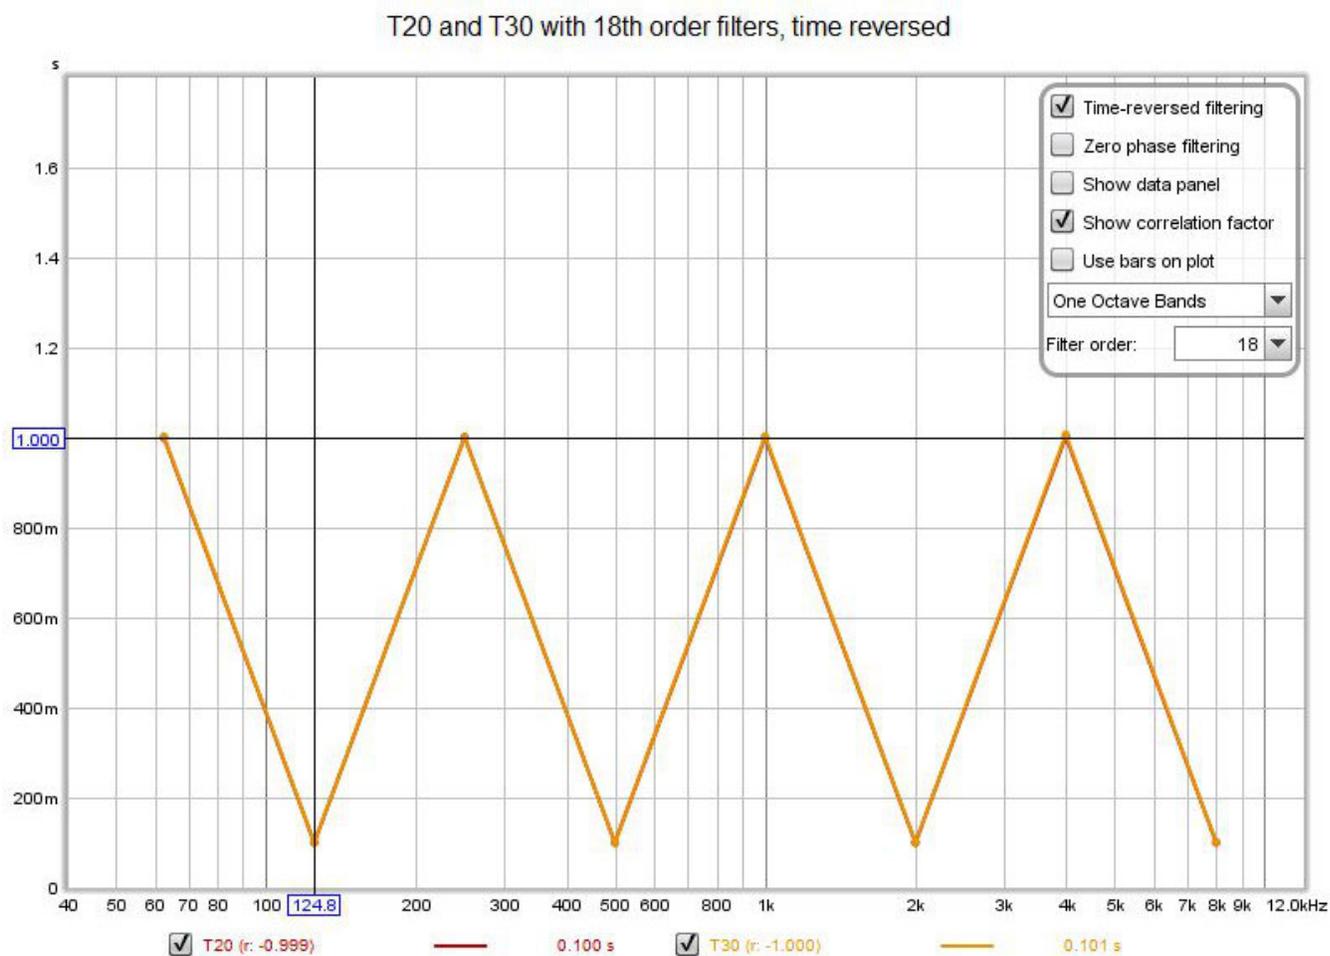
The frequency domain result depends on the window settings used. By default REW picks a 20 ms rise time for this signal, giving the following results (calculated at 96 PPO rather than only at the octave centres).

T20 and T30 18th order time reversed, T60M 20 ms rise time
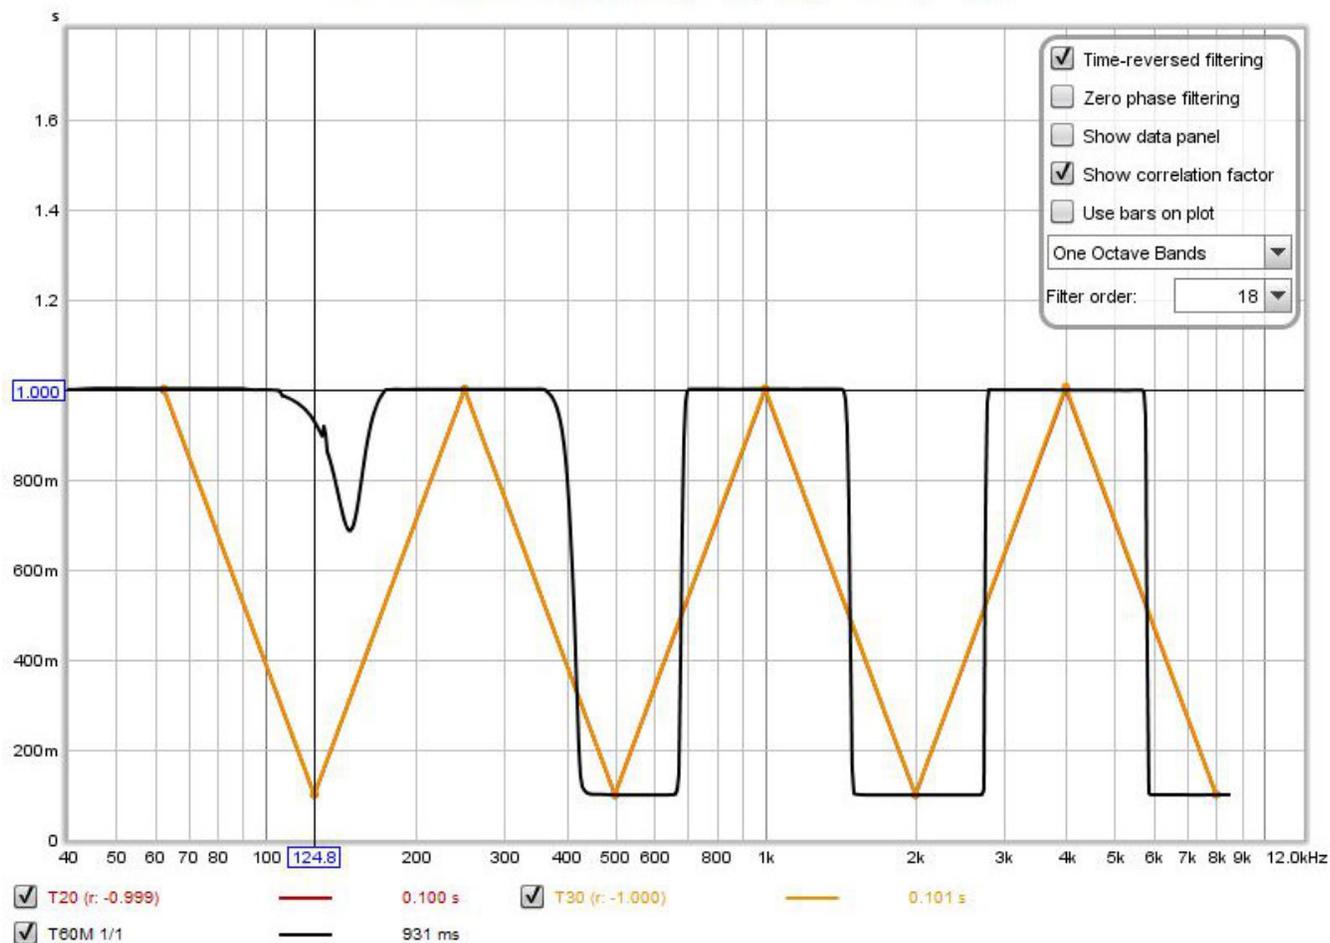
A 100 ms rise time improves the T60M performance, though it still can't quite match the classical approach with 18th order time reversed filters at the 125 Hz octave centre.

T20 and T30 18th order time reversed, T60M 100 ms rise time
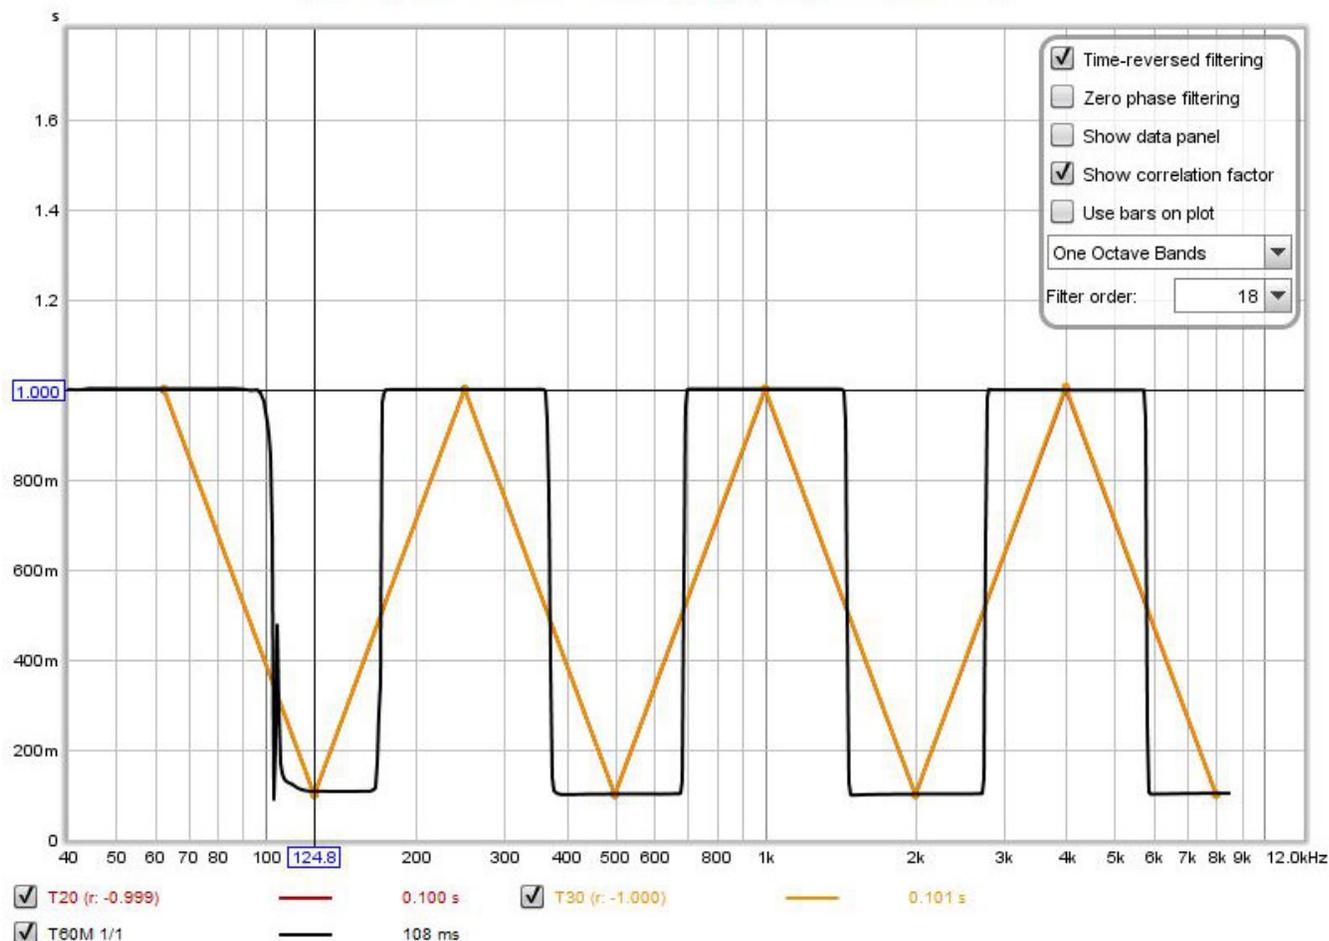
Further increases in T60M rise time improve the selectivity but degrade the accuracy of the RT60 calculation due to the effect using windows wider than the decay time has on the decay data. Nonetheless this is good performance from the alternative approach, particularly if we compare the behaviour of both approaches away from the one octave centres. The plot below uses 1 octave filter bandwidth, 18th order time reversed for T20 and T30, but also evaluates the result at the intervening one-third octave centre frequencies. T60M performs far better.

One octave 18th order time reversed filters at 3rd octave centres vs T60M 100 ms rise time
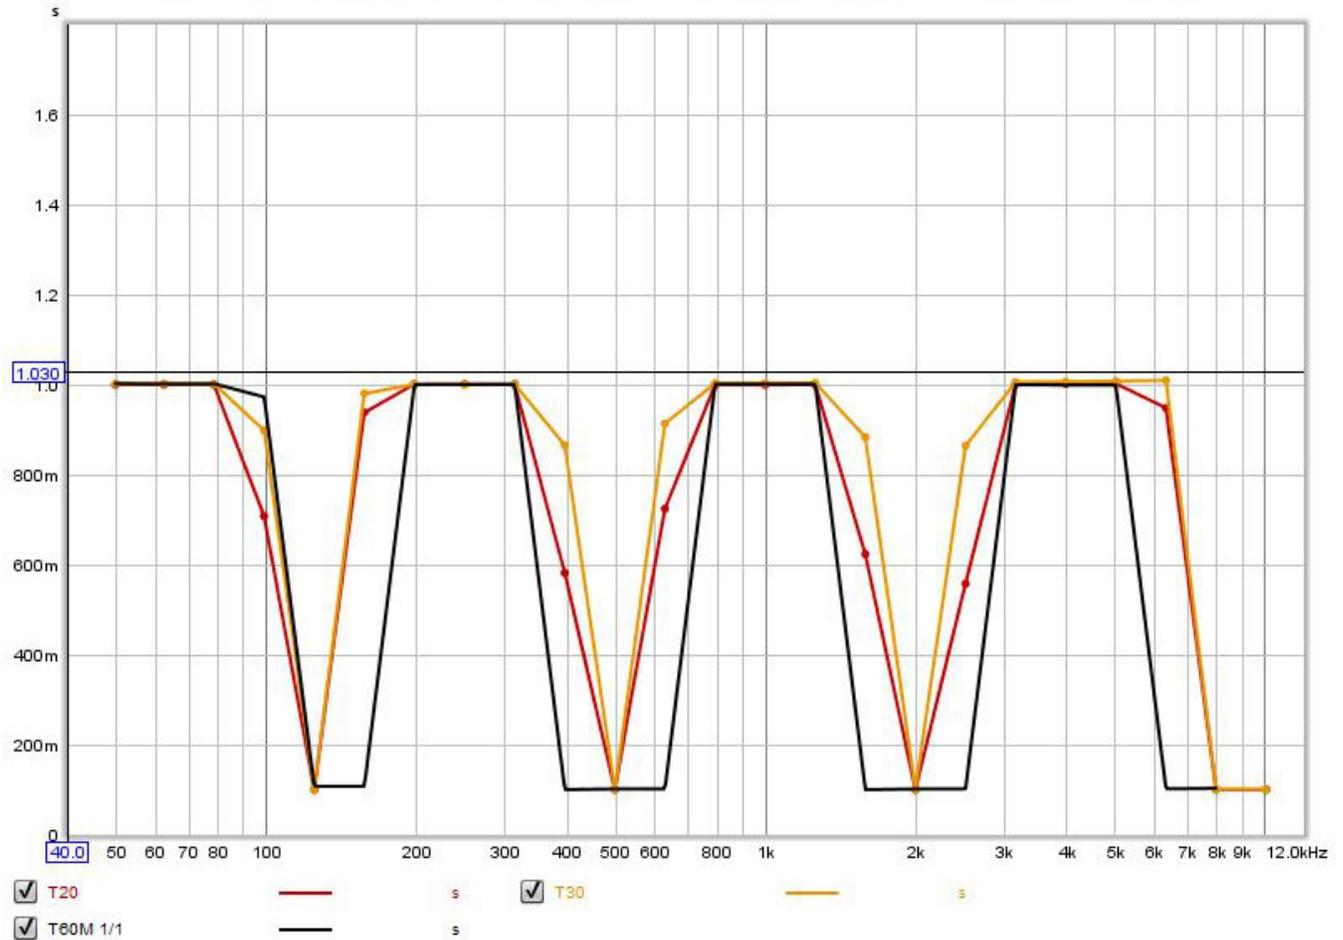
Accuracy comparison

The next test is of accuracy. Another synthetic test signal was generated, again consisting of equal amplitude exponentially decaying cosines at the one octave centre frequencies, but with 60 dB decay times that varied in 0.2 s steps from 1.600 s at 63 Hz to 0.200 s at 8 kHz. Noise was added at two different levels, -80 dB to act as a clean signal reference and -40 dB to test performance in the presence of noise, in both cases relative to the cosine amplitudes. Here are the spectra of the test signals.

Accuracy signal spectrum, noise at -80 dB

Accuracy signal spectrum, noise at -40 dB
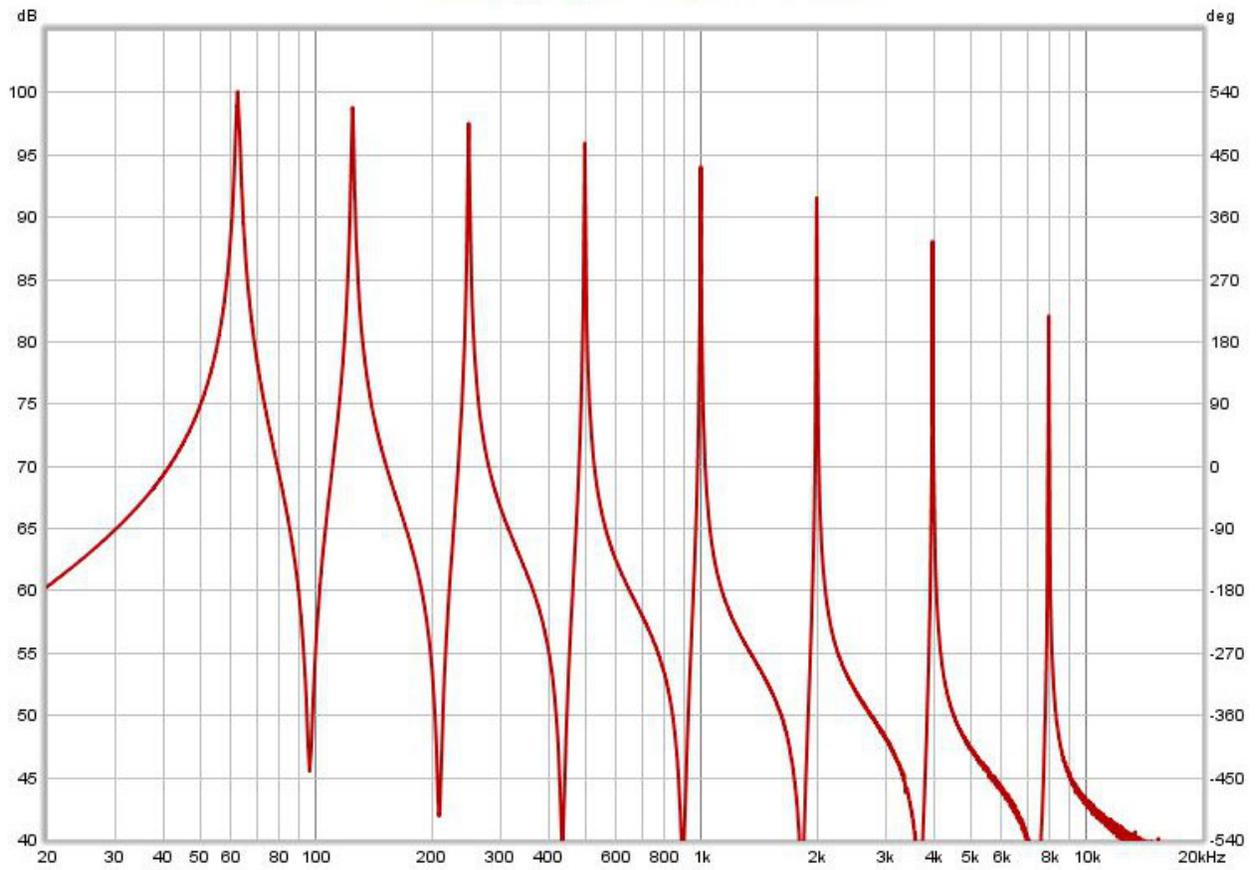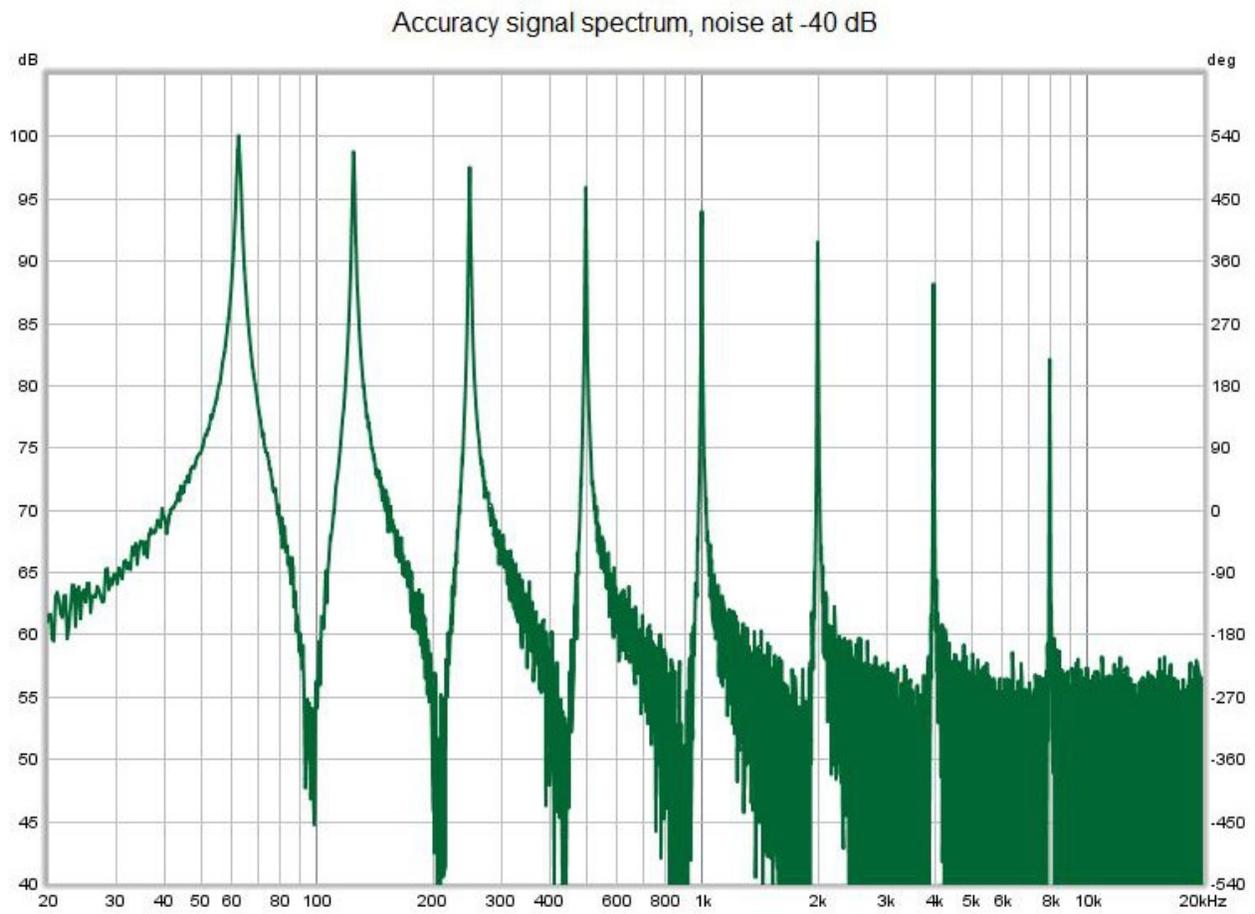
The signal with noise at -80 dB presented no difficulties for the new or the classical approaches.

Accuracy test, noise at -80 dB

☑ T20 (r: -1.000)	——	0.800 s	☑ T30 (r: -1.000)	——	0.800 s
☑ T60M 1/1	——	800 ms			
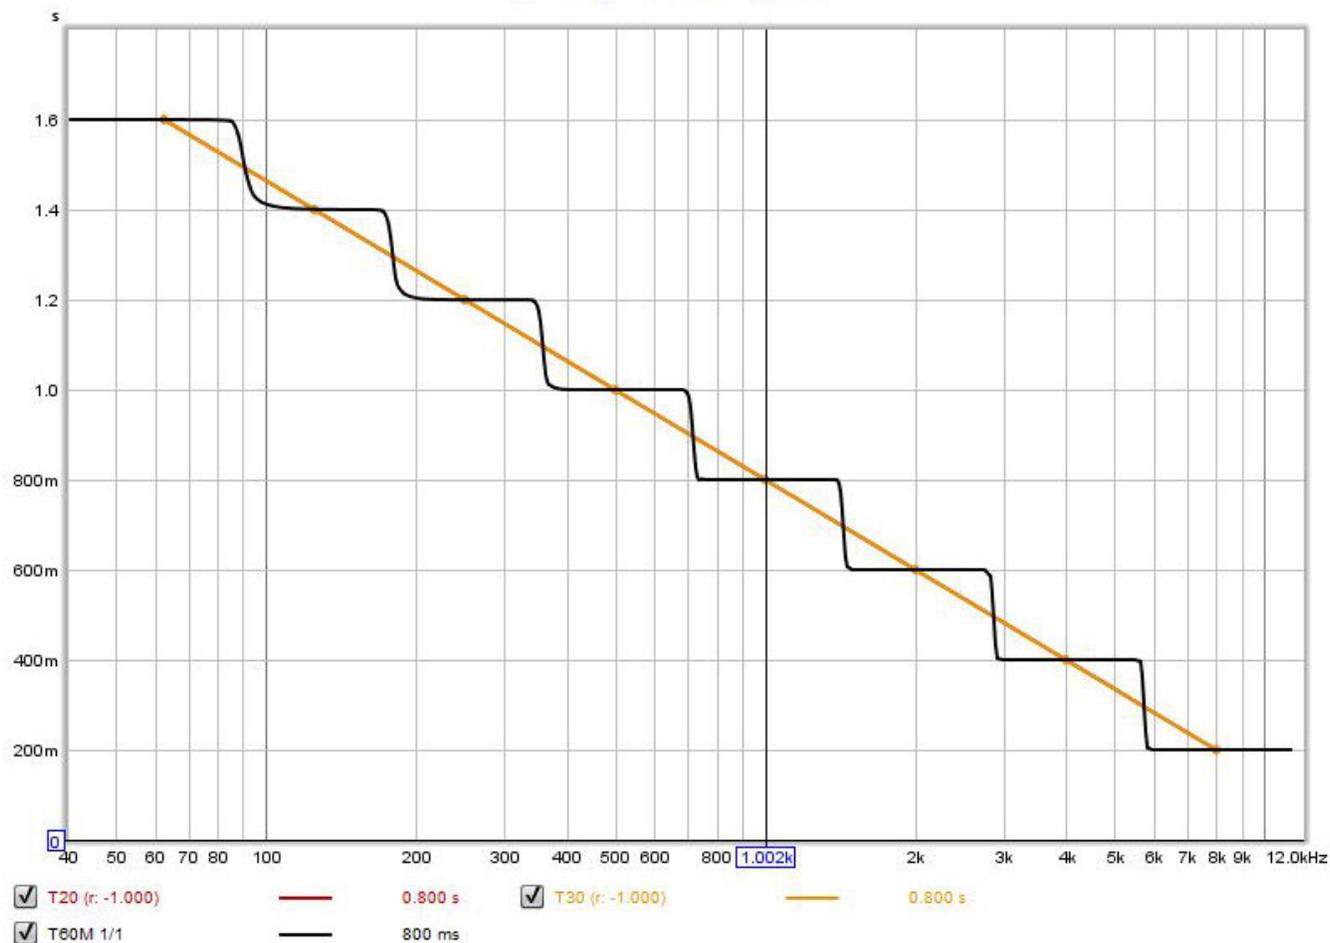
The signal with noise at -40 dB was more challenging for the classical approach, despite successful detection of the noise floor (using Lundeby's method) in all cases. Filter order made negligible difference, 18th order providing a very slight improvement. The results below are for 18th order filters. The T60M frequency domain approach performed significantly better than the classical approach with this signal.

Accuracy test, noise at -40 dB

☑ T20 (r: -1.000)	———	0.822 s	☑ T30 (r: -0.999)	———	0.869 s
☑ T60M 1/1	———	795 ms			
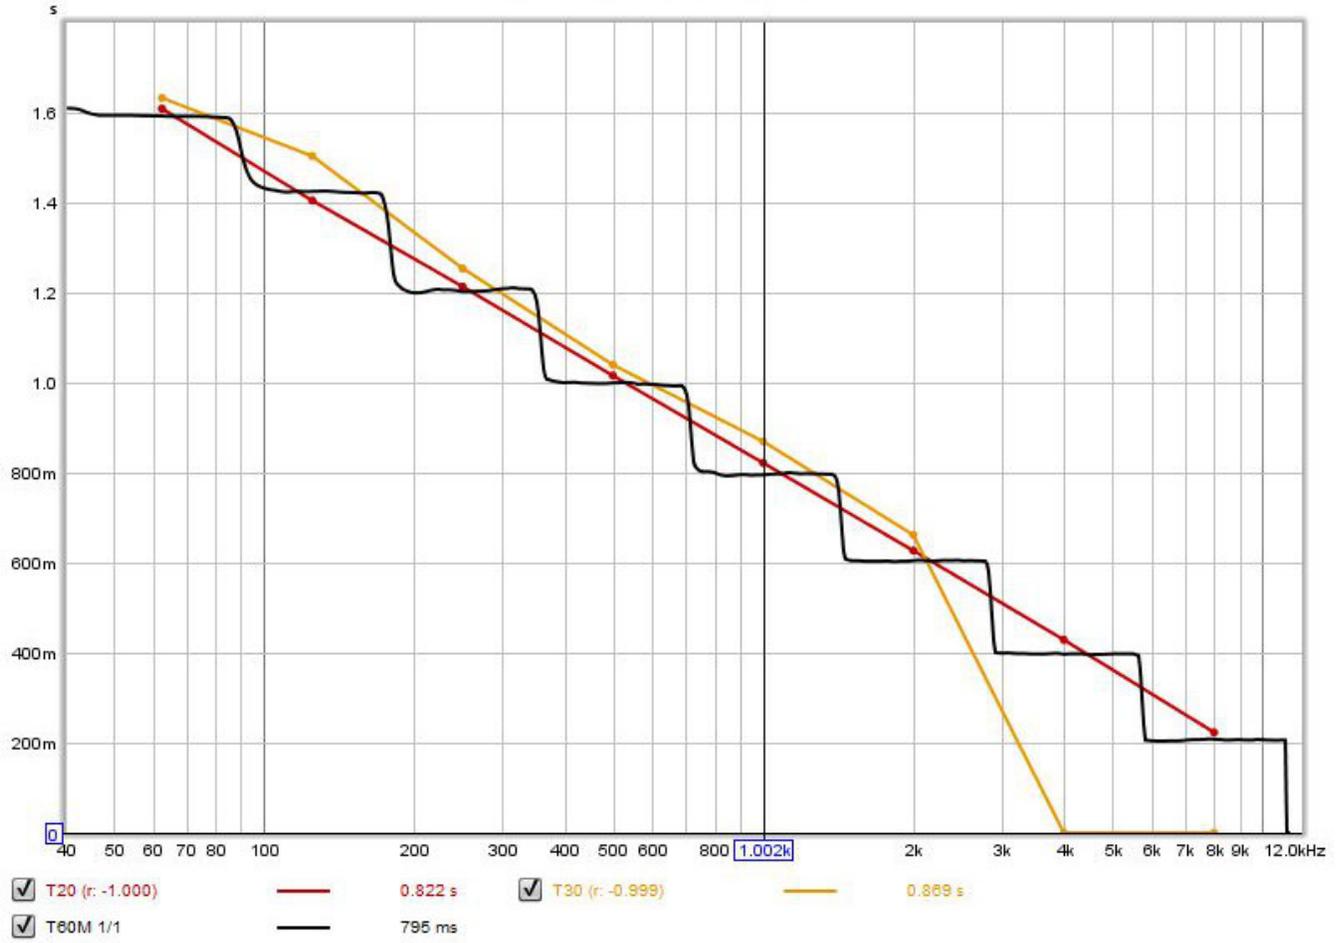
Here are examples of the RT60 Decay graphs for both cases.

RT60 Decay graph, noise at -80 dB

1 kHz 1/1		Classic: 1 kHz 1/1
Model T60	0.800	EDT 0..-10 dB 0.801 s, r = -1.000
Level (dB)	0.4	T20 -5..-25 dB 0.804 s, r = -1.000
Noise (dB)	-73.0	T30 -5..-35 dB 0.805 s, r = -1.000
Model fit	0.00	Topt -5..-26 dB 0.804 s, r = -1.000

1,210 ms range, 500 ms window, 24 ms rise time, 5.0 ms steps, 1.9 Hz resn, t = 1,210 ms

☑ Octave modes -80dB N   -1.7 dB   ☑ T60M 1/1   800 ms
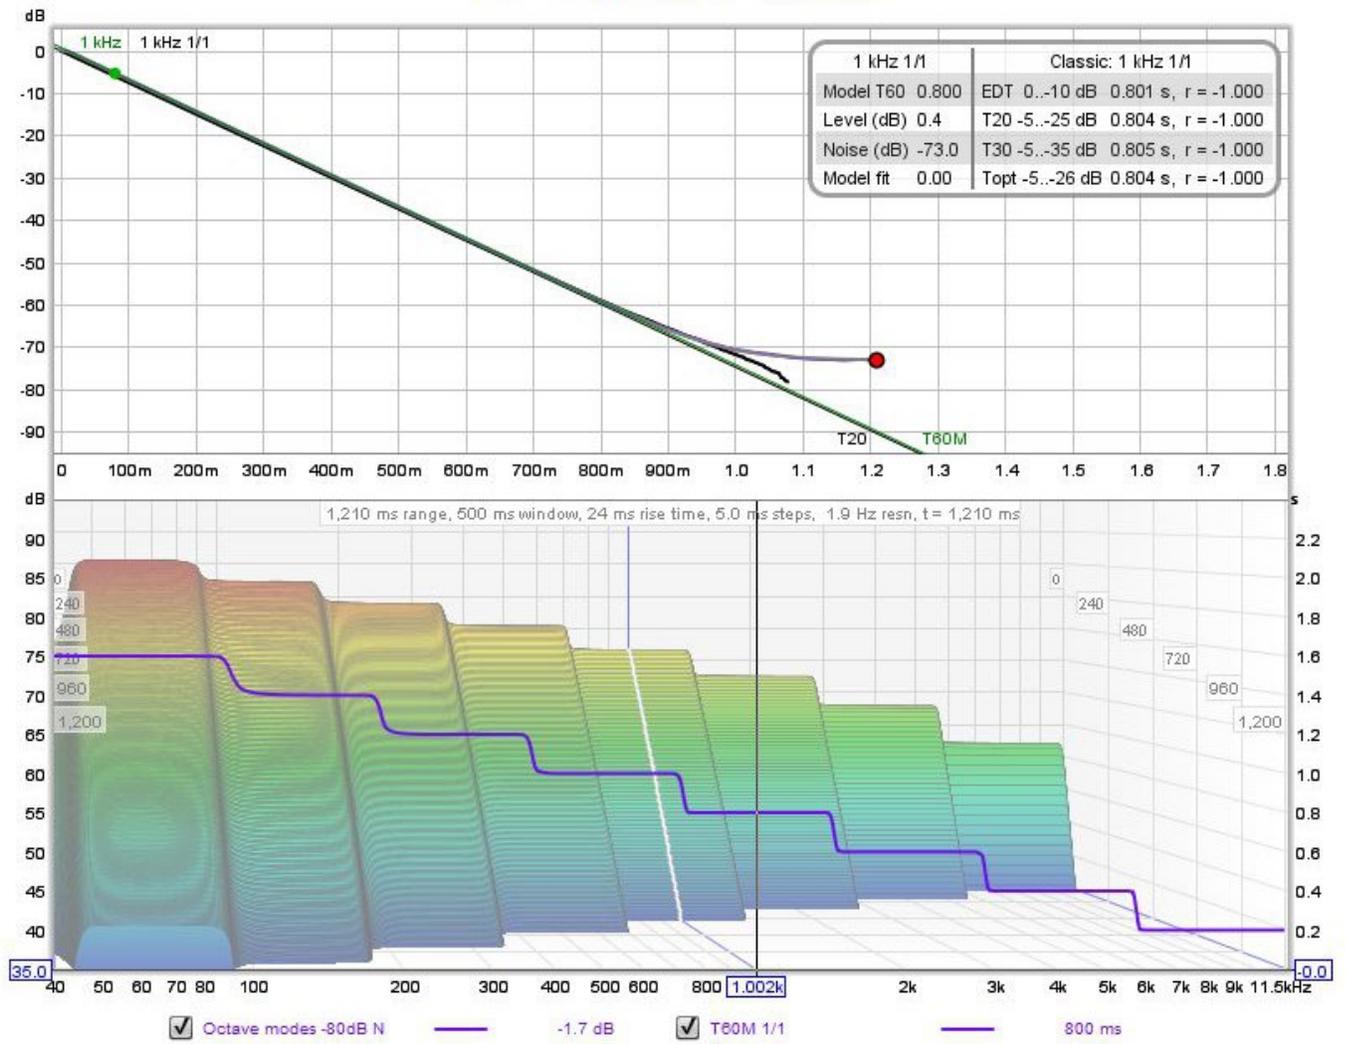
RT60 Decay, noise at -40 dB
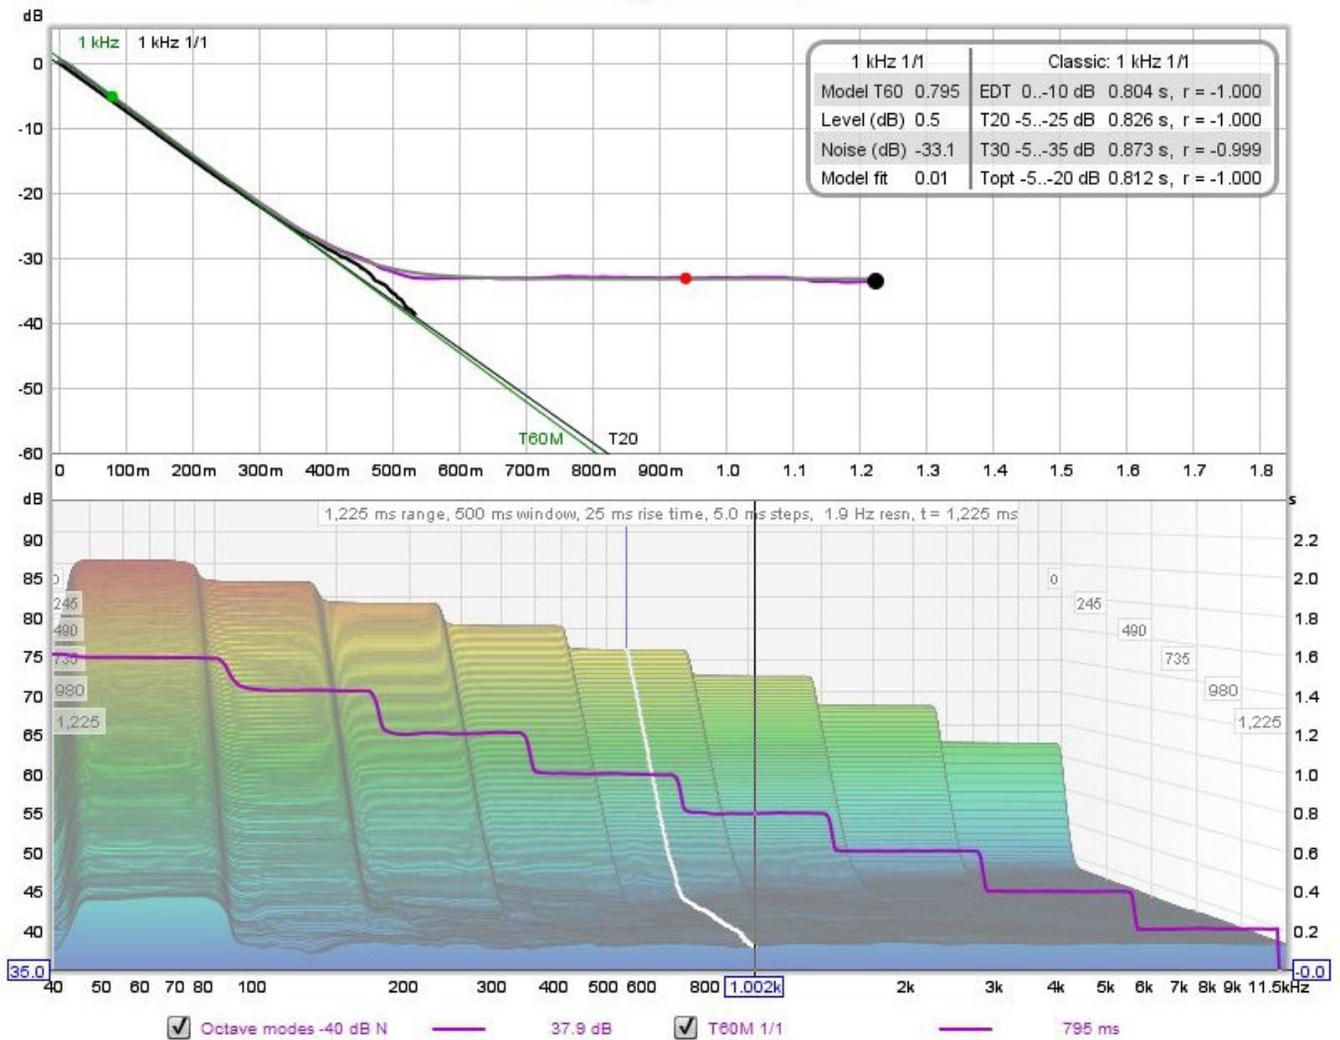
Implementation

Here are some technical details of REW's implementation of the RT60 Decay analysis.

Decay data

The decay data is derived from a series of Fourier Transforms, referred to as slices. The first slice uses the peak of the impulse response as the reference point, each subsequent slice is shifted in time by the slice interval. The total time span covered by the slices is the "Range" setting. Each slice is windowed with a left side Hann window of width equal to the rise time and a right side Hann window equal to the "Window" width. Note that at high frequencies the noise floor of the frequency domain approach is typically higher than the classical approach, due to the FFT window width required for the low frequency range. Narrower windows and shorter rise times could be used to generate high frequency decay data with lower noise floors. A frequency-dependent windowing (FDW) approach could be considered, but FDW implementations typically have symmetric windows which would result in excessive rise time at low frequencies.

The **Automatic** settings derive the time range by carrying out a conventional T20 analysis using 1 octave 6th order filters at 250, 500 and 1000 Hz. The longest T20 figure is used as the time range, with a minimum range of 500 ms. The slice interval is set to provide about 250 slices over the range, with a maximum slice step of 100 ms. The rise time is set to 2% of the range, with a maximum of 400 ms.

The data for each slice is low pass filtered to 1/48th octave then sampled at 96 points per octave to provide a log spaced data set. Subsequent octave band filtering is applied to the squared magnitudes of that log spaced data, using a kernel that is unity across the band width except at each end, where the kernel amplitude is 0.5.

RT60 analysis

The data series formed by the values of each slice at a frequency is analysed for the RT60 calculation. An initial scan of the data is carried out to determine if there are sharp roll offs or rises near the end of the data, if found subsequent data are excluded from the analysis.

The analysis consists of fitting a decaying exponential plus noise function to the data. The function is:

$$f(t) = level*exp(-6*log(10)*t/t60) + noise$$

The three model parameters `level`, `t60` and `noise` are determined by minimising an error function formed by the difference between the function value and the decay data at the slice times. REW uses the mean squared value of the difference between the dB decay data and the dB values of the model function, weighted by an error scaling function that gives less weight to values that are below -20 dB. Initial estimates for the model parameters are based on a least squares line fit to the decay data over the range from -10 to -20 dB.

The model error is evaluated over a sub-section of the decay data. The nominal start point for the model fit is the -5 dB point on the decay curve, but at least 10 ms or the rise time after the decay start, and with an overall limit of not being more than 10 dB down the decay curve. It should be borne in mind that, unlike a Schroeder integral, the decay data is typically not monotonic, so "-5 dB point" refers to the last point in the response where it passes through -5 dB. REW's adaptation of Lundeby's method is applied to the decay squared magnitude data to estimate the noise floor and the intersection of the noise floor and the RT60 decay line, the fit end index is then set twice as far into the decay curve as that intersection point. If the lundeby routine doesn't find a good result the fit end is the end of the valid data.

A Java port of UNCMIN (by Steve Verrill) is used to find the optimal model parameters. UNCMIN is an unconstrained optimiser, constraints are imposed by modifying the minimiser vector entries before they are used in the model function. That constraint process is part of a transformation of parameter values from the optimiser space to the function space. The parameters in the minimiser space corresponding to `level` and `noise` are scaled dB equivalents of the model values. The minimiser parameter corresponding to `t60` is a scaled reciprocal of the T60 value. Limits are placed on the values based on the initial estimates.

Good selectivity at low frequencies requires rise times that are typically larger than the slice interval. The effect of the left window on the decay data must be taken into account before the model fit error is evaluated. The left window is convolved with the decay data, but not with the decay values at each slice themselves but rather with the differences between the decay values for successive slices, since it is that difference in the decay the window is acting upon.

Help Index

Clarity Graph

The C50, C80 and D50 clarity and definition curves at each octave or one-third octave filter centre frequency are displayed on this graph. See below for descriptions of each of these parameters.
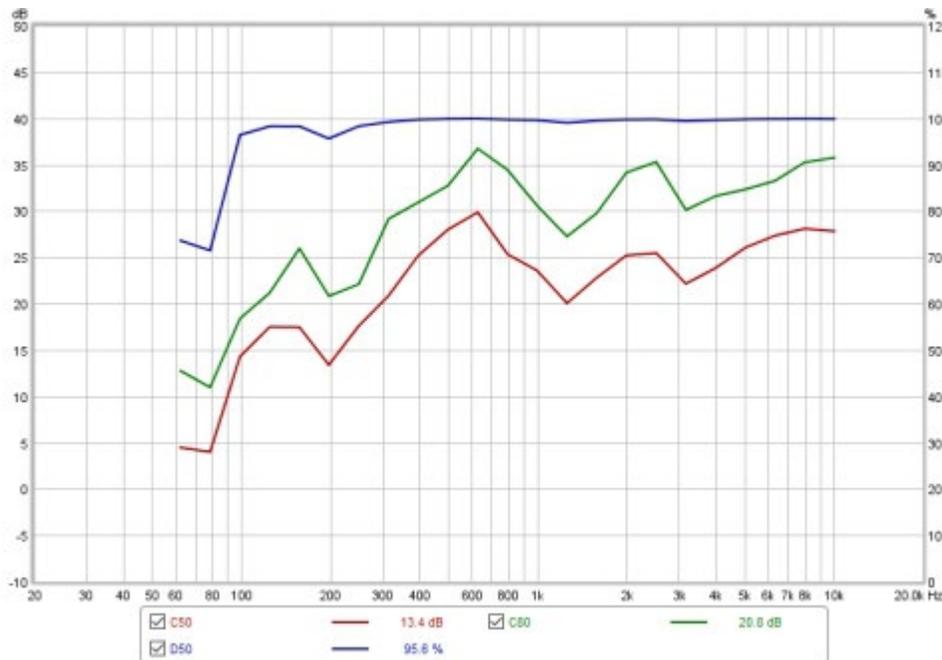
Controls

The control panel for the Clarity graph has these controls:

The band selector chooses either octave or one-third octave filtered results. The filtering settings used are those on the RT60 graph.

The plot can show horizontal bars centred on each filter frequency and spanning the filter's bandwidth, or lines joining the filter centre frequencies, according to the **Use bars on plot** control setting.

The **Trace options** button brings up a dialog that allows the colour and line type of the graph traces to be changed. If a change is made it will be used for all measurements shown on this graph. Traces can also be hidden, which will remove them from the graph and from the graph legend.

The clarity measures available are:

Clarity C50

The early to late energy ratio in dB, using sound energy in the first 50 ms as the 'early' part. C50 is most often used as an indicator of speech clarity.

Clarity C80

The early to late energy ratio in dB, using sound energy in the first 80 ms as the 'early' part. C80 is most often used as an indicator of music clarity.

Definition D50

The early to total energy ratio as a percentage, using sound energy in the first 50 ms as the 'early' part

The parameter values (RT60 and clarity) for the current measurement can be written to a text file using the File → Export → **RT60 data as text** menu entry.

Help Index

Spectral Decay Graph

This graph shows spectral decay traces over the region from 10Hz to the end of the measurement sweep. The plot used logarithmically spaced data at 96 points per octave with 1/48th octave smoothing applied. The Spectral Decay plots are generated by shifting the impulse response window to the right by the slice interval to generate each succeeding slice. Two windows are used, a left side window to taper the data prior to the start of the region being analysed and a right side window that spans the selected window width. The default window type for the left side is Hann, for the right side it is Tukey 0.25, other types may be selected via the Spectral Decay entries in the [Analysis](#) Preferences. The initial reference point for the windows (end of left window/start of right window) is the peak of the impulse response.

The plot is generated automatically when the graph is selected, or may be manually regenerated using the **Generate** button in the bottom left corner of the graph area.

Decay Controls

The time separation of the slices is controlled by the **Slice Interval** setting, the width of the impulse response section that is used to generate the slice is set by the **Window** control. The corresponding frequency resolution is shown to the right.

The **Rise Time** control sets the width of the Left Hand window. Shorter settings give greater time resolution but make the frequency variation less easy to see. The default setting, 100 ms, is aimed at revealing room resonances. When examining drive unit or cabinet resonances with full range measurements a much shorter rise time would be used, 1.0 ms or lower, with time spans and window settings of around 10 ms. CSD mode is often more useful for such measurements as the later part of the impulse response can be noisy, obscuring the behaviour in the later slices.

The **Smoothing** applied to the slices can be increased from 1/48th octave (the minimum, and recommended) to as high as 1/3rd octave.

Use CSD Mode should be selected if the later slices of the decay are contaminated by noise in the measurement. It would commonly be used when examining drive unit or cabinet resonances. CSD mode anchors the right hand end of the window at a fixed point and only moves the left side. This does mean, however, that the frequency resolution reduces (and the lowest frequency that can be generated increases) as the slices progress, as each has a slightly shorter total window width than the previous slice.

The traces for each slice can be drawn as conventional lines or as filled areas, selected by the **Fill slices** check box. The alternative views are shown below.
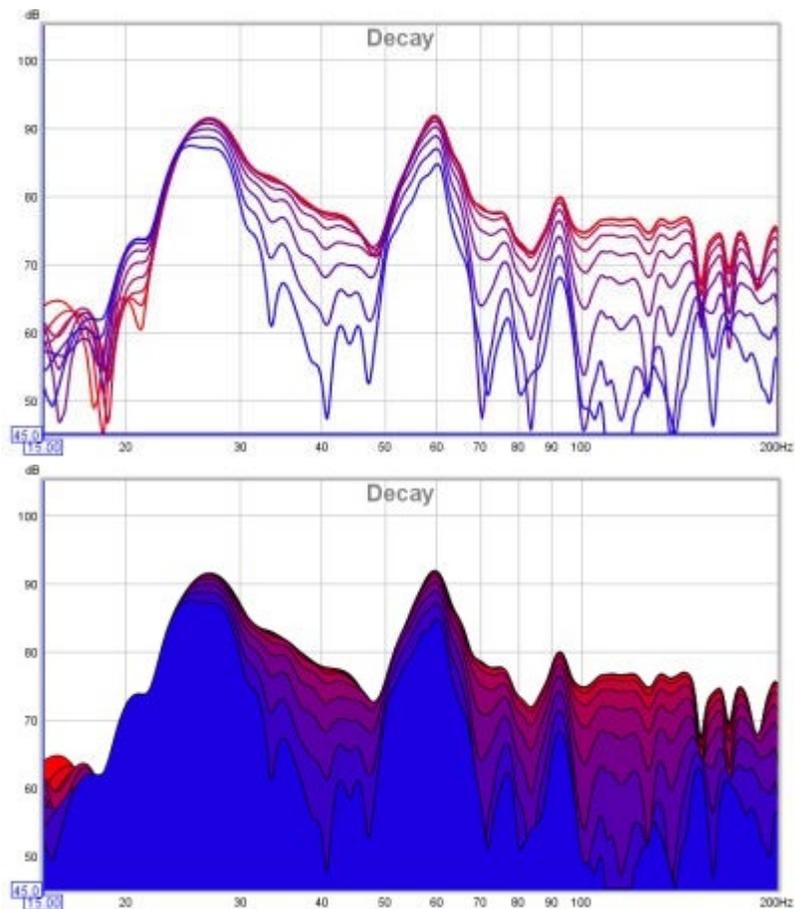
If **Show modal frequencies** is selected the theoretical modal frequencies for the room dimensions entered in the Modal Analysis section of the EQ Window are plotted at the bottom of the graph.

The control settings are remembered for the next time REW runs. The **Apply Default Settings** button restores the controls to their default values.

Help Index

Waterfall Graph

This graph shows a waterfall plot over the region from 10Hz to the end of the measurement. It can be used to view the results of sweep measurements, imported audio files or stepped sine measurements for which the spectrum data has been captured at each measurement frequency. Measurement and audio file plots can be generated in either Fourier or Burst Decay modes. A Fourier plot is classical waterfall, produced by sliding a window along the impulse response or data (a more detailed explanation is provided below). A Burst Decay plot shows how tone burst decay along a decay axis of periods of the decay frequency, making it easier to compare the Q of resonances. The plot is generated automatically when the graph is selected, or may be manually regenerated using the **Generate** button in the bottom left corner of the graph area.
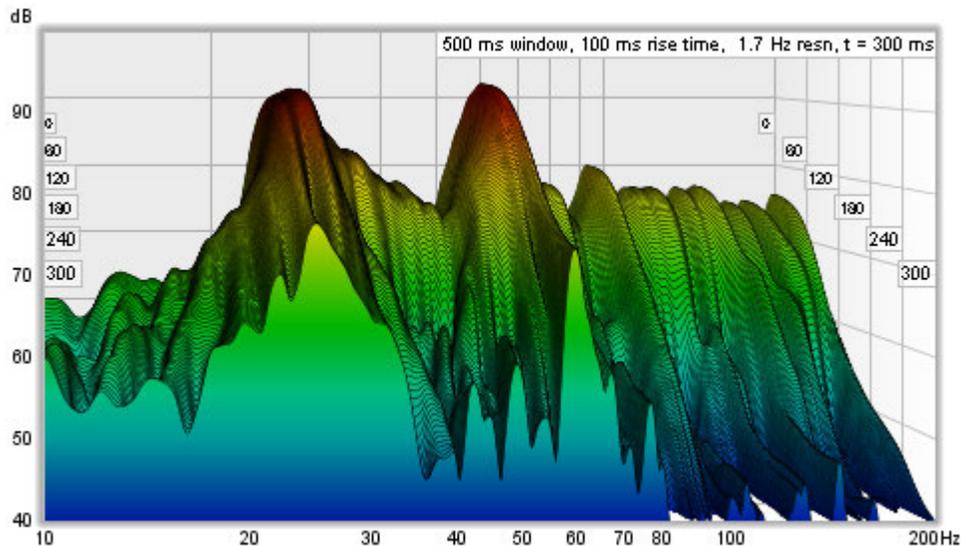
The labels at the sides of the plot show the decay axis values, in either time or, for Burst Decay, periods.
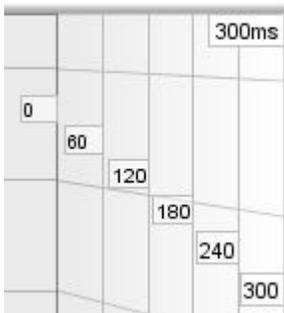
How a Fourier Waterfall Plot is Generated

To understand what the Fourier waterfall plot shows and how its appearance is affected by the various waterfall controls it is helpful to first understand how it is generated.

Each slice of the waterfall plot shows the frequency content of a windowed part of the measurement's impulse response. 'Windowed' means we take the impulse response and multiply each sample in it by the value of a window, which is made up of a left side and a right side whose shapes we can choose (the window types are selected via the Spectral Decay entries in the Analysis Preferences). Here is an example of an impulse response showing the original impulse, the window shape (in blue) and the windowed response.
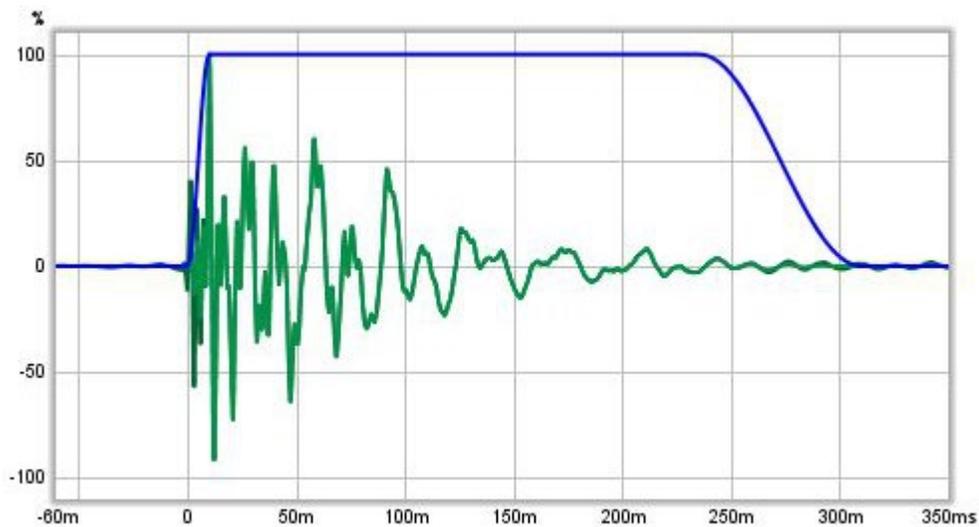
Here is a zoomed in view of the early part, where the effect the windowing has on the windowed (lighter green) trace can be seen.
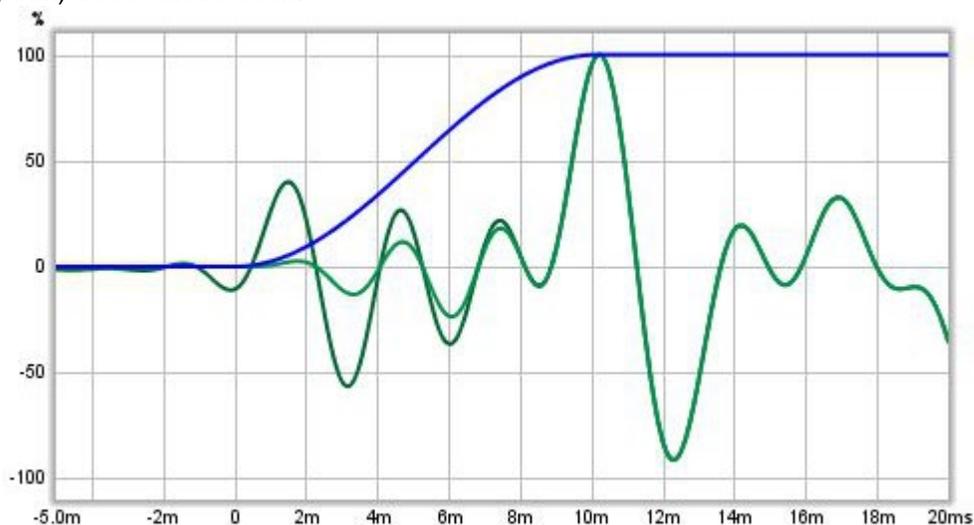
After the frequency content of the first windowed part of the impulse response has been obtained it is plotted as the first slice of the waterfall. The window is then moved along the response and the process is repeated for the next slice. The amount the window moves is determined by the time span of the waterfall and the number of slices that are to be plotted, so that the data for the last slice is from a section of the impulse response that is later than the first slice by the time range - for example, if the time range was 300 ms and there were 51 slices there would need to be 50 shifts of the window (the first slice has no shift) so each slice would be from data obtained after moving the window 6 ms along the impulse (300/50).

The window has a left hand side and a right hand side. In the plots above, the left hand window is a Hann type that ends at the peak of the impulse. The right hand side is a Tukey 0.25 (which means that for 75% of its width it is flat, then the remaining 25% is a Hann window). The overall width of the window (left side plus right side) determines the frequency resolution of each slice of the waterfall. The shape of the window, and particularly the shape and width of the left hand side, affects the way features of the response are smeared out in time.

To understand this, imagine a rectangular window and a perfect impulse, that has one sample at 100% and all other samples zero. As long as that single 100% sample is within the span of the window the frequency response will be a flat line. As soon as the left edge of the window goes past the 100% sample that slice and all slices after it will have no data in them (all the samples will be zero) so the waterfall would disappear off the bottom of the plot. Here is an example of such a waterfall plotted with a 100 ms left hand rectangular

window.
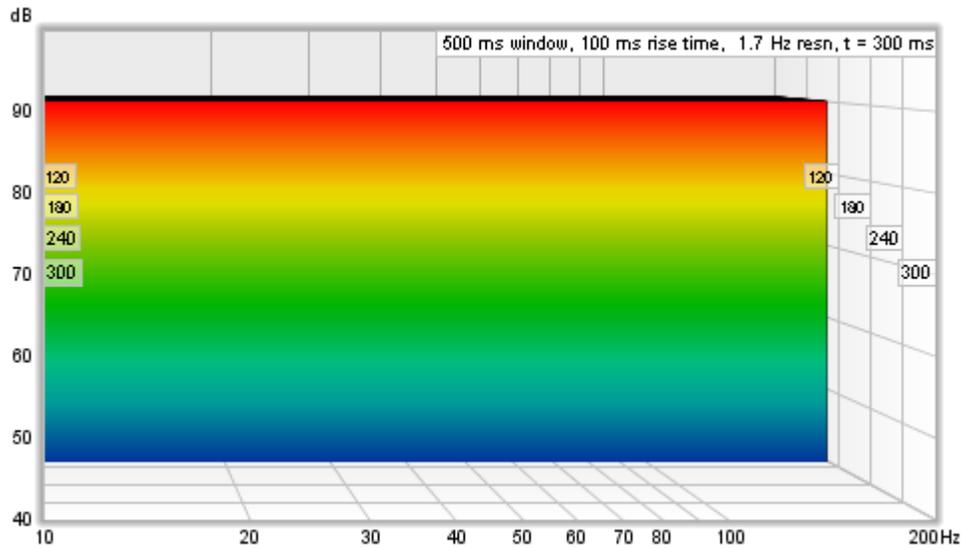
That waterfall is, in the time domain, a faithful representation of how that perfect impulse response looks - and in general for any response a rectangular window gives the best time resolution, but that comes at a price. The price is in the frequency domain behaviour, i.e. the shape of the frequency response in slices of the waterfall. In real impulse responses, that are spread out over time, using a rectangular window creates a sharp step at the left hand edge of the windowed data. That sharp step causes ripples in the frequency response, obscuring the actual frequency content. The waterfall also has an initial period, equal to the width of the left hand window, where the slices are almost identical, creating a flat portion. Here is an example of a measurement with a 100 ms rectangular left hand window, despite its appearance it is the same measurement as shown at the top of this help page, only the shape of the left hand window has been changed.
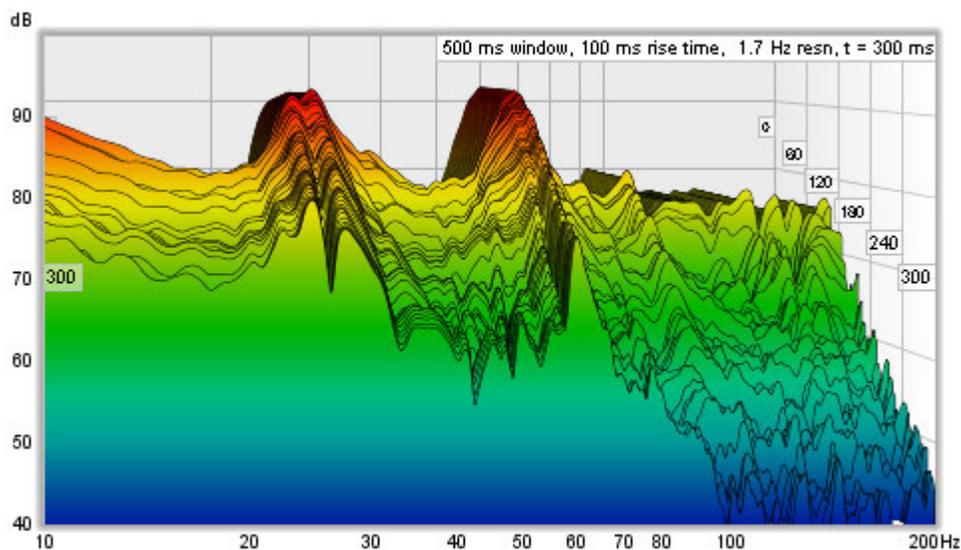
To avoid the damaging effects of that sharp step in the windowed response, a tapered window is used to smoothly attenuate the samples, but now a feature that does actually have a rapid change in the impulse response will linger on in the waterfall, because it will not entirely disappear until the whole left hand window has gone past it. Here are the perfect impulse and the real measurement again, this time with a 100 ms Hann left hand window.
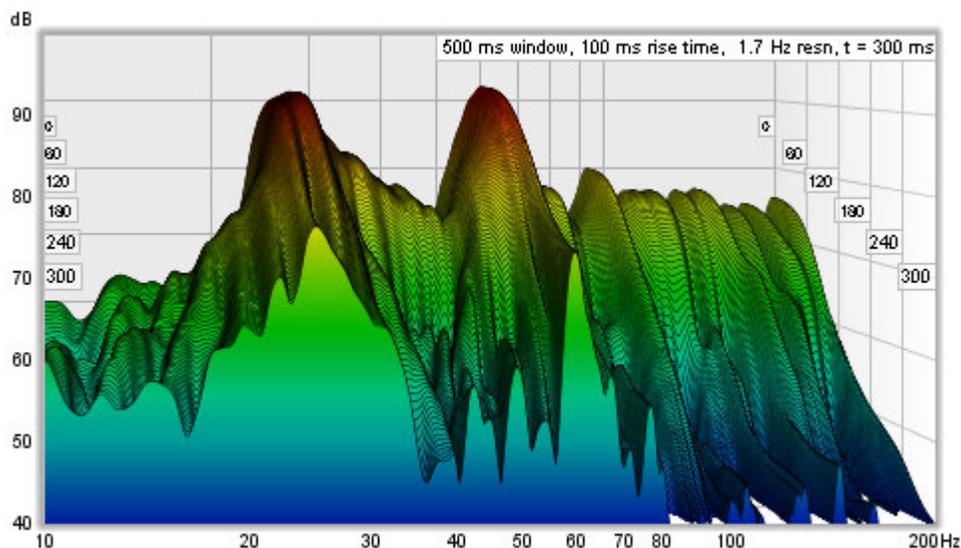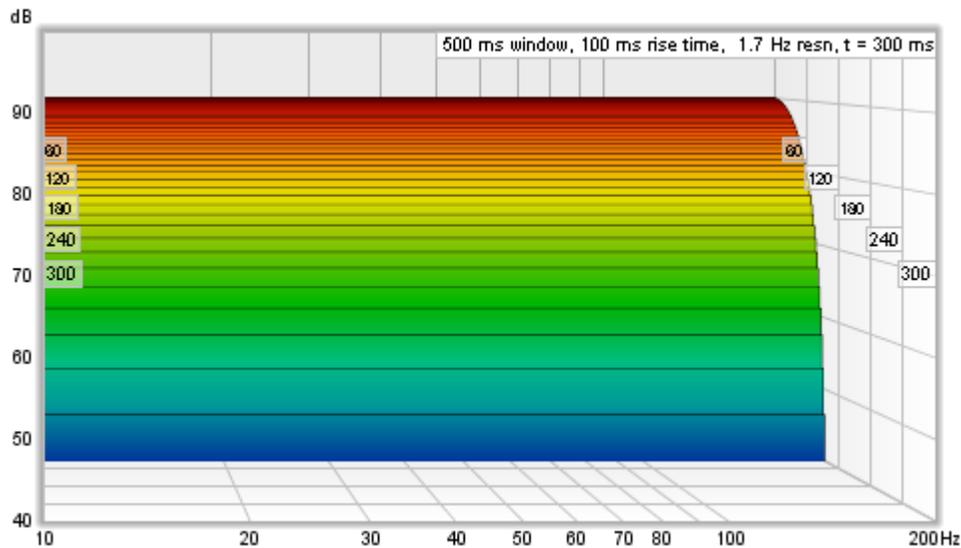
REW's sweep measurement waterfalls are aimed at examining room resonances. To help make those resonances easy to see in the response, a wide left hand window is used. The left hand window width is specified independently, using a setting labelled **Rise Time**. Changing the **Window** setting only alters the Right Hand window, which means that the **Window** setting now controls only the frequency resolution of the waterfall - longer settings give higher resolution - without altering the waterfall's time domain behaviour. There are also controls to select how many slices the waterfall should have and to select the smoothing to apply to each slice.

In addition to the standard waterfall mode, which slides the window along the impulse response, there is a **CSD** (Cumulative Spectral Decay) mode, which anchors the right hand end of the window at a fixed point and only moves the left side, which may be useful when examining cabinet or tweeter resonances over very short time spans if the IR data descends into the noise floor soon after the region being examined. Using CSD mode in those cases prevents the later slices from including increasing amounts of noise floor. This does mean, however, that the frequency resolution reduces (and the lowest frequency that can be generated increases) as the slices progress, as each has a slightly shorter total window width than the previous slice. Note also that it does not make sense to have a time range greater than the window width in CSD mode as the window, with its fixed right hand edge, reaches zero width after stepping along by a time interval equal to the window width and there will be no data for subsequent slices. In CSD mode window width should be greater than the time range. CSD mode is often not required for measurements that have good signal to noise ratio, not using it allows frequency resolution to be maintained throughout the time range of interest.

How a Burst Decay Waterfall Plot is Generated

The **Burst Decay** mode provides a way to more easily distinguish resonances of similar Q but different frequencies. It does this by showing the way a shaped tone burst at each frequency would decay, but along an axis denoted in periods of the frequency rather than time. On a period axis the extent of a decay is the same for resonances of the same Q regardless of the frequency of the resonance.

The Burst Decay plot is produced by convolving the windowed impulse response (using whatever current window settings have been applied to the measurement) with a complex Morlet wavelet analytic signal (a Gaussian windowed complex exponential) to extract the decay envelope then resampling that decay on a period-based scale. The bandwidth of the wavelet can be chosen to be either 1/3rd octave or 1/6th octave. The 1/3rd octave choice favours time resolution, the 1/6th octave choice favours frequency resolution. The convolution is repeated at 48 points per octave across the frequency span of the measurement, with 10 Hz as the lowest frequency allowed and the burst bandwidth below half the sample rate as the highest. Note, however, that an artefact of the period axis is to skew the tail of the decay slightly towards higher frequencies rather than maintaining the symmetry about the resonance's centre frequency that would be seen in a time-based plot.

The slices in a Burst Decay waterfall are separated by an interval equal to the selected span of the plot in periods divided by the number of slices minus 1. For example, a plot spanning 40 periods with 201 slices has a slice interval of 0.2 periods. At 100 Hz that is 2 ms, at 10 kHz just 20 us. It is that varying of the time scale over which the decays are shown which ensures resonances of the same Q decay at the same rate per slice. Here are examples of a Fourier and Burst Decay waterfall of the same system. The Fourier waterfall spans 500 ms, the Burst Decay spans 30 periods, which is 500 ms at 60 Hz where the system has a resonance. It is clear from the Burst Decay that the 60 Hz resonance is higher Q than the 27 Hz resonance.

Burst Decay

Waterfall Controls

Mode selects the type of spectrogram plot that will be produced, which can be either **Fourier** or **Burst Decay**.

The **Total Slices** control determines how many slices are used to produce the waterfall. Fewer slices mean faster processing and lower memory use, but make it less easy to see how the response is varying. The actual number of slices in a Fourier plot may be up to 20% fewer than the total slices requested to allow the slice interval to be an integer number of samples, which speeds up processing.

The width of the impulse response section that is used to generate a Fourier waterfall is set by the **Window**

control (this control sets the Right Hand window width). The corresponding frequency resolution is shown to the right of the window setting. Longer window settings provide better frequency resolution.

The **Time Range** control determines how far the impulse response window is moved from its start position to generate the waterfall.

The **Rise Time** control sets the time duration of the Left Hand window. Shorter settings give greater time resolution but make the frequency variation less easy to see. The default setting, 100 ms, is aimed at revealing room resonances. When examining drive unit or cabinet resonances with full range measurements a much shorter rise time would be used, 1.0 ms or lower, with time spans and window settings of around 10 ms. CSD mode may be more useful for such measurements as the later part of the impulse response can be noisy, obscuring the behaviour in the later slices. The 'rise time' terminology dates back to the late 80s and MLSSA. In MLSSA it referred to the 10% to 90% rise time of a left hand window formed by convolving a window function with a unit step (in essence, the step response of the chosen window function). The actual width of the window was much greater, depending on the window type - about twice the rise time for a Hann window, for example, or about 3 times the rise time for Blackmann-Harris. In REW the term is used to refer to the overall width of the left hand window, somewhat misusing it in the interest of retaining terminology that is in common use for CSD plots whilst adopting a definition that provides a clearer indication of what parts of the response lie inside and outside the chosen window settings. To obtain similar results to the MLSSA-style definition use an REW setting that is twice as long.

Use CSD Mode should be selected if the later slices of the waterfall are contaminated by noise in the measurement. It would commonly be used when examining drive unit or cabinet resonances. CSD mode anchors the right hand end of the window at a fixed point and only moves the left side. This does mean, however, that the frequency resolution reduces (and the lowest frequency that can be generated increases) as the slices progress, as each has a slightly shorter total window width than the previous slice. In CSD mode window width should be greater than the time range (otherwise the window width would be zero at times in the range later than the width of the window).

In **Burst Decay** mode a **Bandwidth** control allows burst bandwidths of 1/3rd or 1/6th of an octave to be selected. The 1/3rd octave choice favours time resolution, the 1/6th octave choice favours frequency resolution. With the 1/6th octave burst resonances are more easily distinguished. With the 1/3rd octave burst reflections become more evident in the plot, showing up as curved lines. Burst decay plots may have visible artefacts near the high frequency limit, though typically more than 40 dB below the peak level. A **Periods** control sets the number of periods the plot will span.

If **Auto generate** is selected the waterfall will be regenerated automatically if a setting is changed, otherwise new settings will not be applied until the **Generate** button is pressed.

If **Show modal frequencies** is selected the theoretical modal frequencies for the room dimensions entered in the Modal Analysis section of the EQ Window are plotted at the bottom of the graph.

The **Smoothing** applied to the waterfall slices can be increased from 1/48th octave (the minimum, and recommended) to as high as 1/3rd octave.

Settings can be copied and pasted between measurements by right clicking on the controls panel.

The control settings are remembered for the next time REW runs. The **Apply defaults** button restores the controls to their default values.

Perspective settings
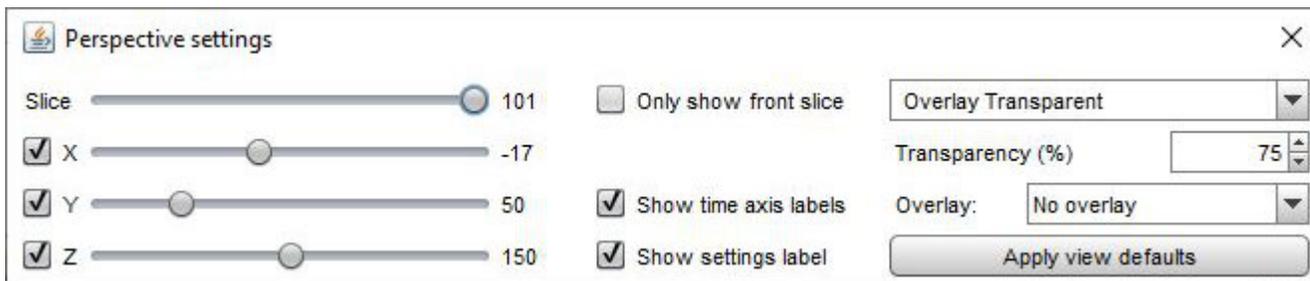
The **Slice** slider selects which slice is at the front of the plot - as the slider value is reduced the plot moves forward one slice at a time. The trace value shows the SPL figure for the front-most slice, the corresponding time for that slice is shown at the top right of the graph.

The x, y and z sliders alter the perspective of the plot, moving it left/right, up/down and forwards/backwards respectively. The check boxes next to the sliders allow the perspective to be disabled in that axis. Disabling the x axis can make it easier to see the frequencies of peaks or dips. Disabling the z axis turns off all the perspective effects which makes the plot like a filled spectral decay. Here is the same plot as above but with the x-axis perspective effect turned off.
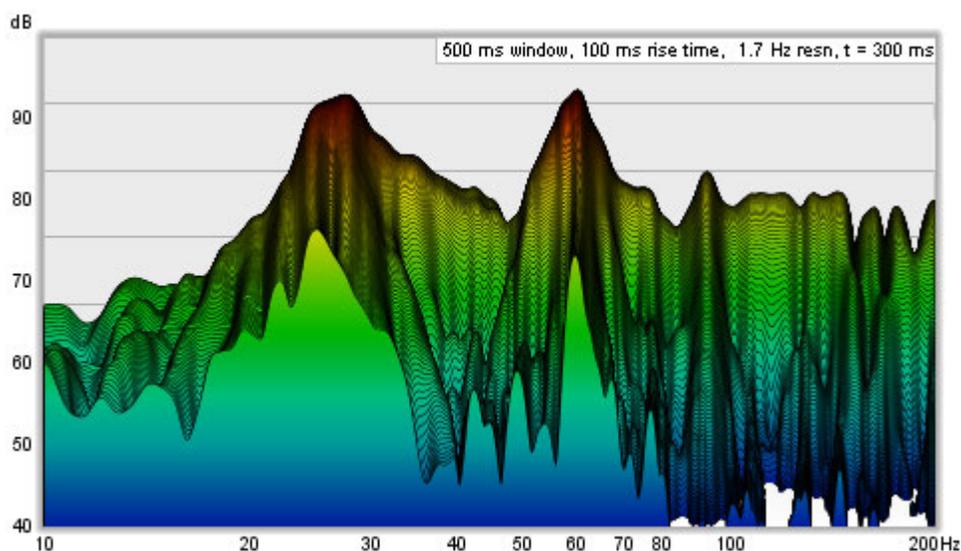
If **Only show front slice** is selected only the slice at the front of the waterfall will be shown.

Show time axis labels controls whether the time (or periods) axis grid lines on the sides of the plot are labelled with the corresponding value. **Show settings label** controls whether to show a label with the plot settings in the top right corner of the plot.

The waterfall allows another measurement's plot to be overlaid on the current measurement. The overlay is generated slice-by-slice, plotting a slice of the current measurement's waterfall, then a slice of the overlay, then the next slice of the current measurement and so on. For this to be possible the waterfalls must cover the same time range and have been generated from measurements with the same sample rate. Overlays are easier to distinguish with the colour scheme set to None. **N.B.** before a measurement is available to overlay it is necessary to generate the waterfall data for it.
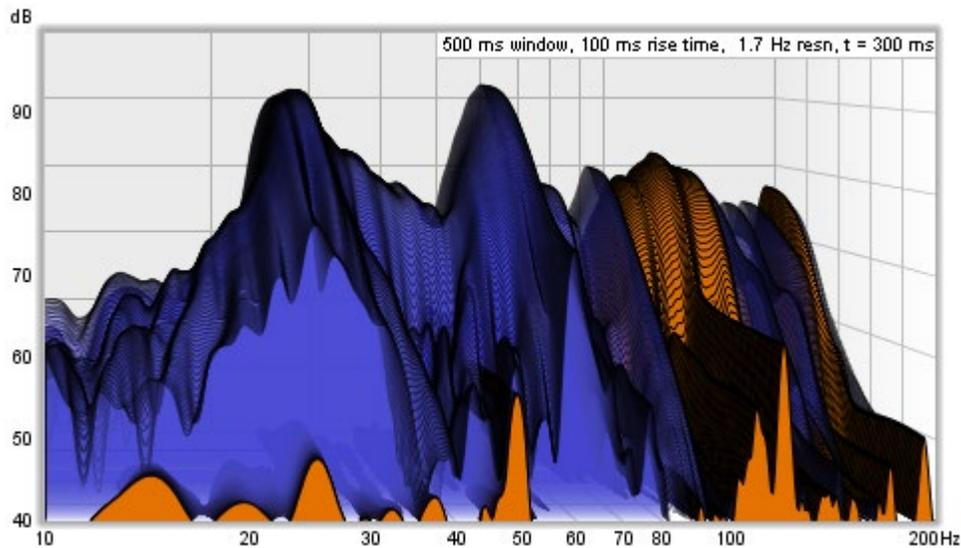
The overlay is selected using the **Overlay** selector. Measurements which do not have waterfall data are shown in grey in the selection list. To generate the data for a measurement select it as the current measurement and use the **Generate** button.

Transparency can be applied to the main plot, the overlay, or both. When transparency is set to 0% both plots are solid. In the image above the main plot is drawn at 75% transparency, allowing the overlay to show through. The transparency mode can be switched between main/overlay/both to ease comparison between the plots.

Colour scheme settings

Waterfalls can be generated with the slices filled in the same colour as the measurement or with a colour gradient that varies with SPL. The **Scheme** selector is set to None to use the measurement colour or to the chosen scheme. Here is an example using the Heat colour scheme:
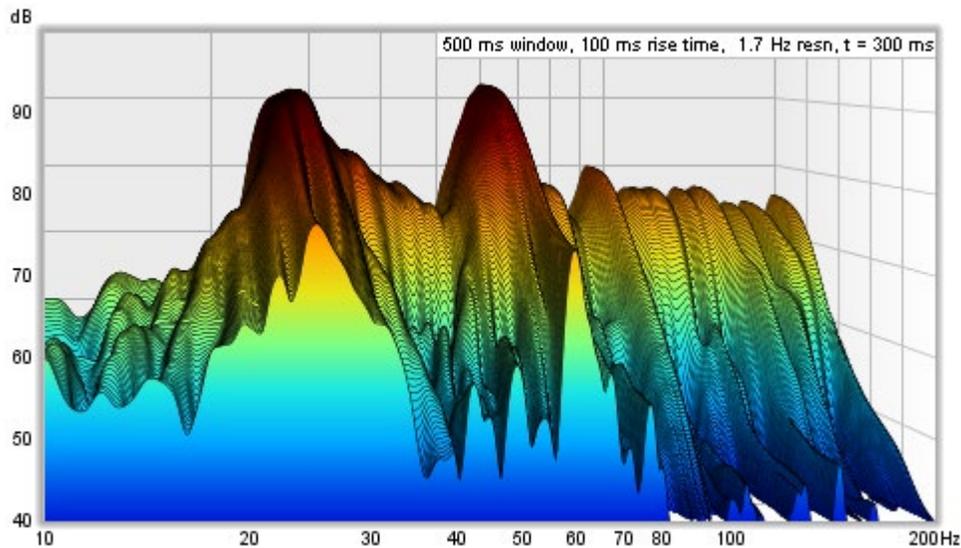
The **Top**, **Bottom** and **Range** controls adjust how the plot colours correspond to the values in the waterfall data. Any values higher than the Top are drawn in the colour at the top of the scale, any values lower than the Bottom are drawn in the colour at the bottom. If the Top setting is changed the Bottom will be adjusted to keep the same Range. If the Bottom is changed the range will be adjusted to keep the same Top. If the Range is changed the Bottom will be adjusted, keeping the same Top.

If **Banded colours** is selected the colour scale has discrete steps rather than a continuous blend from one colour to another - there are 11 colours in that case to provide 10 bands across the scale range.

Show colour scheme bar controls whether a bar showing the SPL range of the selected colour scheme is shown to the right of the plot.

Match top of scale to peak adjusts the Top value so that it corresponds to the highest level found in the data.

Waterfall Controls - Audio Files

The controls are slightly different for imported audio files. There is no Time Range control, the waterfall is generated for the full range of the imported file. There is no rise time control. A single, centred window is used to generate the waterfall, the window type used is selected from the controls.

Waterfall Controls - Stepped Sine Measurements
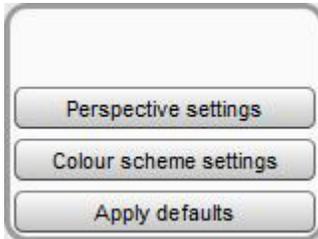
Stepped sine measurements have a reduced set of controls and some additional perspective options that are only used for stepped sine data. Each slice in a stepped sine measurement waterfall shows the spectrum data for a test frequency.

The **Slice** and X, Y, Z perspective controls work in the same way as for sweep measurements. **Only show front slice** hides all slices except the frontmost, which can be selected using the **Slice** slider - this is a convenient way to view the spectrum for a single test frequency. The test frequency for the current slice is shown in the top right corner of the graph and on the left and right side walls. **Reverse slices** switches the order of the slices with the exception of the noise floor, which is the last (frontmost) slice in the set.

Note that waterfalls can only be generated for stepped sine measurements that had the option to **Capture spectrum data at each frequency** selected.
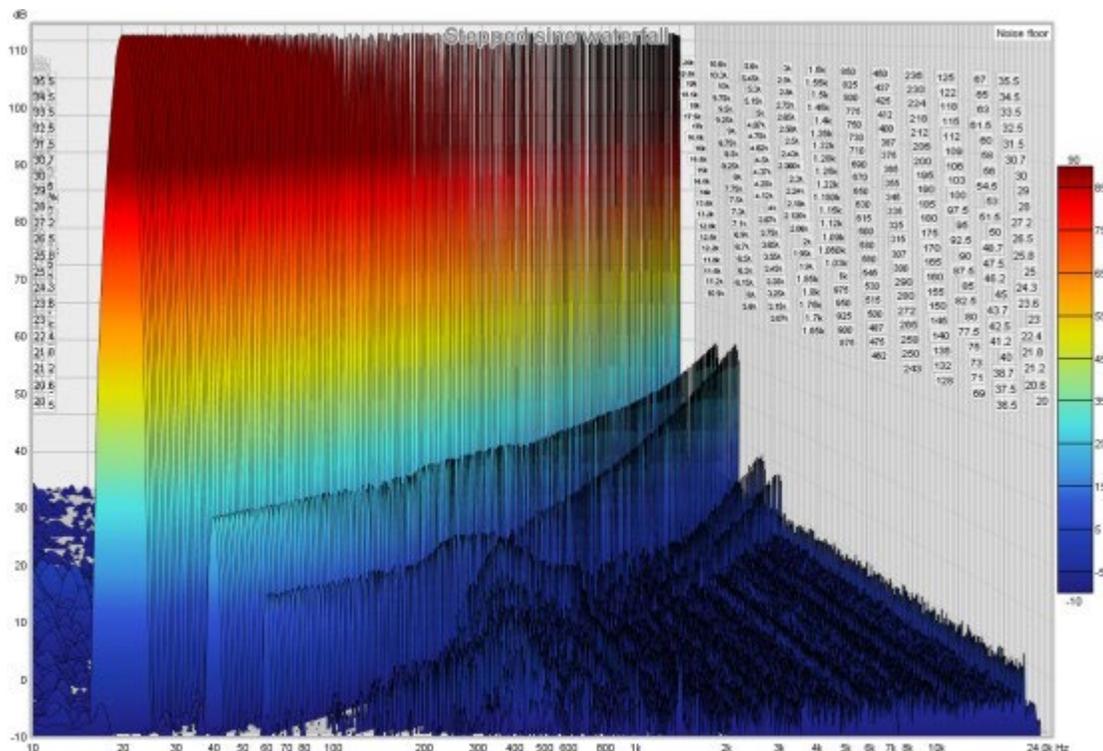
Help Index

Spectrogram Graph

This graph shows a spectrogram plot over the region from 10 Hz to the end of the measurement. It can be used to view the results of sweep measurements, the frequency content of imported audio files or the results of stepped sine measurements for which the spectrum data has been captured at each measurement frequency.

The spectrogram is like a waterfall viewed from above, with the level indicated by colour. The scale showing how colour relates to level is displayed to the right of the plot. In Fourier or Wavelet mode the vertical axis of the plot can show time, increasing towards the top of the plot, or frequency with time on the horizontal axis. In Burst decay mode the time axis is replaced by a periods axis.

When viewing sweep measurements in Fourier or Wavelet mode the time starts before the peak of the impulse so that the onset of the response can be seen. The areas where the response is decaying more slowly show up as streaks along the time axis. The dashed line is the **Peak energy time** trace which shows the peak level in the plot at each frequency. This can highlight variations in peak energy arrival versus frequency - an ideal peak energy time trace would be a straight line with the same time value for all frequencies.
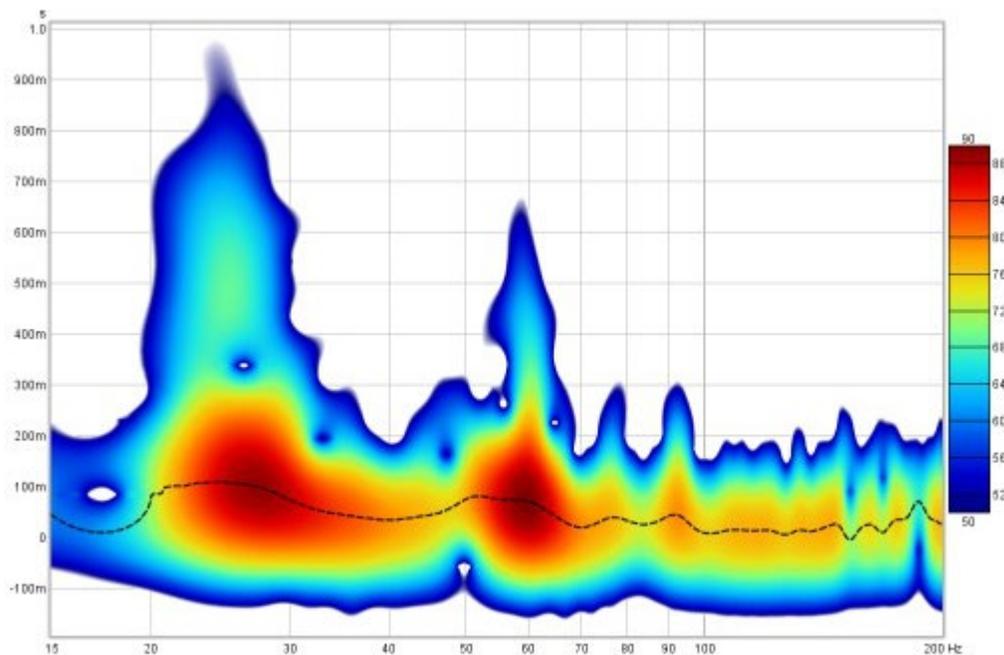
The spectrogram plot in Fourier mode is generated in the same way as the Spectral Decay plot, shifting the impulse response window to the right by a proportion of the time range to generate each succeeding slice. The window type is selected in the graph controls. The plot uses logarithmically spaced data at 96 points per octave.

In Burst decay mode the decay starts from the peak of the impulse and is shown along an axis of the number of periods for each frequency, so that resonances with the same Q show the same rate of decay. Comparing the Fourier spectrogram above with the 1/6th octave Burst Decay spectrogram below it is evident from the Burst Decay that the 60 Hz resonance has a higher Q than the 27 Hz resonance, though the 27 Hz resonance has a longer decay time.
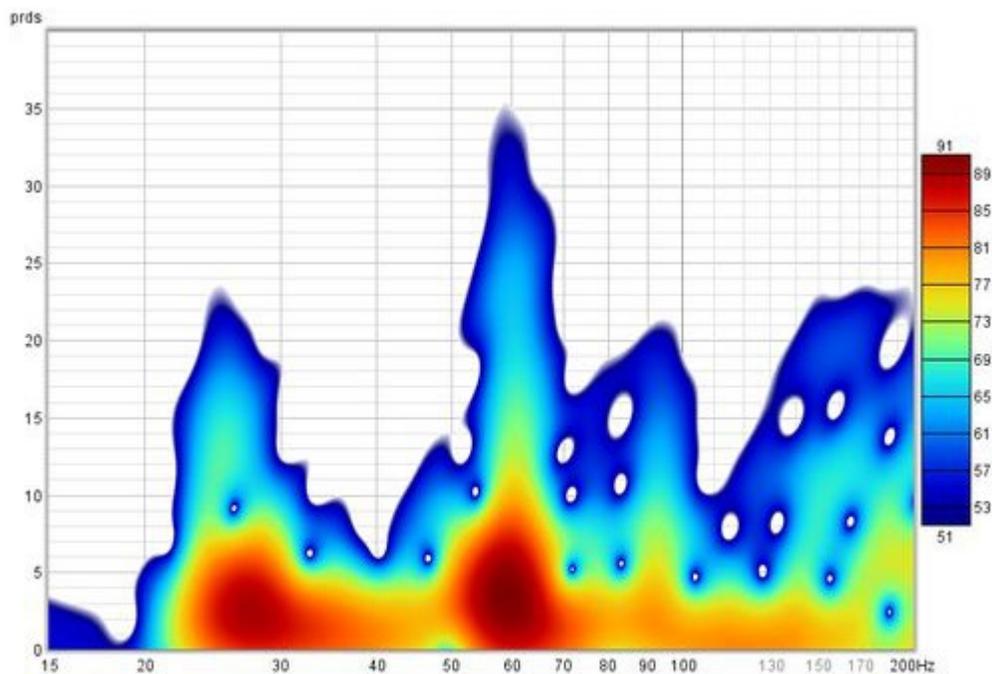
The spectrogram plot is generated automatically when the graph is selected, or may be manually regenerated using the **Generate** button in the bottom left corner of the graph area (shortcut Alt+G). The legend panel shows the plot value at the intersection of the vertical and horizontal cursor lines.

An ideal Spectrogram decays very rapidly off the bottom of the scale range. Here is an example of a plot produced from a soundcard loopback measurement in **Fourier** mode.
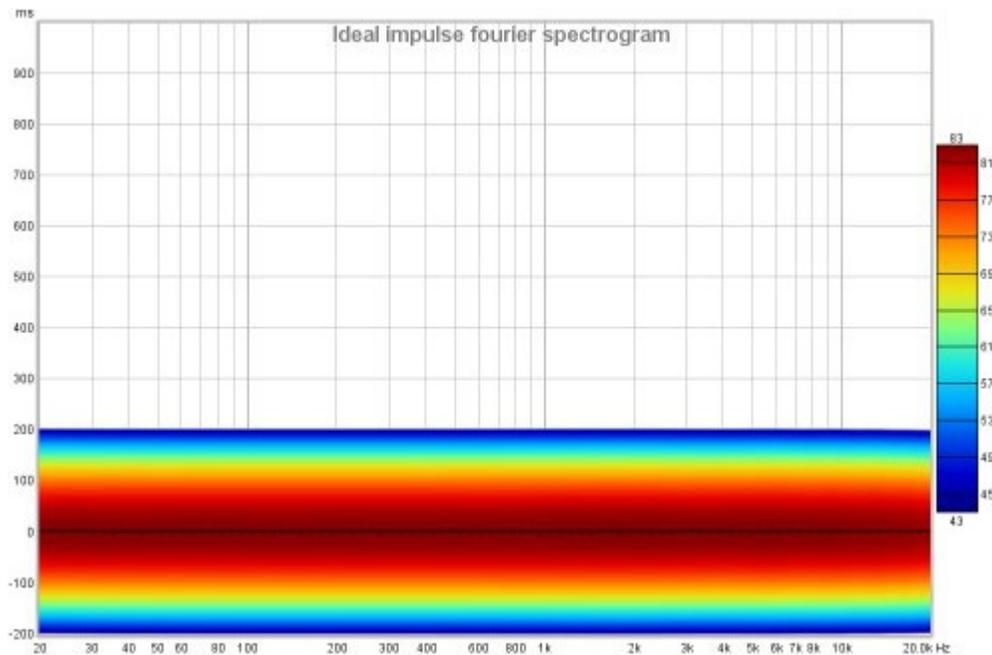
Spectrogram Controls - Sweep Measurements

Mode selects the type of spectrogram plot that will be produced, which can be either **Fourier**, **Wavelet** or **Burst Decay**. In Fourier mode the plot uses fixed width windows, which mean the plot has the same time resolution at all frequencies. If the plot spans a wide range of frequencies this usually means the time resolution is either too low at high frequencies or too high at low frequencies. A 100 ms window, for example, gives 10 Hz frequency resolution. At low frequencies that is a big octave fraction (1/1.4 octaves at 20 Hz), at high frequencies a very, very small octave fraction (1/1386 octaves at 20 kHz). For a time-frequency plot it would be more useful if the trade-off between time and frequency resolution varied with frequency, using a constant octave fraction for frequency resolution rather than a constant number of Hz and so giving higher time resolution at high frequencies and lower at low frequencies. A wavelet transform can achieve that, specifically a constant Q Continuous Wavelet Transform (CWT). A constant Q wavelet transform is mathematically equivalent to using a frequency-dependent window to produce the spectrogram, which is what REW does. This method is faster than typical CWT calculations, but may produce some artefacts in parts of the response that extend to frequencies close to half the sample rate - using a higher sample rate shifts these beyond the usual range of interest.

Here is a 1/6 octave **Wavelet** spectrogram of the same soundcard loopback measurement shown above. It becomes narrower as frequency increases, reflecting the increasing time resolution of the wavelet plot.
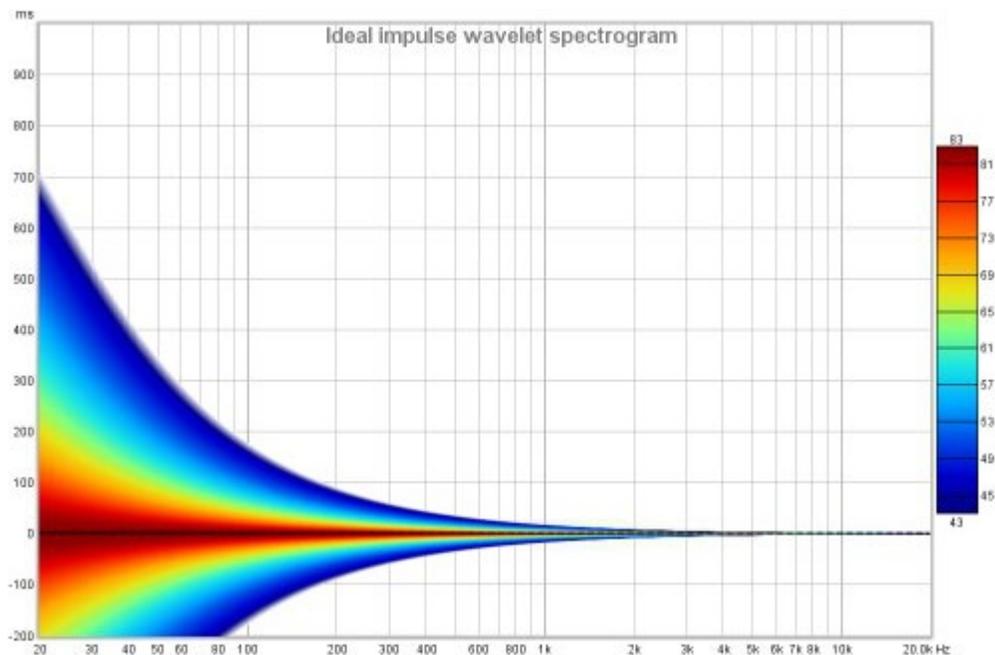
Here is the same measurement from the first image above as a 1/12 octave **Wavelet** spectrogram.
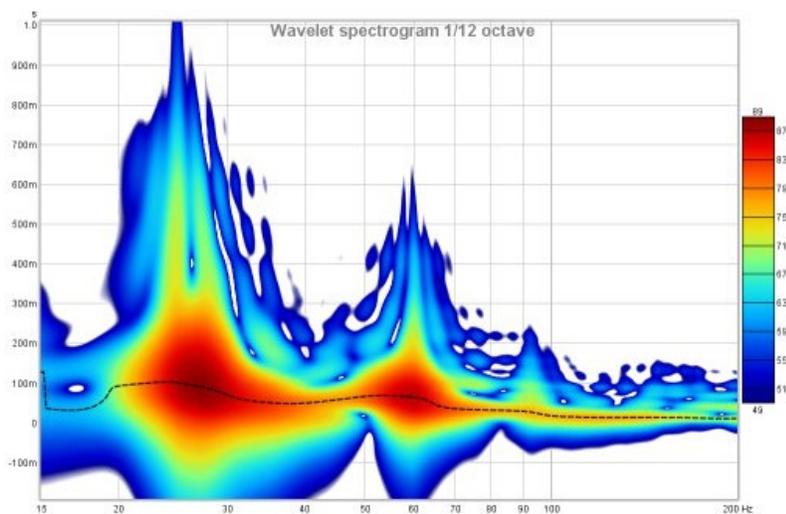
The difference between the Fourier and Wavelet spectrograms can be more easily seen when looking at responses with reflections. Here are two plots of a response which has a series of reflections at 1 ms intervals after the peak. In the Fourier spectrogram, using a 10 ms window and a 10 ms span after the peak, the effect on the frequency response and decay are clearly visible, with peaks at 1 kHz intervals. However, the reflections themselves cannot be distinguished.
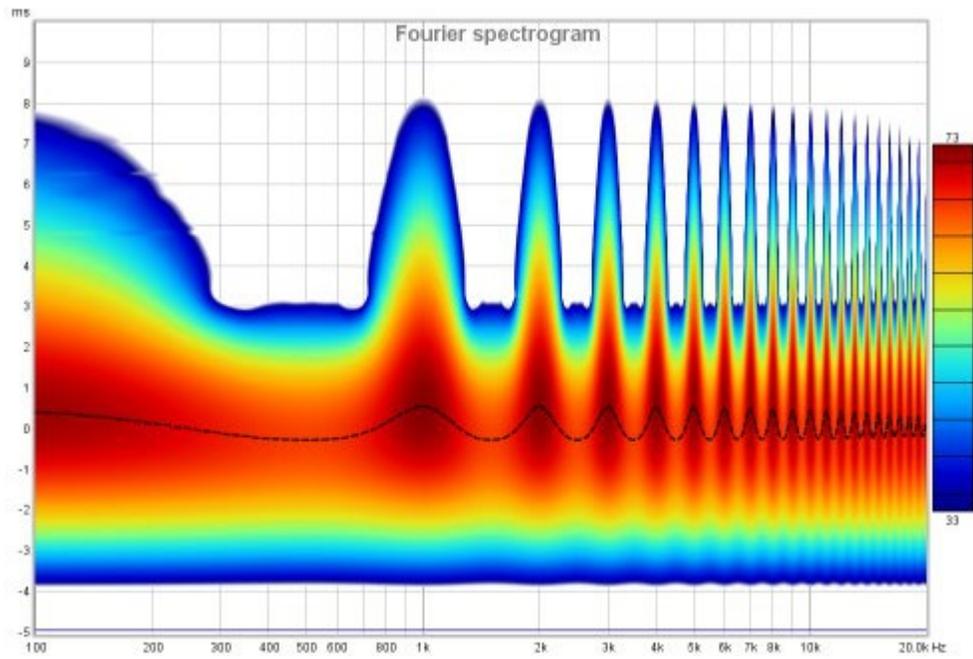
The wavelet plot also shows the frequency response and decay effects, but thanks to its greater time resolution at high frequencies the reflections themselves become visible as horizontal bars.
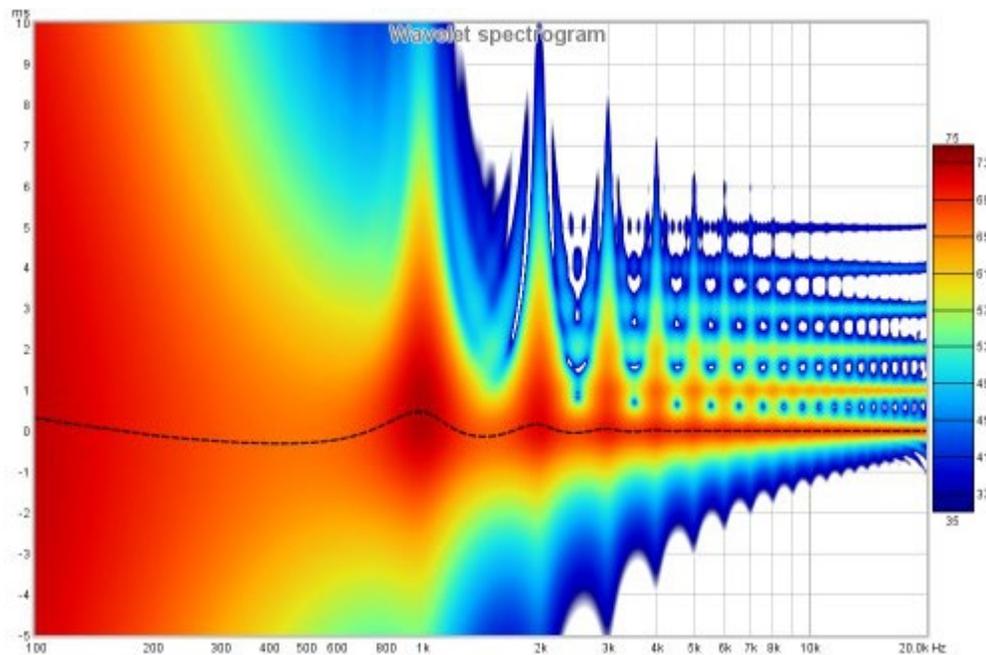
In **Wavelet** mode a **Freq. Resolution** control replaces the **Window** control and allows resolutions between 1 octave and 1/24th octave to be selected.

The **Burst Decay** mode provides a way to more easily distinguish resonances of similar Q but different frequencies. It does this by showing the way a shaped tone burst at each frequency would decay, but along an axis denoted in periods of the frequency rather than time. On a period axis the extent of a decay is the same for resonances of the same Q regardless of the frequency of the resonance. The plot is produced by convolving the windowed impulse response (using whatever current window settings have been applied to the measurement) with a complex Morlet wavelet analytic signal (a Gaussian windowed complex exponential) to extract the decay envelope then resampling that decay on a period-based scale. That is repeated at 48 points per octave across the frequency span of the measurement, with 10 Hz as the lowest

frequency allowed and the burst bandwidth below half the sample rate as the highest. Note, however, that an artefact of the period axis is to skew the tail of the decay slightly towards higher frequencies rather than maintaining the symmetry about the resonance's centre frequency that would be seen in a time-based plot.

In **Burst Decay** mode a **Bandwidth** control allows burst bandwidths of 1/3rd or 1/6th of an octave to be selected. The 1/3rd octave choice favours time resolution, the 1/6th octave choice favours frequency resolution. With the 1/6th octave burst resonances are more easily distinguished. With the 1/3rd octave burst reflections become more evident in the plot, showing up as curved lines. The image below is the same response with reflections examined using the Fourier and Wavelet spectrogram plots above, but using a 1/3rd octave Burst Decay.
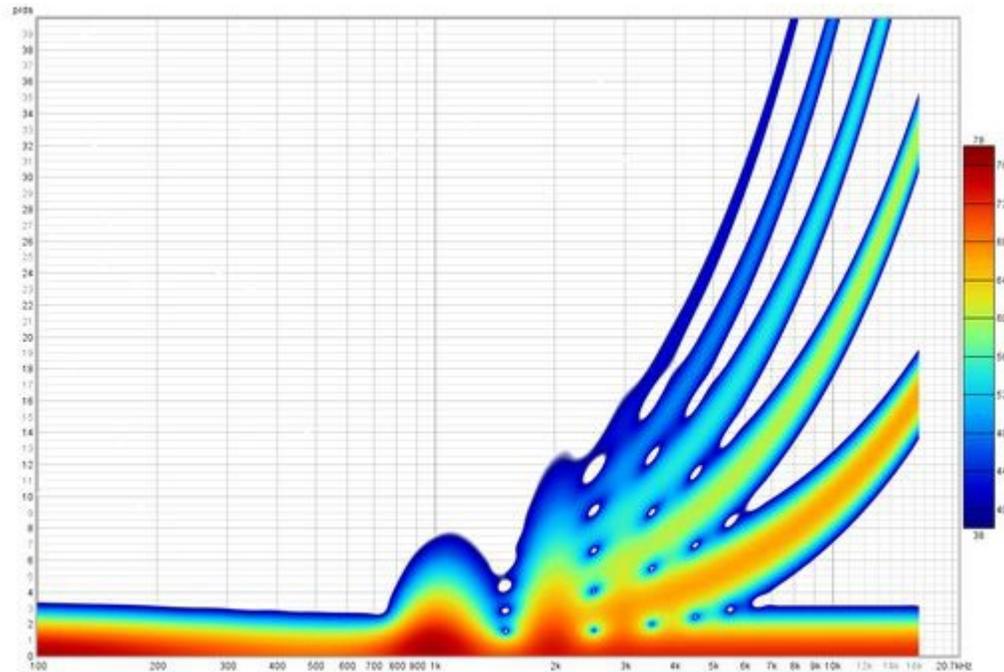
Burst decay plots may have visible artefacts near the high frequency limit, though typically more than 40 dB below the peak level. A **Periods** control sets the number of periods the plot will span.

The **Window type** control selects the window that is used for each slice of a Fourier spectrogram, Hann is well suited to viewing the content of imported audio files, Gaussian provides a more optimal time/frequency trade-off for sweep measurements.

The **Span before peak** and **Span after peak** controls determine how much spectrogram data will be generated around the impulse response peak for a sweep measurement. There are no span controls for imported audio files, the spectrogram is generated for the whole span of the file.

The **Amplitude** control offers a choice between linear and logarithmic scales. The log scales are dB SPL and dBFS, the linear scales are % peak and % FS. Using the linear % peak scale with a Wavelet plot makes it easier to see timing shifts. The dBFS and % FS scales may be useful when viewing imported audio files.

Frequency axis determines whether frequency is along the X or Y axis. Spectrograms of audio data typically have frequency along the Y (vertical) axis, having frequency along the X (horizontal) axis allows easier visual comparison with waterfall plots.

Normalise to peak at each frequency scales (boosts) the plot at each frequency so that it has the same peak value. This can be useful when examining energy decay or the time alignment between drive units as it removes the level differences. Note that using 3D enhancement with normalisation may result in artefacts along the frequency axis.
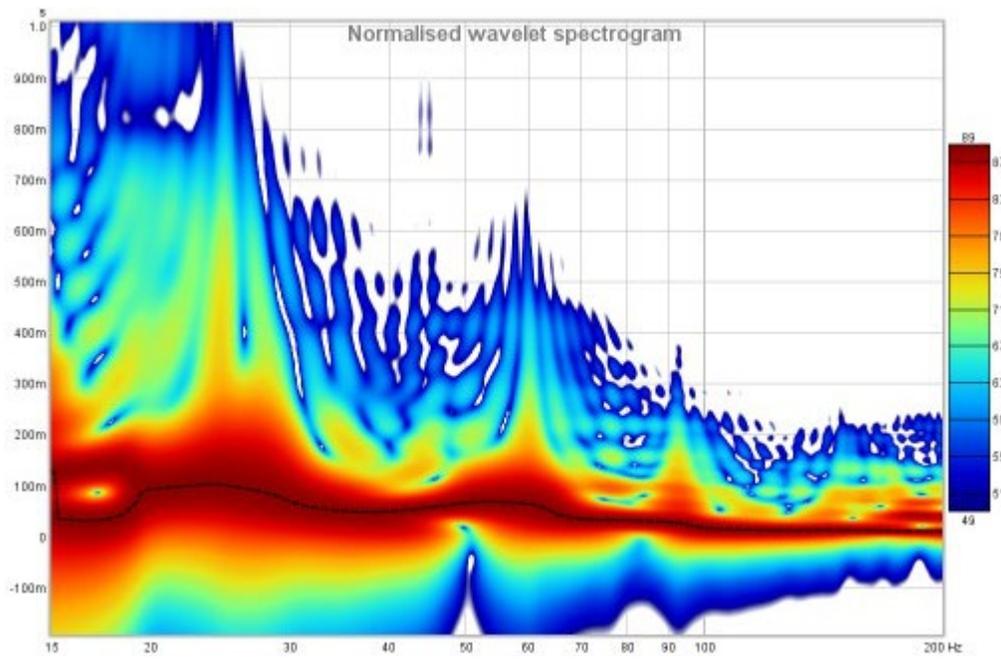
If **Auto generate on settings change** is selected the spectrogram will be regenerated automatically if a setting is changed, otherwise new settings will not be applied until the **Generate** button is pressed.

If **Show slice at cursor** is selected the graph is split with the spectrogram on the top and a plot of the level at the cursor Y axis position underneath. The graph will show either level vs frequency or level vs time depending on whether frequency is on the X axis or the Y axis.
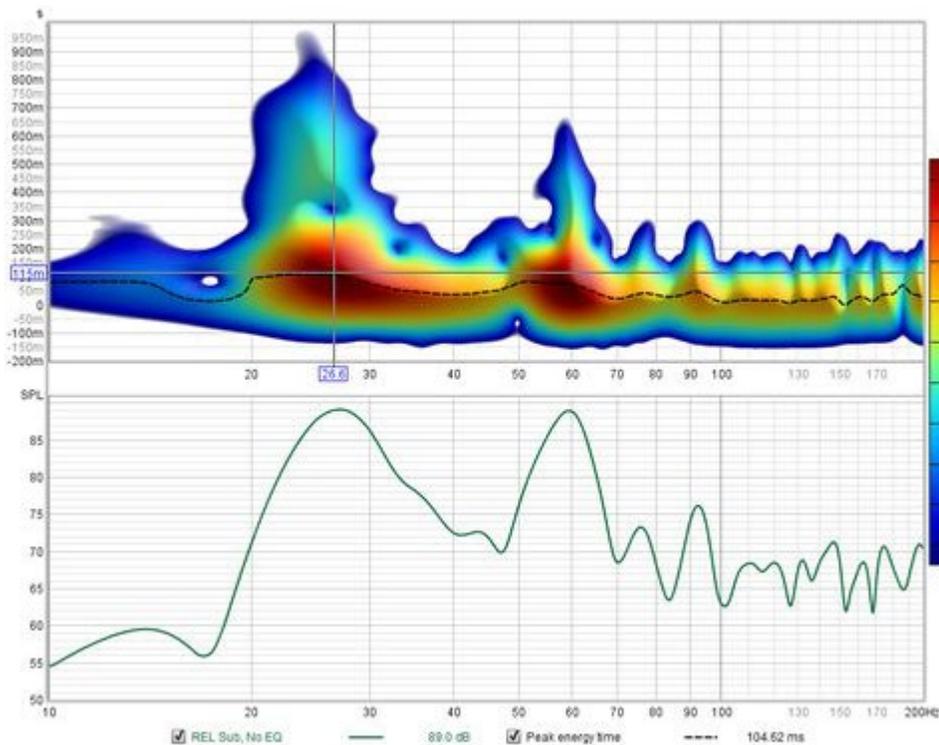
Settings can be copied and pasted between measurements by right clicking on the controls panel.

The **Appearance settings** button opens a dialog with settings that control the appearance of the spectrogram.
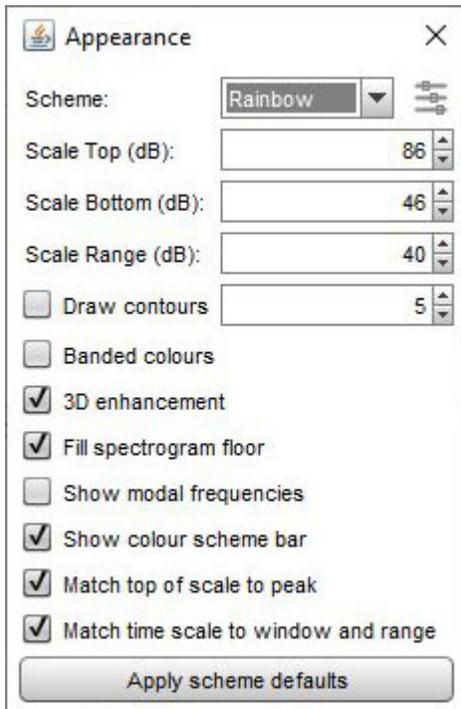
The colour **Scheme** for the plot can be changed, the plots above use the "Heat" scheme, here is a plot using the "Copper" scheme with 3D enhancement active.
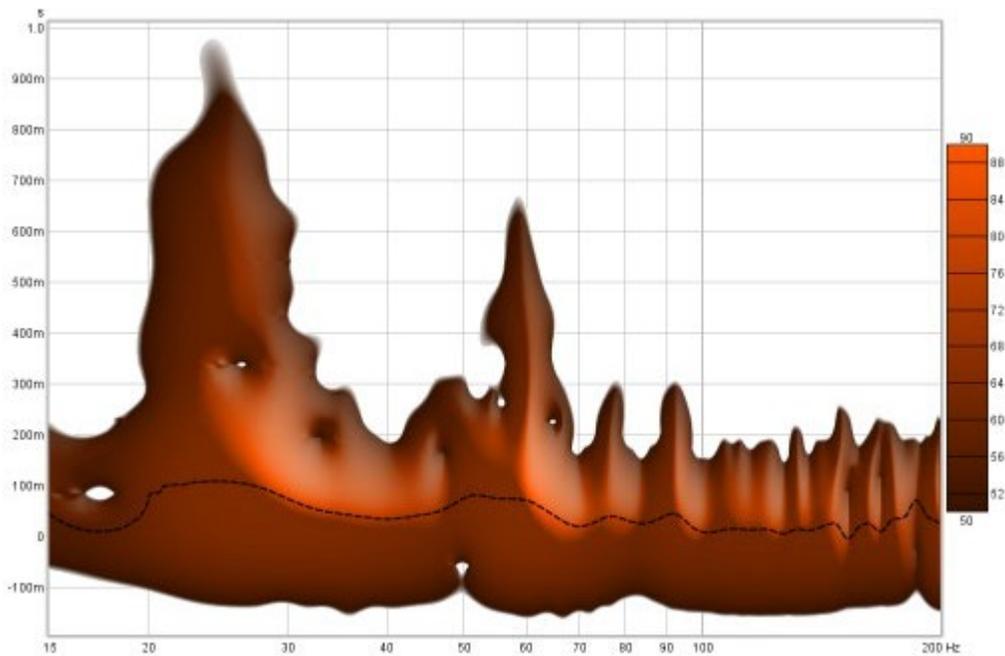
One of the colour schemes is based on cubehelix by Dave Green, see https://www.mrao.cam.ac.uk/~dag/CUBEHELIX/. It is based on a helical path around the diagonal of an RGB colour cube, taking into account the perceived intensity of colours to create a scheme that perceptually has monotonically increasing brightness. The cubehelix scheme can be configured to change its appearance using the settings panel below, which is activated by clicking the icon to the right of the colour scheme selector:

Start hue is the hue in degrees at the base of the plot. **Rotation** is how many degrees the helix travels around the cube diagonal, setting rotation to zero produces a scheme with a single hue. Rotation can be positive or negative. **Hue factor** is a scaling applied to the colours, a factor of 1.0 ensure perceptual uniformity but higher values produce a more colourful scheme. The original scheme covers the whole span from black to white, but the **Min grey** and **Max grey** controls allow starting at a level above black, making the start hue visible, and ending before white, leaving some colour at the top of the scale.

The **Scale Top**, **Scale Bottom** and **Scale Range** controls adjust how the plot colours correspond to the values in the Spectrogram data. Any values higher than the Scale Top are drawn in the colour at the top of the scale, any values lower than the Scale Bottom are drawn in the colour at the bottom. If the Scale Top setting is changed the Scale Bottom will be adjusted to keep the same Scale Range. If the Scale Bottom is changed the Scale range will be adjusted to keep the same Scale Top. If the Scale Range is changed the Scale Bottom will be adjusted, keeping the same Scale Top.

Draw contours adds contour lines at the dB interval set in the adjacent spinner.
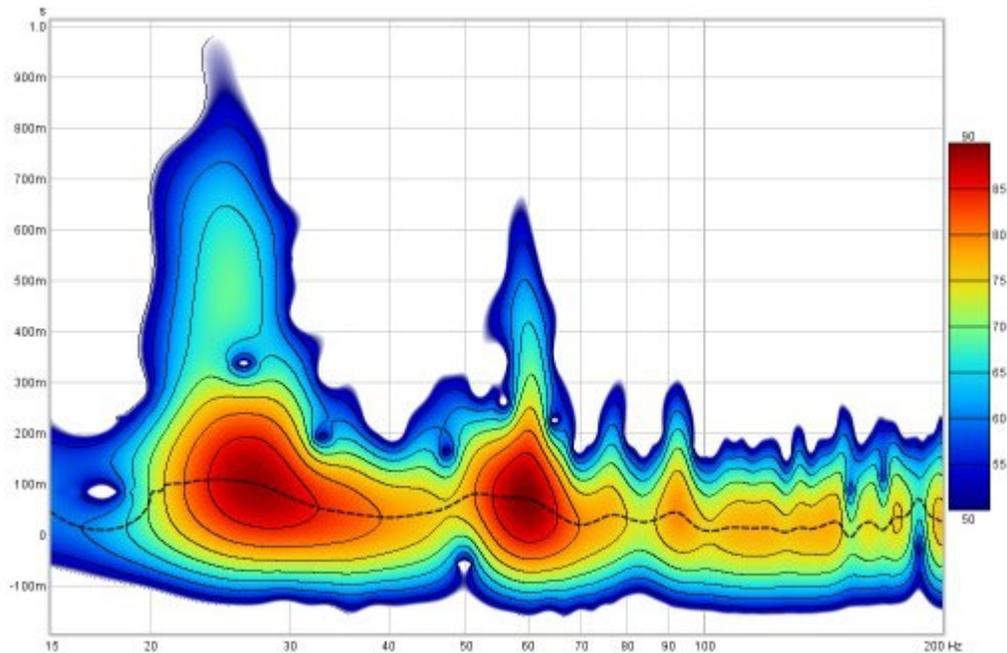
If **Banded colours** is selected the colour scale has discrete steps rather than a continuous blend from one colour to another - there are 11 colours in that case to provide 10 bands across the scale range.
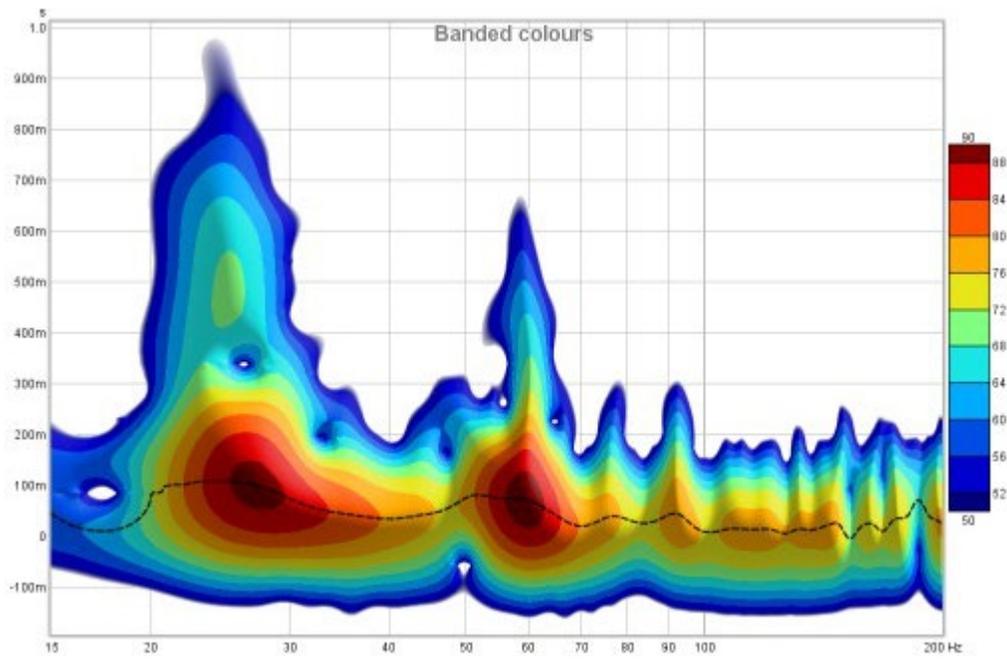
3D enhancement gives the plot a more three-dimensional appearance.
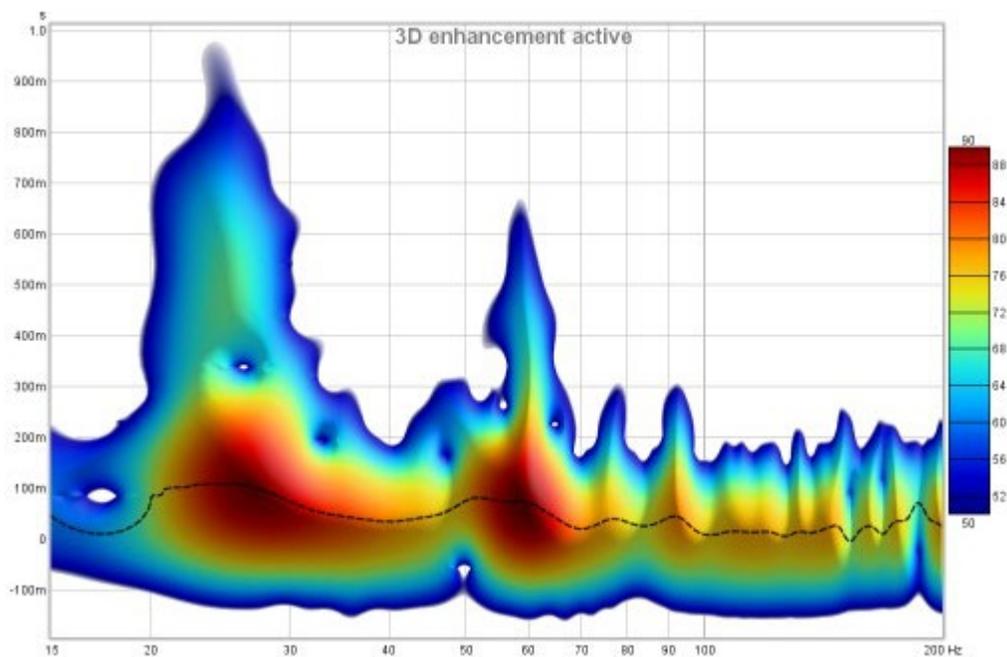
Fill spectrogram floor fills the floor of the plot with the colour at the bottom of the scale range. When the floor is filled the grid is drawn on top of the spectrogram, it can be shown/hidden using the Show/Hide Grid toggle in the Graph menu or using the Ctrl+Shift+G shortcut.
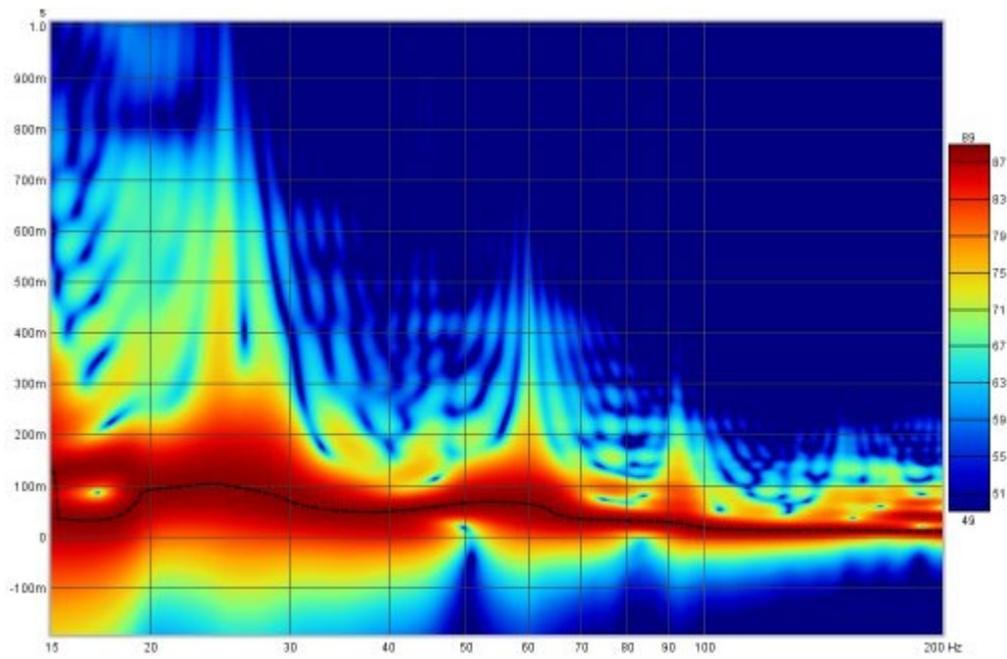
If **Show modal frequencies** is selected the theoretical modal frequencies for the room dimensions entered in the Modal Analysis section of the EQ Window are plotted at the bottom of the graph.

Show colour scheme bar controls whether the bar showing the relationship between colours and SPL is shown to the right of the graph.

Match top of scale to peak adjusts the Scale Top value so that it corresponds to the highest level found in the data.

Match time scale to window and range adjusts the time axis range so that it starts at the Window width before zero (e.g. -300 ms for a 300 ms Window setting) and ends at the Time Range (e.g. 1000 ms for a 1000 ms Time Range) so that the plot shows all the generated data.

The control settings are remembered for the next time REW runs. The **Apply Default Settings** button restores the controls to their default values.

Spectrogram Controls - Stepped Sine Measurements

Stepped sine measurements have a much reduced set of controls, to select the amplitude, frequency axis, colour scheme and the SPL range. The equivalent of the time axis for stepped sine measurements is the test frequency at which the spectrum data was captured, those frequencies are shown along the axis. When a stepped sine measurement is selected the axis is automatically scaled to show all of the test frequencies in the measurement, but it can subsequently be zoomed in or out using the axis zoom buttons. Note that spectrograms can only be generated for stepped sine measurements that had the option to **Capture spectrum data at each frequency** selected.
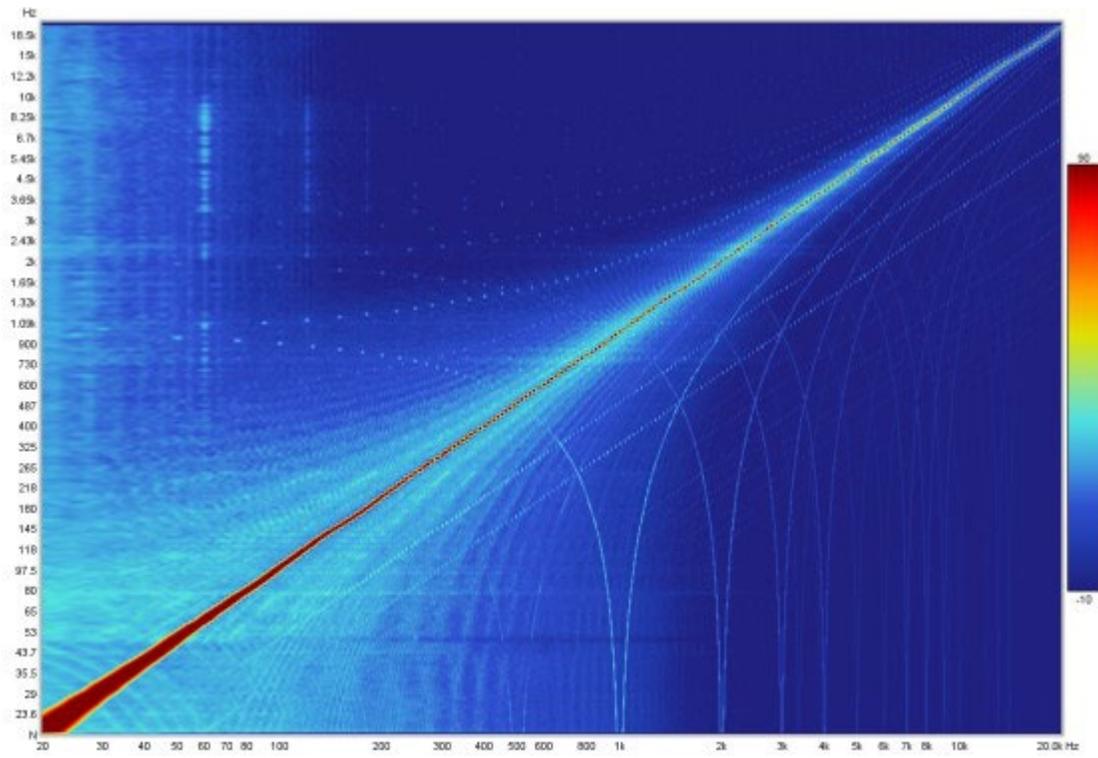
[Help Index](#)

Captured Graph

This graph is available for the LAST generated sweep test signal and the uncorrected system response acquired via the interface. This may be useful for troubleshooting. This is not a live display. Content updates after a sweep is complete. Only the signals for the last measurement are shown. It is not available on measurements loaded from disk. The Y axis is the percentage of digital full scale. If the captured trace reaches ±100%, it is clipping [distorted] and the sweep level, AV processor volume or input volume should be reduced.
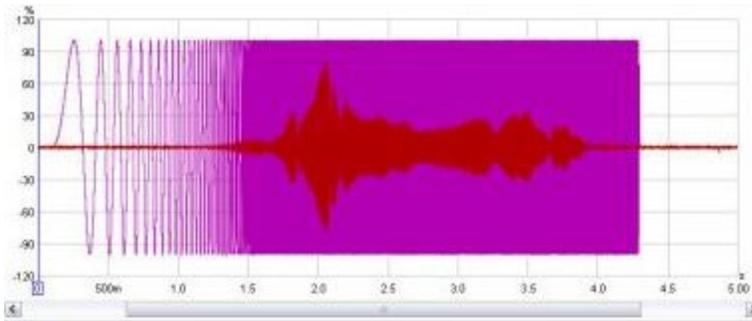
Note:
The Captured signal is offset by latencies in the replay and capture chains.

Captured Controls
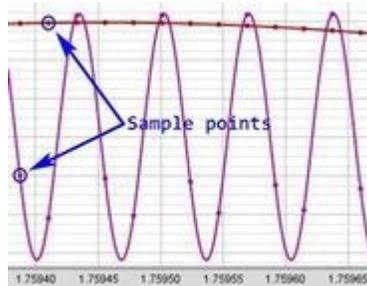
A check box is provided to invert the captured trace for easier comparison with the test signal if the soundcard input is inverting. As a more permanent solution for this select the **Invert** checkbox in the soundcard Input Channel settings.

Show points when zoomed in shows the individual time samples if the horizontal zoom level is high enough to distinguish them.

Help Index

Overlays Window

The overlays window shows plots for all the currently loaded measurements. It is shown by pressing the Overlays button in the toolbar of the main REW window.

Overlays

The overlay plots are selected via the buttons at the top of the graph area.

The various graph types are:

SPL

> All the measurement SPL traces

Predicted SPL

> The predicted SPL for each measurement after applying any EQ filters that have been defined for the measurement in the EQ Window.

Phase

> All the measurement phase traces

Predicted Phase

> The predicted phase for each measurement after applying any EQ filters that have been defined for the measurement in the EQ Window

Distortion

Distortion traces for each measurement, showing the distortion measure selected in the graph controls, which can be THD, any available harmonic up to the 10th or the fundamental

Impulse

All the measurement impulse responses

ETC

All the measurement impulse response envelope traces

Step

All the measurement step responses

GD

All the measurement group delay traces

RT60

All the measurement RT60 traces

Clarity

All the measurement clarity/definition traces

Separate Traces

The basic controls for the overlay graphs are described in the main Graph Panel help, but the Overlays window has one additional button:

The **Separate Traces** button to the right of the graph selector offsets each trace downwards from the preceding trace to make it easier to distinguish individual features when the traces are at similar levels. When the button has been pressed a slider next to it allows the amount of separation to be adjusted.

Graph Controls

The SPL graph has controls to apply smoothing to all the currently selected traces, a control to offset any of the traces and a box to select whether data points should be plotted. The trace offset moves the graph

position, but does not alter the data so the legend values do not change. If the **Add offset to data** button is pressed the current offset value is transferred to the measurement data and the legend readings will update accordingly. If **Show points when zoomed in** is selected the individual points that make up the measured phase responses are shown on the graph when the zoom level is high enough for them to be distinguished (which may only be over part of the plot).

The **Trace options** button brings up a dialog that allows the line type of the graph traces to be changed.
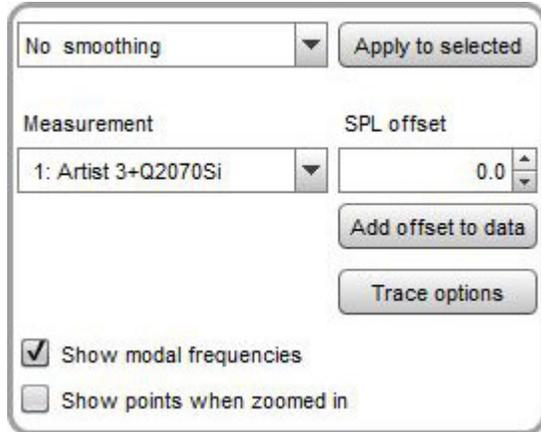
The Predicted SPL, Phase, Predicted Phase and Group Delay overlays also have a smoothing control. The Phase and Predicted Phase overlays have additional controls to wrap or unwrap the currently selected phase traces. The Phase, Impulse, Step and Group Delay overlays have a control to show data points when zoomed in. All have the trace options control.

Legend area

Right clicking in the legend area of an overlay graph brings up a small menu that allows all traces to be selected, all selections to be cleared or the selections to be toggled. There are also menus to select or clear traces for measurements that have been loaded from a particular file.

Hovering the cursor over the name of a measurement in the legend panel will bring up a tool tip showing the measurement notes.

[Help Index](#)

RTA Window
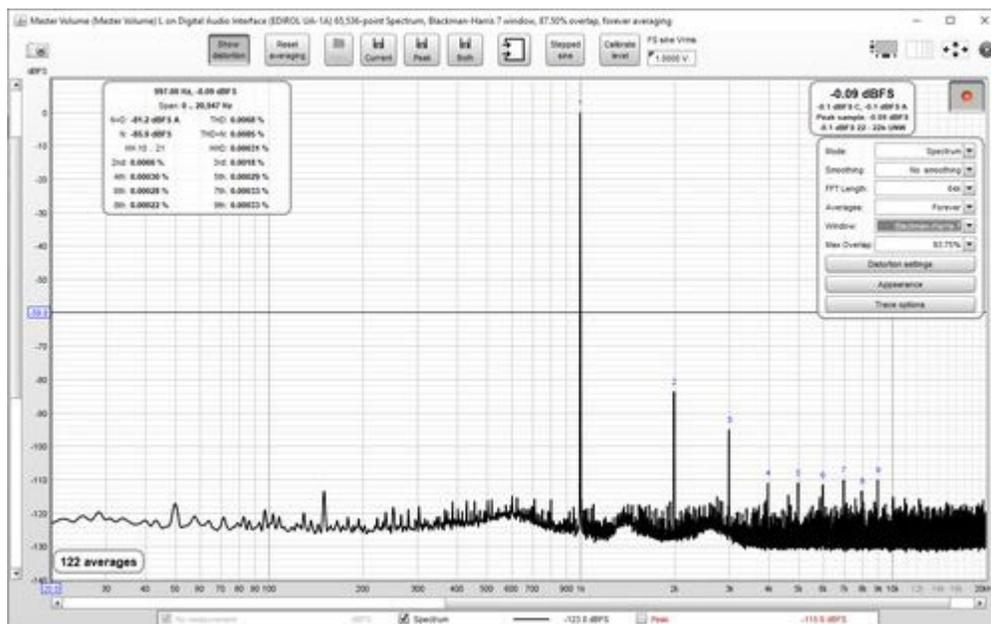
The RTA window allows Real Time Analyser (RTA) or spectrum analyser plots to be generated, updating as the input signal is analysed. It is shown by pressing the RTA button in the toolbar of the main REW window.

The RTA trace is activated by pressing the record button in the top right hand corner of the graph area, after which it will continuously analyse blocks of input samples and display the frequency spectrum of each block. If the RTA settings are such that the update interval is more than 1 second the record button will show a percentage progress figure.

Sometimes the analyser is used without a test signal, for example to look at the frequency content of background noise, but more often it is used together with the REW generator or an external generator or signal source. If the generator is playing a pink noise signal (or even better, pink Periodic Noise) the RTA display will show the frequency response of the room, updated live so that the effects of changing EQ settings can be immediately seen.

Playing a sine wave test tone on the generator allows the levels of the tone and its harmonics to be observed on the analyser and distortion percentages to be calculated, whilst using the dual tone generator allows intermodulation distortion measurements.

The RTA plot shows the currently selected measurement as a reference and the live RTA or spectrum. A Peak trace is also available, which is reset by the **Reset averaging** button. If Inverse C compensation is being applied the icon is shown after the trace value. If Mic/Meter calibration file or soundcard calibration file have been loaded they are applied to the results.
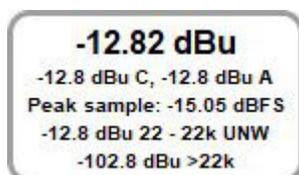
The current Input RMS value is shown to the left of the record button, in dB SPL, dBFS, dBu, dBV, dBW, volts or watts according to the setting of the Y axis. This figure excludes any DC content in the signal. A and C weighted values are shown below the unweighted rms figure. **All three figures are calculated over the range specified by the distortion LP and HP settings, if either or both are enabled.** The peak sample value in the last RTA block is shown in dBFS. Below that are the in-band (22.4 Hz to 22.4 kHz) and, for sample rates above 44.1 kHz, out of band (> 22.4 kHz) unweighted RMS levels. If clipping is detected in the input the RMS value turns red. If the RTA is in one octave mode Noise Criterion and Noise Rating figures will be shown below the RMS values.

A dBc Y axis option is offered which places the peak level of the input at 0 dBc, or places the peak of the fundamental at 0 dBc when the distortion panel is active. There is also a volts per √Hz Y axis option to display the amplitude spectral density, this is only meaningful when using the spectrum view.

Input selection

The current RTA input can be changed by clicking the Input selection button:

That brings up an input selection dialog, which also has a button to show the currently active mic cal files.
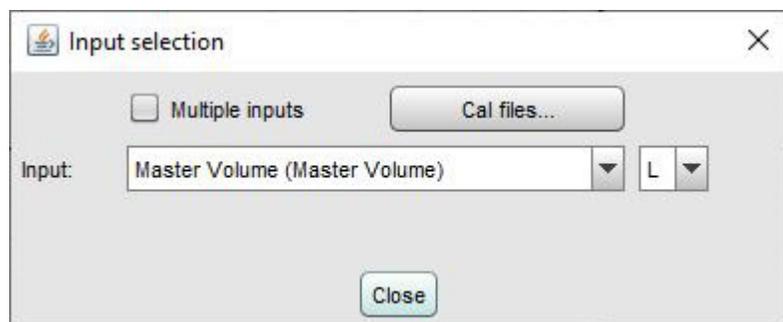
When using multi-input capture (Pro upgrade) there is an option to select multiple inputs and choose the range of inputs that will be averaged by the RTA to produce the graph traces. Up to 16 individual input traces can be displayed along with the average and peak, according to the setting in the View preferences.
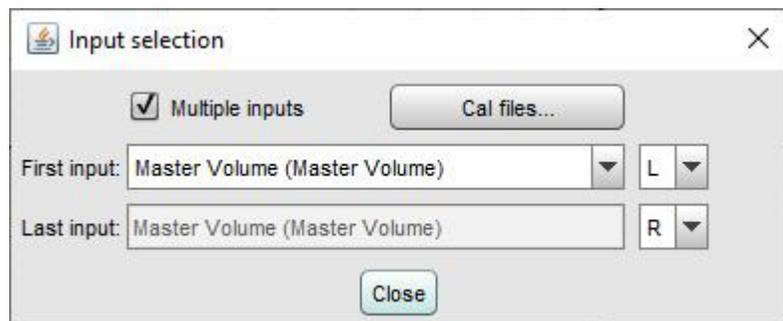
Calibrating the input level

The input level is calibrated by pressing the **Calibrate level** button above the graph while the RTA is running. A signal with a known rms voltage level should be applied to the input and that rms level entered in the calibration dialog.
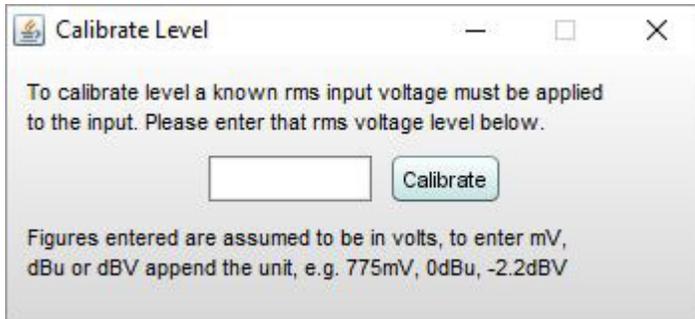
After the value is entered a confirmation message will be shown, stating the maximum peak and rms input voltages the input can accept before clipping at digital full scale. If any volume control setting along the input path is altered the calibration will need be done again. If the input sensitivity is known (including the effect of any volume control settings) it can be typed directly into the text field below the **FS sine Vrms** label. Clicking in the triangle in the upper left corner brings up a list of preset values that may be entered, the labels and values can be changed as required.
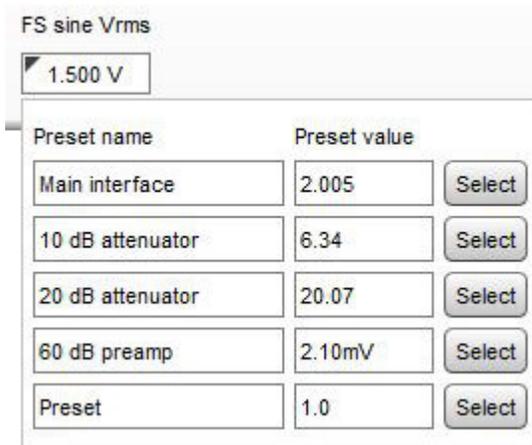
Processing WAV files

The RTA can be used to analyse WAV files (singly or in groups) by dragging and dropping them on the RTA window or by clicking the Open WAV button (shortcut Alt+O) to select a WAV file to process. The dialog below appears to determine how each file will be processed, at the end of processing the current, peak or both results will be saved as measurements according to the selection made.
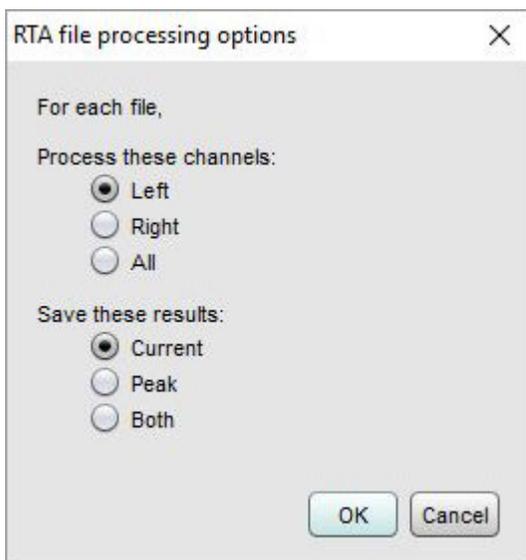
Spectrum/RTA controls

The controls for the plot are shown below.
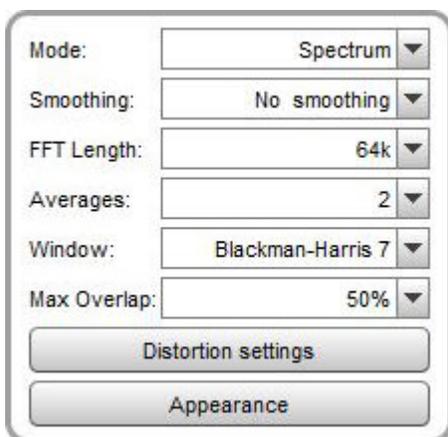
The **Mode** can be set to Spectrum for a spectrum analyser plot or to various RTA resolutions from 1 octave to 1/48 octave. The difference between spectrum and RTA modes is how the information is presented. In spectrum mode the frequency content of the signal is split up into bins that are all the same width in Hz. For example, with a 64k FFT length and 48 kHz sample rate the bins are 0.732 Hz wide. The plot shows the energy in each of those bins. In RTA mode the bin widths are an octave fraction, so their width in Hz varies with the frequency. For example, a 1 octave RTA plot has bins that are 70.7 Hz wide at 100 Hz (from 70.7 Hz to 141.4 Hz) and 707 Hz wide at 1 kHz (from 707 Hz to 1.414 kHz). The plot shows the combined energy at each frequency within each bin. This is closer to how our ears perceive sound. The different presentations mean signals with a spread of frequency content will look different on the plot. The best known examples are white noise and pink noise. White noise has the same energy at each frequency. On a spectrum plot, which shows the energy at each frequency, the white noise plots as a horizontal line. On an RTA plot it appears as a line that rises with increasing frequency, as each RTA bin gets wider it covers more frequencies and so has more energy. The bin widths double with each doubling of frequency so the energy also doubles, which adds 3 dB on the logarithmic plots we use to show level. White noise sounds quite 'hissy', we perceive it as having more energy at higher frequencies.

Pink noise has energy that falls 3 dB with each doubling of frequency. On a spectrum plot it is a line that falls at that 3 dB per octave rate, on an RTA plot it is a horizontal line as the energy in the signal is falling at the same rate as the bins are widening. We perceive pink noise as having a uniform distribution of energy with

frequency.

Single tones are a special case, they will appear at the same level on either style of plot as their energy is all at one frequency, so on a spectrum plot they show as a vertical line, on an RTA plot they show (typically) as a bar of the width of the bin at their frequency, but the height of the bar is the same as the height of the line on the spectrum as all the energy is at that one frequency.

In Spectrum mode smoothing can be applied to the trace according to the setting of the **Smoothing** box. Smoothing is not applicable for RTA modes.

FFT Length

The **FFT Length** determines the basic frequency resolution of the analyser, which is sample rate divided by FFT length. The shortest FFT is 8,192 (often abbreviated as 8k) which is also the length of the blocks of input data that are fed to the analyser. An 8k FFT has a frequency resolution of approximately 6Hz for data sampled at 48kHz. As the FFT length is increased the analyser starts to overlap its FFTs, calculating a new FFT for every block of input data. The degree of overlap is 50% for 16k, 75% for 32k, 87.5% for 64k and 93.75% for 128k. The overlap ensures that spectral details are not missed when a Window is applied to the data. The maximum overlap allowed can be limited using the **Max Overlap** control below to reduce processor loading at higher FFT lengths.

Averaging

The plot can be set to show the live input as it is analysed or to show the result of averaging measurements, according to the selection in the **Averaging** control. Selecting a number for averages results in that many measurements being averaged to produce the result, with the oldest measurement being removed from the average as each new measurement is added. There are several **Exponential** averaging modes, which give greater weighting to more recent inputs. The figure shown in the selection box is the proportion of the old value which is retained when a new measurement is added, the higher the figure the more heavily averaged the display becomes. There is also a **Forever** averaging mode which averages all measurements with equal weight since the last averaging reset. After starting the RTA or changing the FFT length averaging does not begin until a full FFT length of data has been received, plus the lengths of the input and output buffers.

The **Reset averaging** button above the graph restarts the averaging process (keyboard shortcut Alt+R). Averaging is needed when measuring with pink noise or when there is noise in the signal being measured. Note that if measuring a response using pink noise the best results are obtained using REW's periodic noise signals, which can be exported as wave files from the signal generator to produce a test disc for the system to be measured if direct connection to the PC running REW is not possible.

The **Save current** button converts the current display into a measurement in the measurements pane (keyboard shortcut Alt+S). It is converted in the current mode of the analyser, so if the analyser is in Spectrum mode the measurement shows the spectrum, if it is in RTA mode it shows the RTA result. The saved measurements can be used as references for subsequent spectrum/RTA measurements. If distortion data is available it is copied to the comments area of the saved measurement. Peak data can similarly be saved using the **Save peak** button, or both saved at once using the **Save both** button.

Window

The FFT resolution is also affected by the **Window** setting. Rectangular windows give the best frequency resolution but are only suitable when the signal being analysed is periodic within the FFT length or if a periodic noise signal is being measured. The Rectangular window should always be used with the REW periodic noise signals. Most other signals, e.g. sine waves from the REW generator or test tones on a CD or random noise, typically would not be periodic in the FFT length. Using a rectangular window when analysing such a tone would generate spectral leakage, making it difficult to resolve the frequency details - the plot below shows an example of a 1kHz tone from an external generator with a Rectangular window.
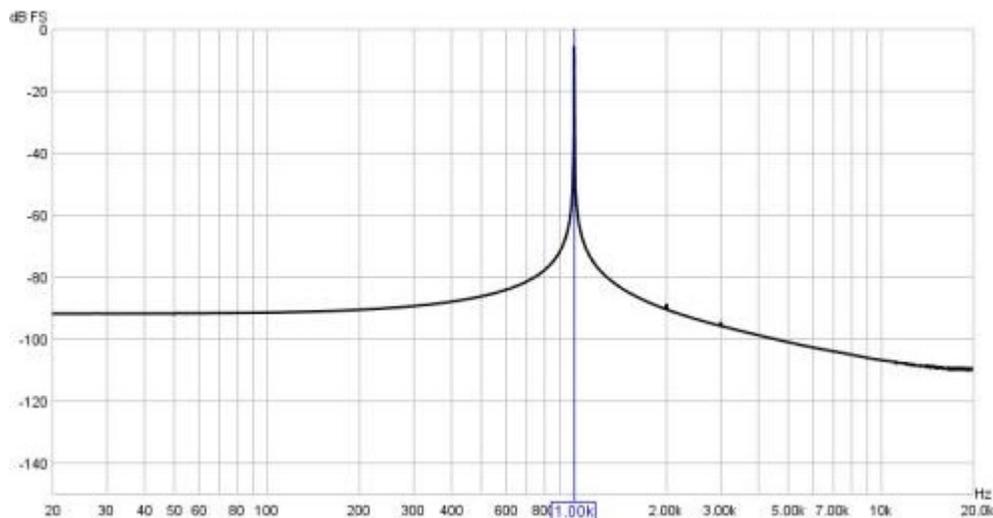
Here is the same tone analysed with a Hann window.
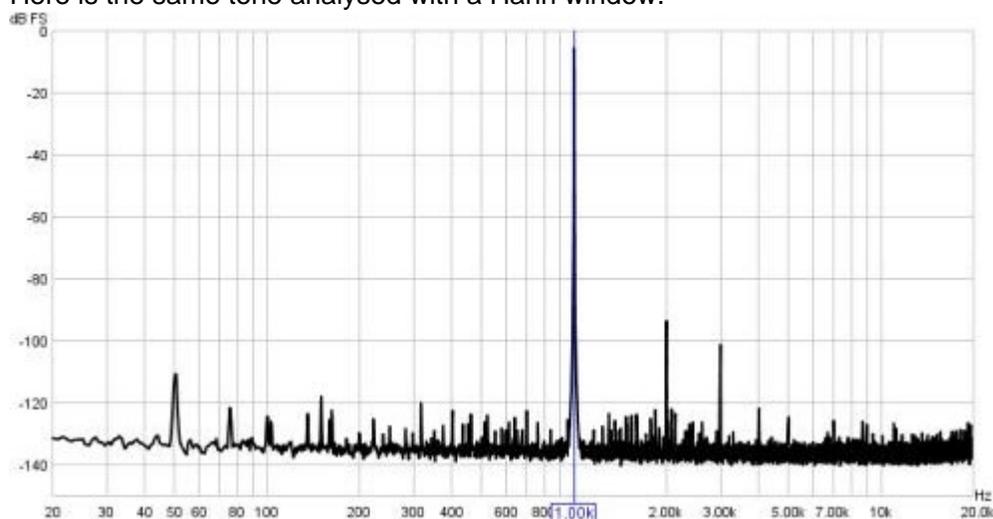
The window allows the harmonics of the tone to be resolved. However, the tradeoff is that windows cause some spreading of the signal they are analysing, which reduces the frequency resolution. To use a rectangular window with the REW signal generator use the generator's Lock frequency to FFT option.

The Hann window is well suited to most measurements, offering a good tradeoff between resolution and shoulder height. If very high dynamic range needs to be resolved (very small signals close to very large signals) use the 4-term or 7-term Blackman-Harris windows. If the spectral peak amplitudes must be accurately measured use the Flat Top window, this will provide amplitude accuracy of 0.01 dB regardless of where the tone being measured falls relative to the bins of the FFT. The other windows only show the spectral amplitude accurately if the tone is exactly on the centre of an FFT bin, if the tone falls between two bins the amplitude is lower, with the maximum error occurring exactly between two bins. This maximum error is 3.92dB for the Rectangular window, 1.42dB for Hann, 0.83dB for the 4-term Blackman-Harris and 0.4dB for the 7-term Blackman-Harris.

Max Overlap

The spectrum/RTA plot can be updated for every block of audio data that is captured from the input, overlapping sequences of the chosen FFT length. This can present a significant processor load for large FFT lengths. The processor loading can be reduced by limiting the overlap allowed using this control.

Distortion settings

Distortion High Pass and Low Pass

The Distortion High Pass and Low Pass are used to set the lowest and highest frequencies that will contribute to the calculation of THD and THD+N. They are only applied when the **Enable high pass** and **Enable low pass** boxes are selected and the **Distortion** button is pressed. Either can be enabled individually. When they are active the region of the plot which is excluded from the calculations will be greyed out and the THD and THD+N figures will show the range over which they have been calculated.

Manual fundamental

A value for the fundamental level may be entered for use in harmonic distortion calculations when the fundamental is attenuated by a notch filter. The level is in dBFS and should take account of any gain introduced after the notch. That ensures REW uses the correct fundamental level when calculating distortion. To compensate for the effect of the notch filter on harmonic levels its response can be loaded as a calibration file. For example, make a sweep measurement of a loopback connection firstly with the notch filter in place, then another without it, then use the Trace arithmetic feature of the All SPL graph to generate (notch response)/(no notch response) and export that result as a text file to be loaded as a mic cal file. Alternatively just measure the notch response, offset the measurement so that the dB values correctly reflect the notch filter loss at the harmonics (the fundamental is not critical since Manual fundamental deals with that) and export that offset notch response as text and load it as the mic cal file. The harmonic levels will then be corrected to allow for the filter's attenuation. The fundamental levels shown by the graph trace will typically be in error due to shifts in the centre frequency of the notch, but they will not be used if **Manual fundamental** is selected.

Units for distortion display

Distortion ratios may be shown as either percentages or in dB according to the selection made.

Show phase of harmonics

If this option is selected the distortion data panel includes the phase of each harmonic relative to the fundamental.

Coherent averaging

This option is only applicable to harmonic distortion measurements. If it is selected the FFT data is phase

aligned according to the phase of the fundamental before averaging, this can lower the noise floor substantially compared to magnitude averaging without needing very long FFTs - in fact shorter FFTs (e.g. 64k) will allow faster averaging and more quickly lower the noise floor. The noise level drops by about 3 dB for each doubling of the number of averages. Mains frequency components will be suppressed along with any other noise not harmonically related to the fundamental, so this option is **only** suitable for examining harmonic levels. The various noise-related values in the distortion panel are not relevant when using coherent averaging due to the way the averaging suppresses noise and are shown in orange. Note that if the harmonic levels are varying coherent averaging will converge to their average level whilst magnitude averaging will converge to their rms level.
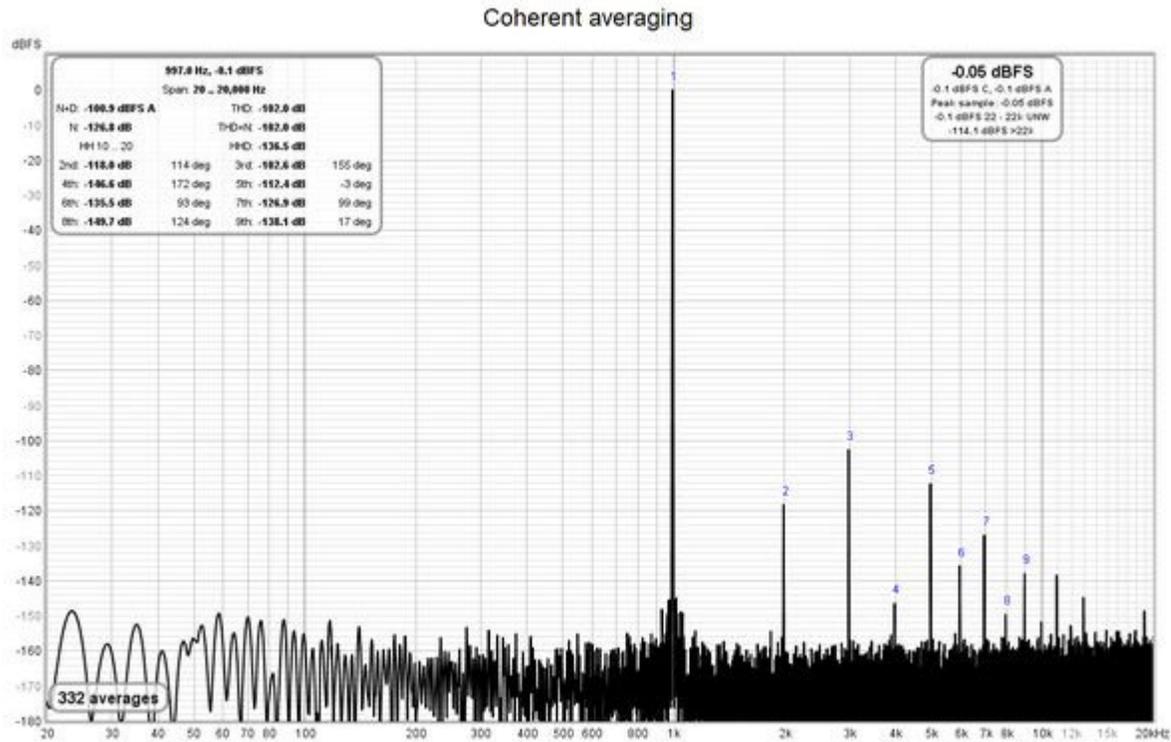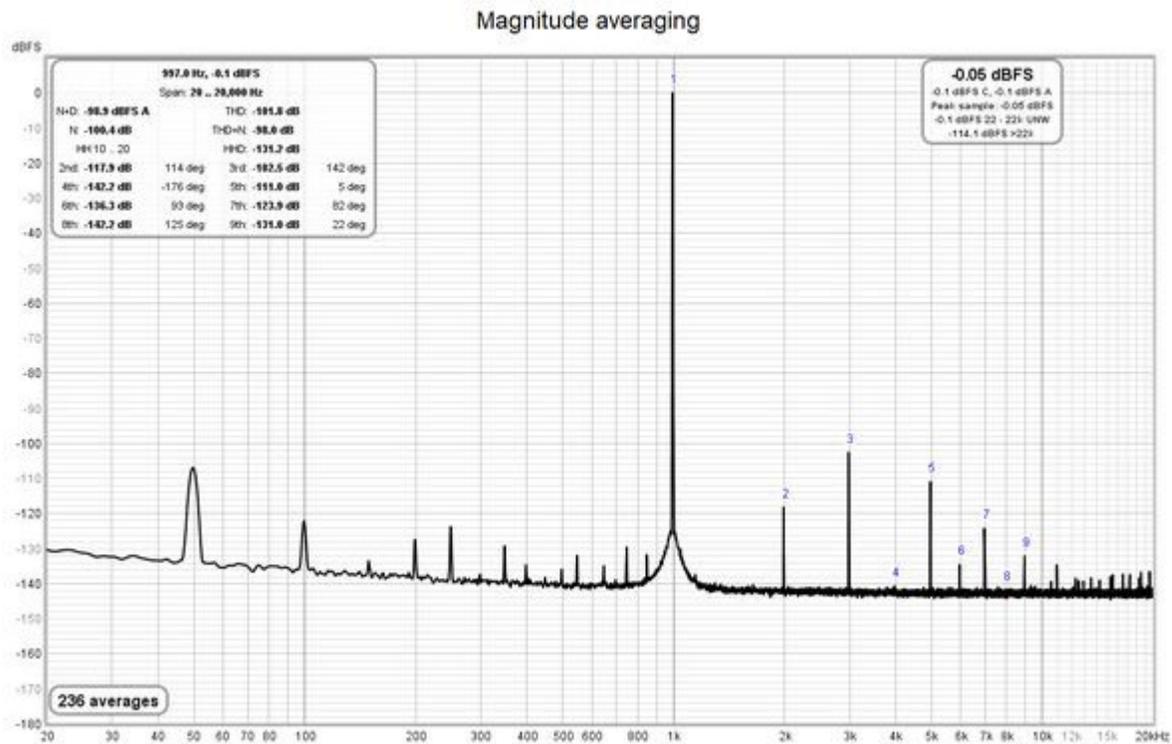
Coherent averaging

Appearance controls

Update Interval

The spectrum/RTA plot is updated by default for every block of audio data that is captured from the input. This can cause a significant processor load, particularly if the RTA window is very large or for large FFT lengths. The processor loading can be reduced by updating the plot less often, which is set by the Update

Interval control. An update interval of 1 redraws the trace for every block, an interval of 4 (for example) only updates the trace on every 4th block.

Peak Hold and Peak Decay

The Peak Hold and Peak Decay controls set how long, in seconds, a peak value is held and how quickly, in dB per second, the peak values decay. If Peak Hold is set to 0 the peak values are not held at all. If Peak Decay is set to 0 the peak trace does not decay.

Show noise curves

If the **Show noise curves** option is selected the chosen noise curves will be drawn on the graph when the RTA is in 1 octave mode and the Y axis is showing dB. The options are Preferred Noise Criterion (PNC), Balanced Noise Criterion (NCB), Noise Criterion (NC) and Noise Rating (NR). Note that if **Adjust RTA levels** is selected the curves will be offset by the same amount as the RTA trace, but this will not alter the noise criteria results.
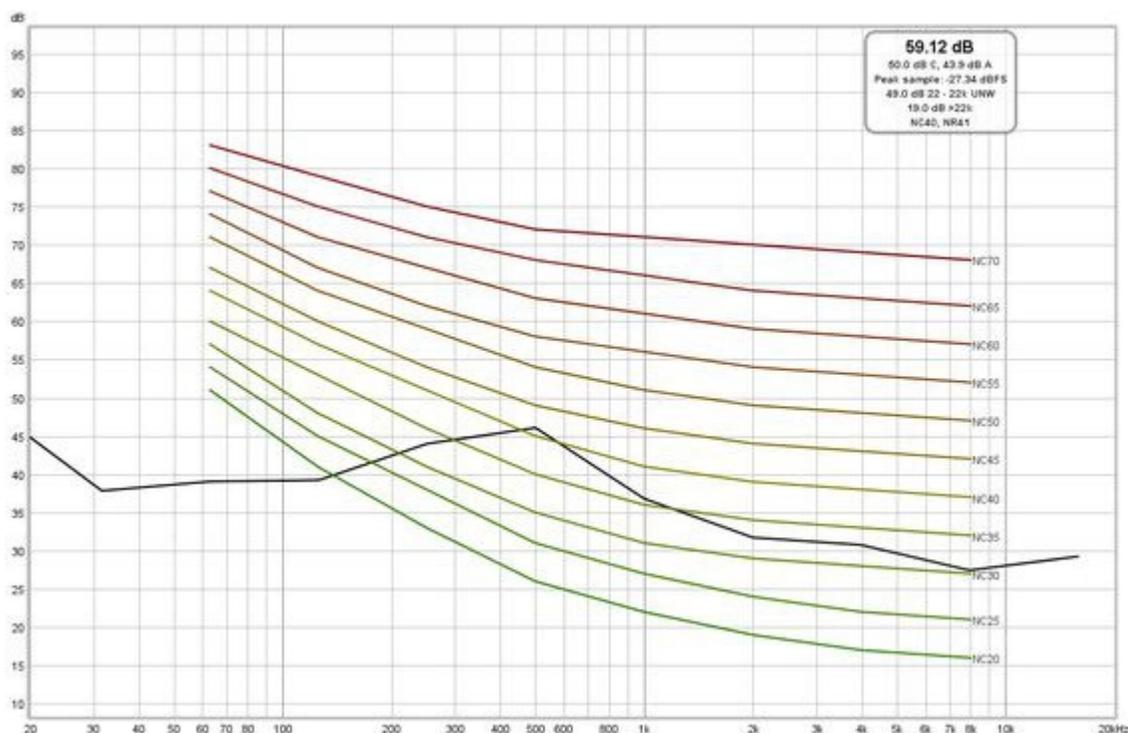
Ref resistance for dBW

The watt and dBW values are calculated from the voltages assuming a reference load resistance, this control sets the value of that resistance.

CEA2010 max SPL limit

The thresholds for distortion plus noise when using the CEA2010 burst test signal can be selected as those used for the CEA-2010 Standard Method of Measurement for Powered Subwoofers or the CTA-2034-A Standard Method of Measurement for In-Home Loudspeakers.

Lines or bars

In Spectrum or RTA modes the plot can either draw lines between the centres of the FFT bins or draw horizontal bars whose width matches the FFT bin or RTA octave fraction width, this is controlled by the **Use**

bars on spectrum and **Use bars on RTA** check boxes.

Adjust RTA Levels

The RTA plot shows the energy within each octave fraction bandwidth. As the RTA resolution increases, from 1 octave through to 1/48 octave, the octave fraction bandwidths decrease and, for broadband test signals such as pink noise, the energy in each octave fraction decreases correspondingly. Whilst the RTA is correctly showing the actual level within each octave fraction, this variation of trace level with RTA resolution can be awkward when using the RTA with a pink PN noise signal to adjust speaker positions or equaliser settings. The **Adjust RTA Levels** option offsets the levels shown on the RTA plot to compensate for both the bandwidth variation as resolution is changed and the difference between a sweep measurement at a given sweep level and a pink PN RTA measurement at the same level, allowing direct comparison between RTA and sweep plots. Whilst the levels shown are not the true SPL in each octave fraction, they are more convenient to work with. **N.B.** This option should only be used with broadband test signals, such as pink noise or pink PN.

Use 64-bit FFT

If this option is selected the RTA uses a 64-bit FFT to process the incoming data instead of 32-bit. This is useful when analysing purely digital 24-bit data paths to view behaviour below -160 dBFS. It has no visible effect when analysing signals that have an analog connection at any point along the data path or when dealing with 16-bit data, as in those cases noise and quantisation effects far exceed any numerical limitations of 32-bit processing. Here are some examples showing the difference the 64-bit FFT makes when analysing undithered and dithered 24-bit data over an S/PDIF loopback connection from REW's signal generator producing a 1 kHz sine wave at -20 dBFS. Note that the 2nd harmonic spike at -173 dBFS in the dithered data appears to be an artefact of data handling within the S/PDIF loopback connection (via Windows 10). The vertical divisions are at 20 dB intervals, the bottom of the plot is at -220 dBFS.

Undithered 24-bit S/PDIF data, 64-bit FFT

Dithered 24-bit S/PDIF data, 32-bit FFT

Dithered 24-bit S/PDIF data, 64-bit FFT
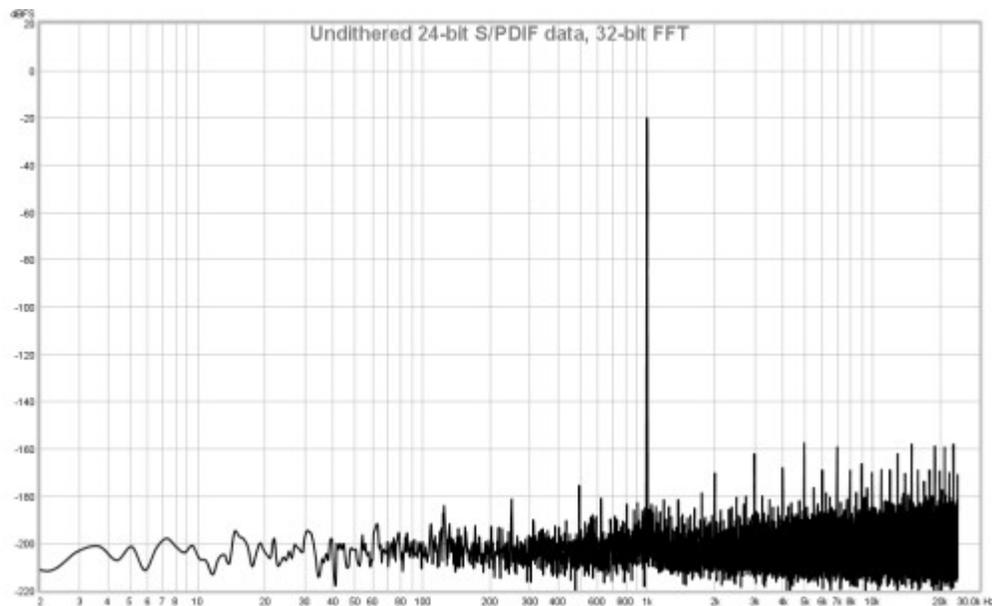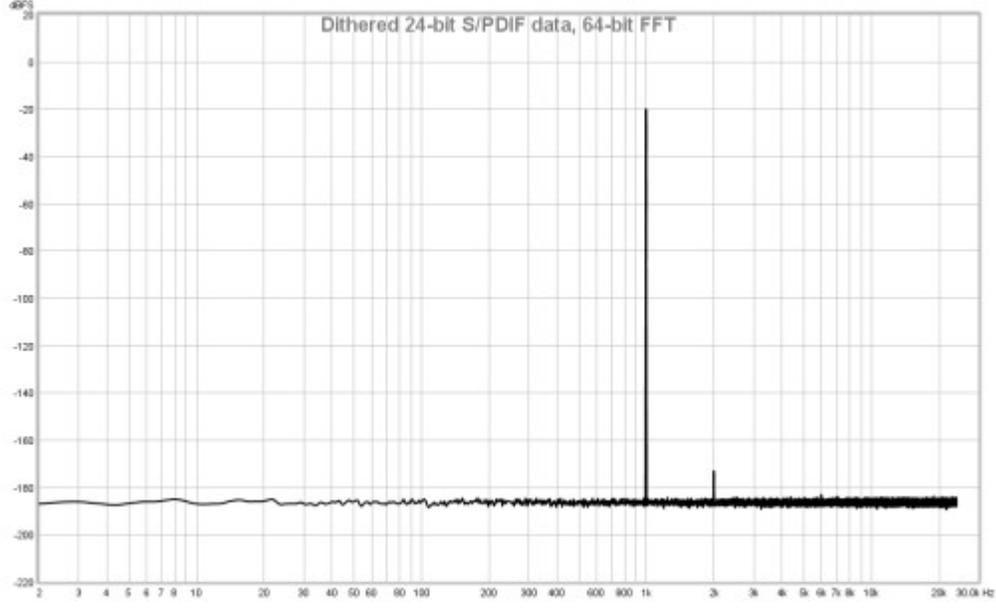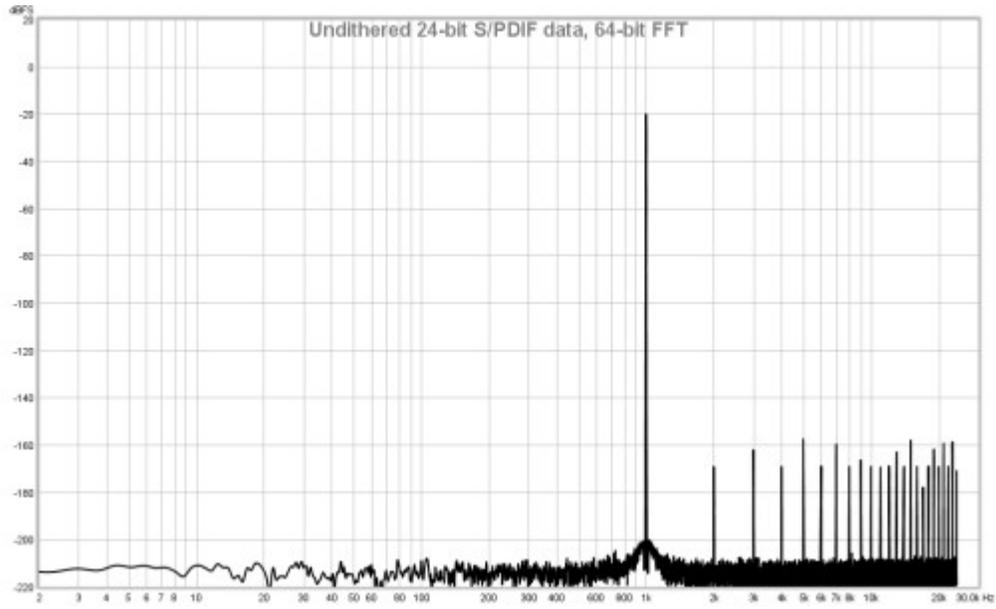
Show tone burst peak SPL

If this option is selected, the generator is playing a tone burst and the graph is in spectrum mode the peak SPL in the input data will be shown on the plot. This is similar to the CEA-2010 peak SPL but without 1/3 octave filtering around the fundamental frequency.

Distortion measurements

When the **Distortion Panel** button (keyboard shortcut Alt+D) is selected the analyser calculates and displays harmonic or intermodulation distortion figures for the input, including THD, N+D (A-weighted noise plus distortion), N (noise and non-harmonic distortion), THD+N, HHD (higher harmonic distortion for harmonics from the 10th up to at most the 50th) and the relative levels of the 2nd to 9th harmonics. N and N+D are displayed in the current Y axis units, harmonic distortion can be displayed as percentages or dB relative to the fundamental depending on the Distortion settings.

Harmonic distortion

Harmonic distortion results are only valid when the system being monitored is driven by a tone or tone burst at a single frequency. If the REW signal generator is playing a sine signal or tone burst the generator frequency is used as the fundamental frequency of the input, otherwise the highest peak is used to determine the fundamental. The fundamental and its level are displayed.

When the stimulus is a tone burst the tone burst envelope has a strong spectral spreading effect, which (depending on the envelope shape) means only relatively high levels of distortion can be measured. The distortion calculations use the level from the FFT bin closest to the fundamental and harmonic frequencies. Noise-related results are not calculated as they would not be meaningful.

When calculating the power for the fundamental and harmonics for a continuous tone the energy in the FFT bins within the relevant span of the nominal frequencies appropriate for the RTA window selection is summed and then corrected according to the window's equivalent noise bandwidth. To obtain accurate results the window should have low side lobes. Good choices in order of reducing side lobe level are Dolph-Chebyshev 150 (side lobes 150 dB down), Blackman-Harris 7 (side lobes 180 dB down), Dolph-Chebyshev 200 (side lobes 200 dB down) and Cosine sum 9-235 (side lobes 235 dB down). Hann, Blackman-Harris 4 and Flat-Top are not recommended. If using the REW signal generator the option to lock frequency to the RTA FFT allows a rectangular window to be used.

Dither should be enabled on the generator at the bit width the system is using, check the bottom left corner of the REW main window for the bit width in use. On Windows using Java drivers only 16 bit is supported, for 24-bit use ASIO drivers. On macOS 24-bit is supported, make sure the devices are configured to operate at 24-bit in Audio Midi setup.

The THD figure is based on harmonics up to at most the 50th or the number of harmonics whose levels are displayed and is calculated from the sum of those harmonic powers relative to the power of the fundamental. A separate HHD (higher harmonic distortion) figure is calculated from the sum of harmonic powers for the 10th up to at most the 50th harmonic. The THD figure includes the HHD contribution. Individual harmonic figures are also calculated from their power relative to the power of the fundamental.

The THD+N figure is calculated from the ratio of the input power minus the fundamental power to the total input power (note that it is possible for THD+N to be lower than THD using these definitions). Note that the reciprocal of THD+N is SINAD. The N figure is calculated from THD+N minus THD, either as a ratio to the total input power if the Y axis units are dBc or as an absolute figure for other Y axis settings.

The upper limit for data used in distortion calculations is 95% of the Nyquist frequency (i.e. 95% of half the sample rate) or the Distortion Low Pass, if enabled. The lower limit is the first FFT bin (DC is excluded) or the distortion High Pass, if enabled.

The example below shows data for a 1 kHz sine input. The positions of the harmonics are shown on the spectrum or RTA plot. The Distortion High Pass and Distortion Low Pass have been set to 20 Hz and 20 kHz respectively, hence results are based on data from the span 20 Hz to 20 kHz.

An A-weighted noise plus distortion (**N+D**) figure is also shown alongside the THD data, in the current Y axis units. Using this figure together with the maximum input level (at distortion of better than -40 dB) allows a dynamic range figure to be generated. For a meaningful N+D result per AES17-2015 the system should be driven with a 997 Hz sine wave at 60 dB below the maximum input level.

Intermodulation distortion

Intermodulation distortion results are only valid when the system being monitored is driven using REW's [Dual Tone](#) test signal. The signal can be generated live from the generator, or saved to a file and played back on the system being measured. If playing from file the generator must still be showing with the same signal selected so that REW knows what it should calculate.

The generator provides preset signals for SMPTE, DIN, CCIF and AES17-2015 intermodulation measurement signals and a 'Custom' option allowing a user-selected pair of frequencies at a 1:1 or 4:1 amplitude ratio. Signals at 1:1 ratio will begin to clip at -3.0 dBFS (with the View option **Full scale sine rms is 0 dBFS** selected, 3 dB lower otherwise). Signals at 4:1 ratio will clip at -1.8 dBFS.

Dither should be enabled on the generator at the bit width the system is using, check the bottom left corner of the REW main window for the bit width in use. On Windows using Java drivers only 16 bit is supported, for 24-bit use ASIO drivers. On macOS 24-bit is supported, make sure the devices are configured to operate at 24-bit in Audio Midi setup.

To obtain accurate results the RTA window should have very low side lobes. Good choices in order of reducing side lobe level are Dolph-Chebyshev 150 (side lobes 150 dB down), Blackman-Harris 7 (side lobes 180 dB down), Dolph-Chebyshev 200 (side lobes 200 dB down) and Cosine sum 9-235 (side lobes 235 dB down). Do not use the Hann, Blackman-Harris 4 or Flat-Top windows.

The CCIF figure is calculated from the level at f2 - f1 (also called the 2nd order Difference Frequency Distortion or DFD2). The 3rd order DFD3 figure based on the levels at 2*f1 - f2 (18 kHz) and 2*f2 - f1 (21 kHz) is also shown. The reference level for the DFD figures is the sum of the level at f1 and the level at f2. An IMD_{pwr} figure is also shown, which is the ratio of the rms sums of the IMD components to the rms sum of f1, f2 and the IMD components.

The TDFD IMD_{TDFD} is calculated from the rms sum of 2nd order and 3rd order components, the reference level for the percentage figure is the sum of the levels at f1 and f2. For TDFD Phono and TDFD akl d2L (f2 - f1) and d3L (2*f1 - f2) are used, for TDFD Bass d2H (f2 + f1) and d3H (2*f2 - f1) are used.

The AES17 DFD IMD_{AES} figure is based on the levels at f2 - f1 (2 kHz), 2*f1 - f2 (16 kHz) and 2*f2 - f1 (22 kHz), the reference level for the IMD_{AES} percentage figure is the level at f1 (18 kHz). In all cases levels are measured across a 500 Hz bandwidth centred on the component being measured, per the AES17-2015 specification. DFD2 and DFD3 are also shown.

The AES17 MD IMD_{AES} is calculated from the rms sum of the 2nd order (d2) components, the reference level for the percentage figure is the level at f2. REW displays the overall IMD figure and the combined 2nd order (MD2 = d2L + d2H) and 3rd order (MD3 = d3L + d3H) figures.

For signals other than AES17 MD with an f2/f1 ratio > 7 (including SMPTE and DIN) the IMD_{DIN} figure is calculated from the rms sum of the 2nd order (d2) and 3rd order (d3) components, the reference level for the percentage figure is the level at f2. REW displays the overall IMD figure and the combined 2nd order (MD2 = d2L + d2H) and 3rd order (MD3 = d3L + d3H) figures.

In all cases REW also shows a total distortion + noise percentage, TD+N. This figure is the square root of the ratio of the noise and distortion powers to the power of the tones.

DFD Components		MD Components	
Component	Freq	Component	Freq
d2L	f2 - f1	d2L	f2 - f1
d2H	f2 + f1	d2H	f2 + f1
d3L	2*f1 - f2	d3L	f2 - 2*f1
d3H	2*f2 - f1	d3H	f2 + 2*f1
d4L	3*f1 - 2*f2	d4L	f2 - 3*f1
d4H	3*f2 - 2*f1	d4H	f2 + 3*f1
d5L	4*f1 - 3*f2	d5L	f2 - 4*f1
d5H	4*f2 - 3*f1	d5H	f2 + 4*f1

Multitone Total Distortion

When the system being monitored is driven using REW's Multitone test signal the RTA shows a figure for total distortion + noise, TD+N. This figure is the square root of the ratio of the power of the noise and distortion components to the power of the tones. The signal generator and signal capture must be operating at the same sample rate for correct results, through the same device or devices with synchronised sample clocks. The RTA window type must be set to Rectangular for this calculation, REW will select that if it is not already selected.

A signal-to-noise ratio figure is also displayed if the FFT is two or more times the signal length, in which case the tones will only occupy even bins and anything in odd bins will be noise. The noise figure is obtained by

rms summing the odd FFT bins and multiplying by sqrt(2), SNR is the ratio of the total rms level to that unweighted noise level.

TD+N: **0.0070 %**

SNR: **86.7 dB**

Stepped sine distortion

When the **Stepped Sine Measurement** button is pressed a dialog appears to configure and run a stepped sine distortion measurement, stepping in either level or frequency.

REW's signal generator is used to produce the measurement frequency. **Note** that the currently applied signal generator settings are used, so dither will only be applied if selected (it is selected by default and is recommended). The frequency can be stepped in intervals of between 1 and 96 points per octave over the selected span, level can use steps down to 0.1 dBFS. The nominal test frequencies will be the preferred values at the chosen measurement PPO that are within the span. To avoid scalloping loss effects the test frequencies use the closest FFT bin frequency, that ensures the peaks of the fundamental and all harmonics

are captured in the plots.

At each step all the distortion data is captured, when all points have been captured a new measurement is generated. The levels of the fundamental (brown), 2nd harmonic (red), 3rd harmonic (orange) and THD are shown on the RTA graph while the measurement progresses. If the **Stop** button is pressed a measurement is generated from the data collected up to that point. Stepped sine measurements typically take many minutes. The **Pause** button pauses the measurement, turning off the signal generator. Press it again to resume the measurement. The **Back** button removes the last point measured and re-measures it. **Back** can be used when the measurement is running or paused. **Cancel** discards the measurement.

When stepping in frequency a 96 PPO log spaced copy of the spectrum data at each measurement step will be captured if **Capture spectrum at each step** is selected. That spectrum data can be viewed on the [Waterfall] or [Spectrogram] graphs or [exported] to a text file. Note that exported spectrum data should not be used to try and calculate distortion levels as it does not have sufficient frequency resolution to accurately calculate the energy in each harmonic.

If **Stop measurement if heavy input clipping occurs** is selected the measurement will be stopped if more than 30% of the samples in an input block are clipped. That corresponds to an input level about 2 dB above the clipping threshold.

If **Stop if THD percentage exceeded** is selected the measurement will be stopped if the THD goes beyond the limit entered. To arm this THD must first drop below the limit, that allows stepped level measurements to start at low levels, where the THD figure may be high purely due to noise when the signal level is very low.

When stepping in frequency the minimum start frequency depends on the FFT length and the sample rate - for example, for an 8192 point FFT and 44.1 kHz sample rate the minimum start frequency is approx 60 Hz, for a 32768-point FFT at 44.1 kHz the minimum start is approx 15 Hz. The spreading effect of the RTA window would obscure the 2nd harmonic level at frequencies lower than the minimum and prevent a valid reading of distortion. The measurement frequencies are chosen so that they correspond to bin frequencies for the selected FFT length, this prevents window scalloping loss from affecting the amplitudes of the fundamental or harmonics.

At the beginning of each measurement the noise floor is captured and used to (optionally) mask distortion results that are below the noise floor (see the [distortion graph help]).

The progress bar shows the approximate time remaining to complete the measurement. Stepped sine measurements are faster when using ASIO drivers as the input and output buffers are smaller than when using Java drivers, hence less time is required for the buffers to flush through when changing frequency. The distortion results can be viewed on the Distortion graph. If the [Analysis Preference] [Apply cal files to distortion] is selected the results will include corrections for the cal file responses (as is the case for the RTA distortion figures). Applying the cal files provides more accurate results in regions where the fundamental or harmonics are affected by interface roll-offs but boosts the noise floor in those regions. This should be borne in mind when viewing the results. If large cal file corrections are required make sure the [Analysis Preference] [Limit cal data boost to 20 dB] is not selected.

Although much, much slower than a log sweep the stepped sine measurement captures N (noise and non-harmonic distortion) and THD+N (neither is available with a log sweep) and can measure low distortion levels more accurately than a sweep, particularly at high frequencies and for higher harmonics. This makes it well suited to measuring the distortion of electronic components. The plots below show a soundcard loopback measurement at -12 dBFS measured with stepped sine (64k FFT, 24 ppo) and with 8 repetitions of a 1M log sweep. Note the rise in the levels of harmonics with frequency when measuring with the sweep. This reflects the rise in the noise floor of the device (as can be viewed with the RTA in RTA mode), the sweep cannot separate distortion from noise as well as the stepped sine measurement.

[Help Index](#)
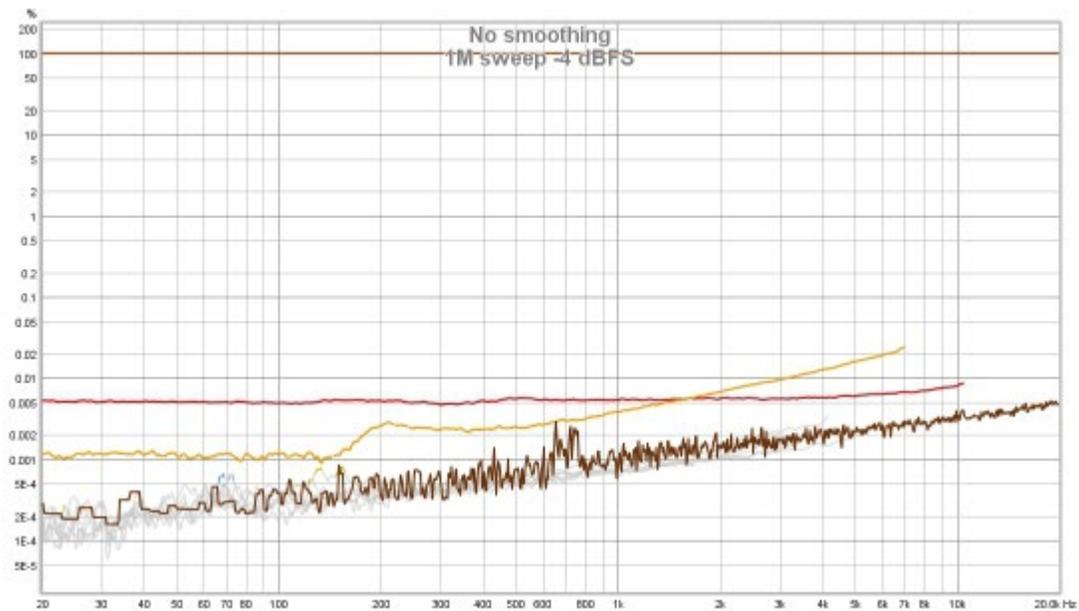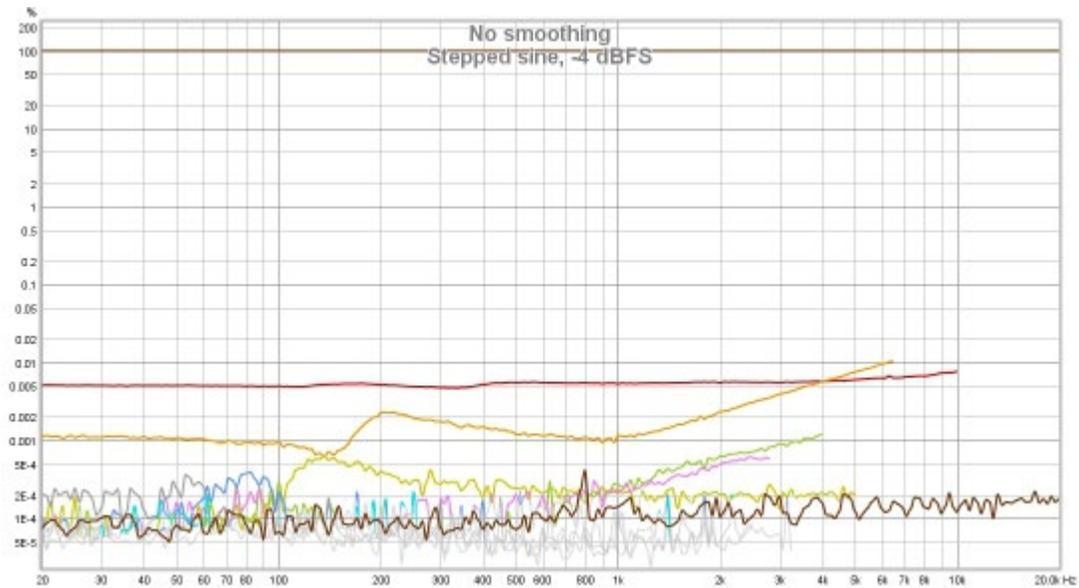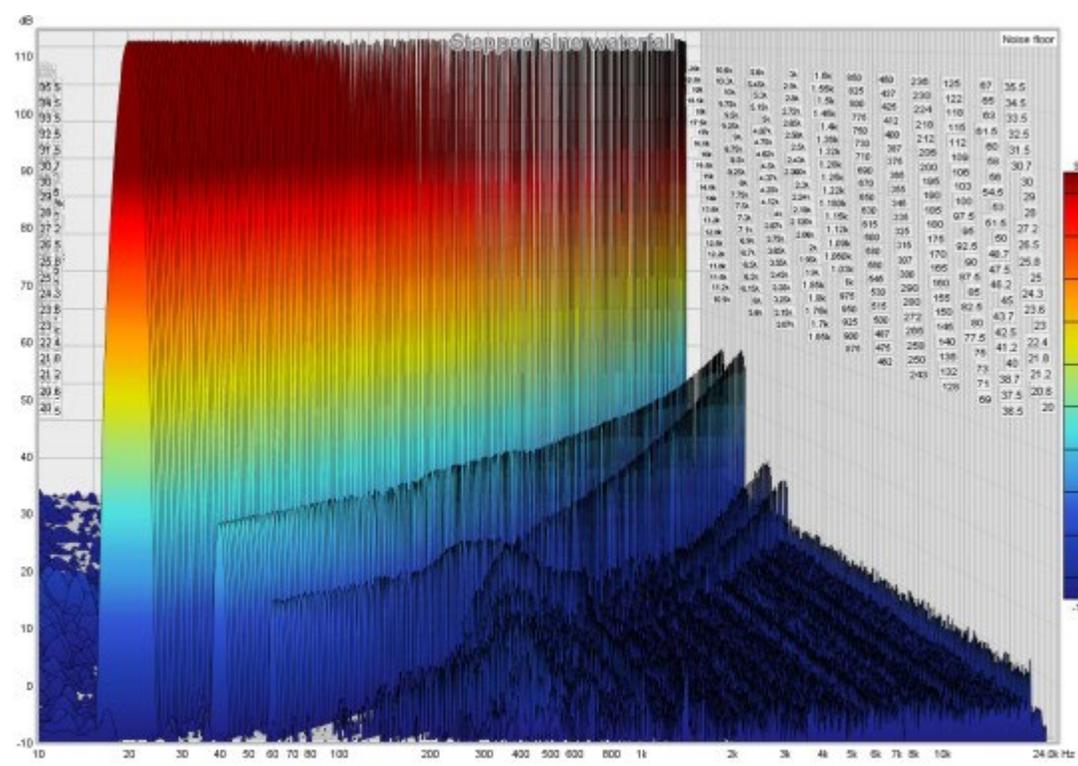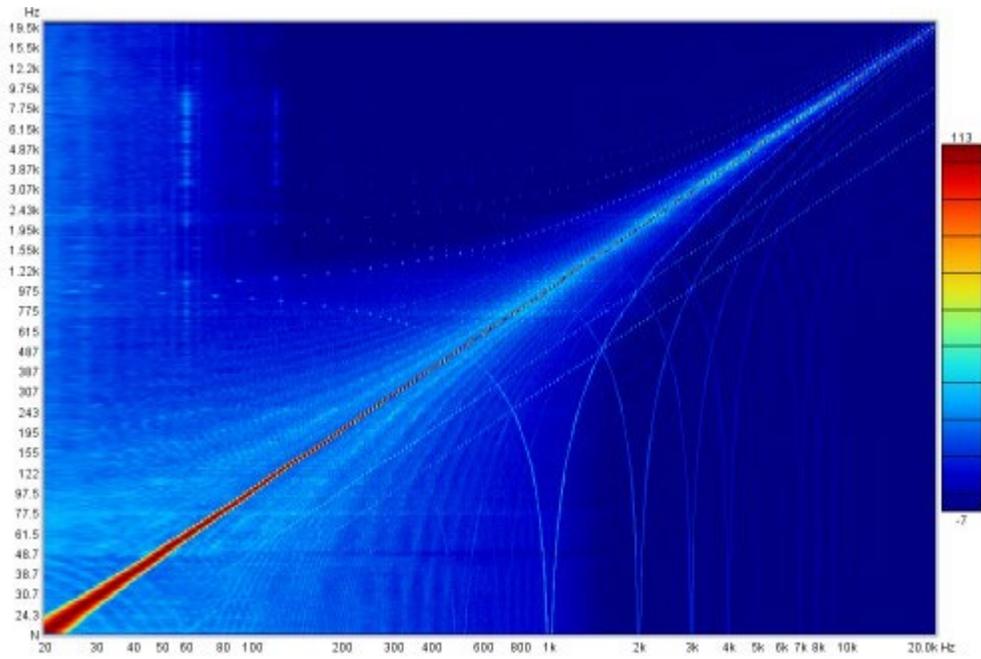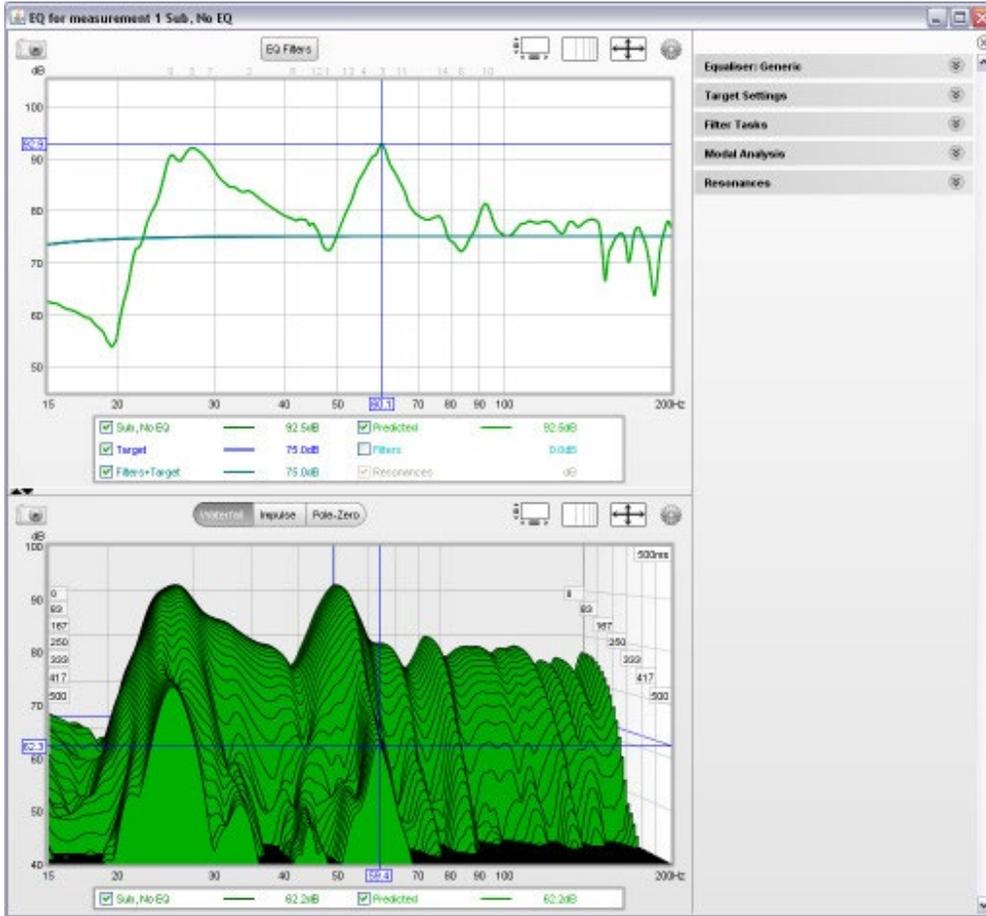
EQ Window

The EQ Window is used to determine what EQ filters to apply to a response and to see the effect those filters would have on both the frequency and time domain behaviour. It always shows the response currently selected in the main REW window, which can be changed from the REW main window or by pressing ALT + a measurement number (e.g. Alt+3 selects the third measurement) or using ALT+UP/ALT+DOWN to move through the measurements.

The window has 3 main areas: a "Filter Adjust" graph of frequency responses, a second graph area showing the impulse response and waterfall, and a panel on the right with various settings related to the EQ functions and modal analysis. The right hand panel can be hidden/shown using the button at the top of the scroll bar.

Filter Adjust

The Filter Adjust plot shows the measured and predicted (equalised) response for the current measurement along with the target response and the response of the equaliser filters with and without the target. This plot, in common with all plots that have a frequency axis, also optionally shows where any filters have been defined, displaying the filter's number along the top margin of the plot at the position corresponding to its centre frequency.
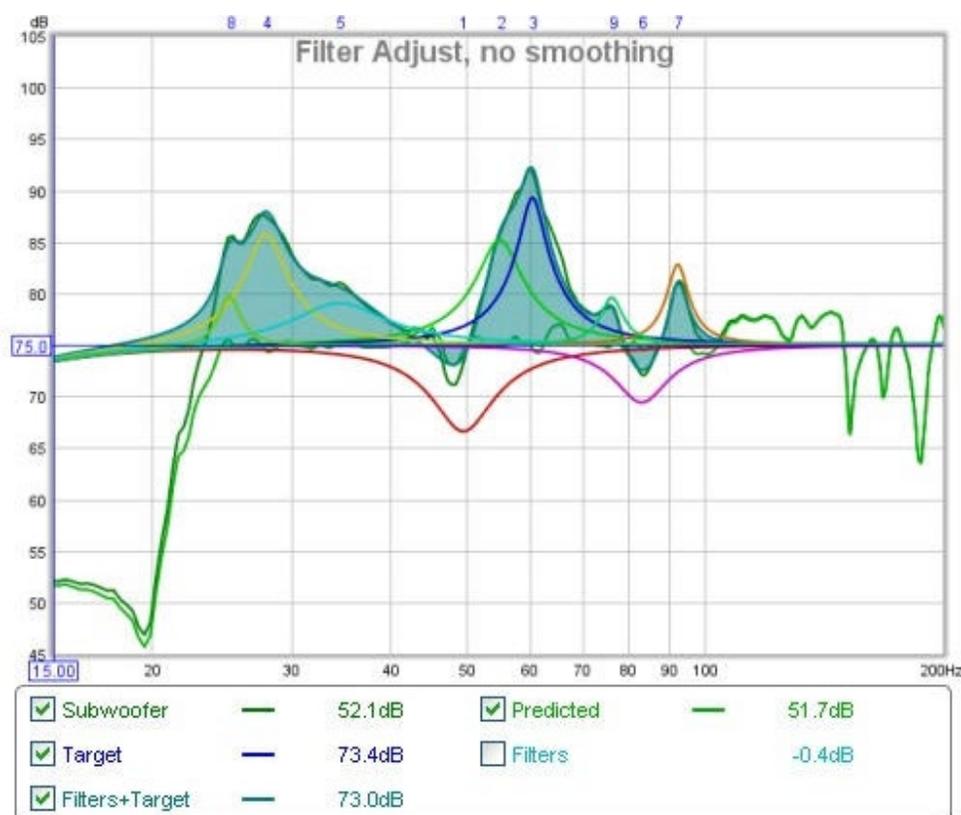
The frequency response of the measurement is labelled with the measurement name. The **Predicted** response shows the predicted effect of the measurement's filters. The **Target** trace shows the target frequency response for the measurement, including any desired House Curve response shape. If a House Curve has been loaded the ⌂ symbol will be displayed by the trace value. The Target response includes the Bass Management curve appropriate to the speaker type selected for the measurement in the Target Settings. The **Filters** trace shows the combined frequency response of the filters for this measurement, along with the individual filter responses if this has been selected (see Filter Adjust Controls below). The **Filters+Target** trace shows the frequency response of the filters overlaid on the Target response. Selecting the filter responses to be drawn inverted and adjusting the filters so that this curve matches the measured response will result in the predicted response matching the target.

Filter Adjust Controls

The control panel for the Filter Adjust graph has these controls:

The smoothing selector operates in the same way as those on the other graphs. When **Invert filter responses** is selected the responses of the filters are drawn inverted. This is useful for graphically matching the shape of a filter to the shape of the peak it is being used to correct, when the shapes match the overall response in that region will be flat. **Fill filter responses** fills the overall filter response. **Show each filter** draws the individual

filter response shapes separately in different colours. **Fill each filter** fills the individual responses. **Show filter numbers** controls whether the numbers of each filter are shown along the top of the graph, while **Show modal frequencies** controls whether the positions of the modes corresponding to the room dimensions entered in the Modal Analysis panel are shown along the bottom of the graph.

EQ Filters Panel

The EQ Filters panel is displayed by clicking the button at the top of the EQ window.

EQ Filters

Waterfall

The Waterfall plot shows a waterfall for the measurement and for the predicted result of applying the current filters to the measurement. The Predicted waterfall can be configured to update automatically as filters are adjusted (see Waterfall Controls below).
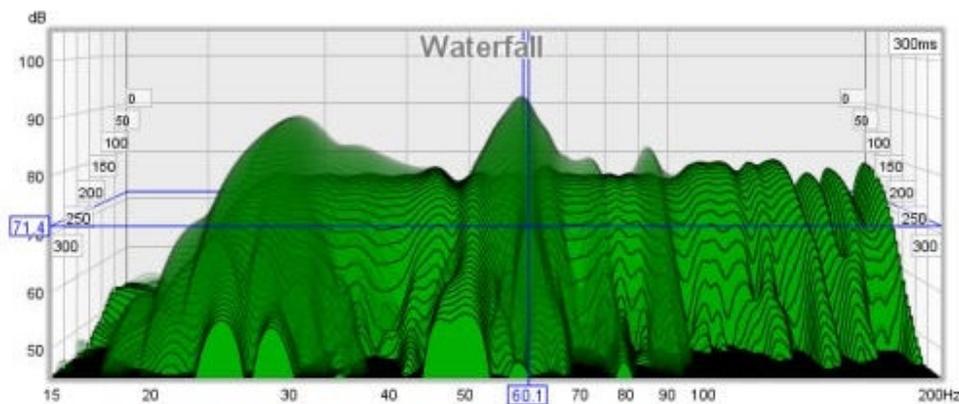
Waterfall Controls

Total slices:	101
Window (ms):	500
Time Range (ms):	300
Rise Time (ms):	100

☑ Auto generate
☐ Show modal frequencies
☑ Predicted Waterfall Live Update

Smooth 1/48

Perspective settings
Colour scheme settings
Apply defaults

The Predicted plot can be overlaid on the current measurement. The overlay is generated slice-by-slice, plotting a slice of the current measurement's waterfall, then a slice of the overlay, then the next slice of the current

measurement and so on.

If **Predicted Waterfall Live Update** is selected the waterfall will be regenerated as filters are adjusted - it may take a few seconds for the update to appear, depending on the speed of the computer and the frequency span of the measurement.

The **Total Slices** control determines how many slices are used to produce the waterfall. Fewer slices mean faster processing, but make it less easy to see how the response is varying over time.

The width of the impulse response section that is used to generate the waterfall is set by the **Window** control (this control sets the Right Hand window width). The corresponding frequency resolution is shown to the right of the window setting. Longer window settings provide better frequency resolution.

The **Time Range** control determines how far the impulse response window is moved from its start position to generate the waterfall.

The **Rise Time** control sets the width of the Left Hand window. Shorter settings give greater time resolution but make the frequency variation less easy to see. The default setting, 100 ms, is aimed at revealing room resonances. When examining drive unit or cabinet resonances with full range measurements a much shorter rise time would be used, 1.0 ms or lower, with time spans and window settings of around 10 ms. CSD mode is often more useful for such measurements as the later part of the impulse response can be noisy, obscuring the behaviour in the later slices.

The **Smoothing** applied to the waterfall slices can be increased from 1/48th octave (the minimum, and recommended) to as high as 1/3rd octave.

The **Perspective settings** control the appearance of the plot including a **Transparency** setting which can be applied to the main plot, the Predicted overlay, or both. When transparency is set to 0% both plots are solid. The transparency mode can be switched between main/overlay/both to ease comparison between the plots.

The control settings are remembered for the next time REW runs. The **Apply Default Settings** button restores the controls to their default values.

Impulse

The Impulse plot shows the impulse response of the measurement and of the predicted result of applying the current filters to the measurement.
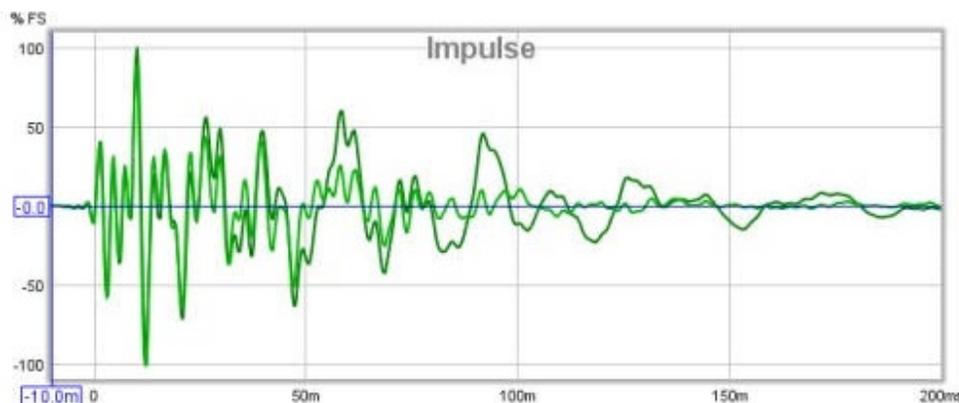
EQ Settings

The area to the right of the graphs contains a group of collapsible panels containing settings that affect the EQ functions.

Equaliser Panel

The Equaliser panel is used to select the type of equaliser that will be applied to the current measurement. Changing the equaliser type updates the filter panel, applying the settings appropriate to the selected equaliser. Filters already defined are retained where possible, but parameter values will be adjusted if necessary to comply with the ranges and resolutions of the chosen equaliser. The currently selected equaliser is shown in the panel title and in the EQ Filters panel. Details of the various equaliser types can be found here.

Target Settings

The Target Settings panel is used to tell REW what you expect or want the response to look like, so it knows what to aim for when applying EQ. The first selection (Target Type) should reflect the kind of speaker the measurement is from. If it's a Full Range ("Large") speaker the target is flat with a configurable low end roll off frequency (LF Cutoff) and slope (LF Slope). Setting the LF Cutoff to zero results in a target response that remains flat to 0Hz.

If it's a Bass Limited ("Small") speaker the target includes the effect of the bass management filter, the Bass management slope setting lets you tell REW how steep the bass management filter is and Bass management cutoff is the frequency it is set to, typically 80Hz in Home Theatre systems.

The Subwoofer target is similar, except the high frequencies are rolled off and there is a configurable low end roll off frequency and slope. For a 'Full Range' speaker the LF cutoff might be 40 Hz, for a sub it might be 20 - 30 Hz. You can usually see from the measurement where it is rolling off and adjust these settings for the target until it matches.

The bass management slope would typically be 24dB/octave for a subwoofer and 12dB/octave for a bass limited speaker, however the 12dB/octave figure for a speaker is used because the speaker itself is expected to have around a 12dB/octave acoustic roll-off, hence the overall effect of the filter and the speaker's roll-off is around 24dB/octave - the 24dB setting may be a better match to the measured response in those cases.

There is also a **Speaker driver** target type to use with measurements of individual drive units. Low pass and high pass crossover filters of up to 8th order can be configured to form the target.

An Add room curve option allows the target to to include typical room effects at the listening position and, if desired, a boost at low frequencies. HF Fall is used to reflect the downward tilt at high frequencies which is normal for most listening position speaker measurements, the result of the room's absorption and the speaker's power response. LF Rise allows that tile to be extended to low frequencies, to target a higher bass level which may be subjectively preferred. The target curve will rise below the LF Rise start frequency at the slope selected until reaching the LF Rise end frequency. Similarly, the target curve will fall above the HF Fall start frequency at the slope selected. The room curve effect can be turned on and off using the Add room curve box.

The Target Level control lets you move the whole target response up or down until it sits in the right place relative to your measurement. When the target is right the bits that go above it are the peaks you want to tame and it usually runs more or less through the middle of the measurement. **Calculate target level from response** automatically adjusts the level of the target response to provide a good match to the measurement over a suitable frequency range for the target type, but don't be afraid to manually adjust the level to suit.

Generate measurement from target shape creates a new measurement whose response matches the currently configured target. That measurement may be used on the RTA or the ALL SPL graphs to act as a reference.

The default target type, bass management slope, cutoff, LF rise and HF fall to use for new measurements are specified in the Equaliser Preferences.

Filter Tasks

The **Filter Tasks** panel is used to control REW's automatic filter adjustment feature. REW can automatically assign and adjust filter settings to match the Predicted response to the target response.

The Match Range defines the frequency span over which REW attempts to match the target response, and within which filters will be assigned. REW can apply filters anywhere across the band, but it is usually best to limit filters to low frequencies (less than 200 Hz or so) unless you are compensating for some general characteristic in the speakers (an example might be a dip in the mid range or a bit too much HF) - that is using

EQ as a fancy tone control.

Individual Max Boost sets the maximum boost that REW will allow for any individual filter. This can be set to zero to prevent REW assigning any boost filters.

Overall Max Boost sets the maximum boost that REW will allow for the combined effect of all the filters. This can be set to zero to prevent REW allowing any overall boost, but individual filters may still have boost. If the filter set has any overall boost the headroom this requires will be shown on the EQ filters panel.

In addition to the gain limits, boost filters are subject to Q limits to avoid inadvertently creating artificial resonances. The Q of boost filters is not allowed to exceed a value which would cause the filter's 60dB decay time to exceed approximately 500 ms (the actual Q limit value depends on the filter's gain).

The Flatness Target controls how tightly REW tries to match the Predicted response to the Target Response. The lower the Flatness Target, the more filters will be required.

Allow narrow filters below 200 Hz determines whether target match uses filters narrow enough to counter modal resonances at low frequencies. This should be selected when applying EQ to a room measurement, but is best not selected when applying EQ to a device response (headphone EQ, for example). When this option is not selected the highest Q used will be 5.0.

Match Response to Target starts REW's automated filter assignment and adjustment process. REW assigns filters to match the Predicted response to the Target Response, beginning with the area within the Match Range where the measurement is furthest from the target. After assigning filters, REW adjusts the settings of the filters to get the closest match. It is best to apply the 'variable' smoothing to the response before running the target match.

For best results it is essential to first ensure the shape of the target response is correctly selected to suit the type of speaker whose response is to be equalised and set the Target Level so that REW does not end up applying filters to try and correct a level difference - equalisers are not volume controls!

Note that by default REW will not apply filters below the frequency at which the measurement first exceeds the target or above the frequency at which the measurement last drops below the target to prevent trying to boost a response beyond its natural roll-offs, if you wish to lift the low or high end response this can be done with manually applied filters but beware of exceeding the excursion limits or headroom of the woofer or power handling limits of the tweeter.

The Filter Tasks panel also includes a set of controls to optimise the settings of the current filters. Note that only filters that lie within the **Match Range** will be adjusted. **Optimise gains** will adjust the gains of all 'Automatic' PK and modal filters to best match the target response. **Optimise gains and Qs** will adjust the gains and Qs of all 'Automatic' PK filters and the gains of all 'Automatic' modal filters. **Optimise gains, Qs and frequencies** will adjust the gains, Qs and centre frequencies of all 'Automatic' PK filters and the gains of all 'Automatic' modal filters - it is equivalent to **Match response to target** without the automatic assignment of filters. Centre frequencies will be adjusted to within 10% of their initial setting and will remain within the match range.

If REW can read from the equaliser **Retrieve filter settings from equaliser** will be enabled, selecting it will read the settings either directly from the equaliser or from a file exported by the equaliser, depending on the equaliser type. **Send filter settings to equaliser** will transfer the current filter settings to the equaliser if the equaliser offers that capability. **Save filter coefficients to file** writes the biquad coefficients for the current filters at the selected sample rate in a form the equaliser can import. **Export filter settings as text** generates a text file with the filter types and settings. **Reset filters for current measurement** will clear all the filters.

Generate measurement from predicted creates a new measurement whose response matches the predicted effect of any EQ filters. It has the same name as the current measurement prefixed by "EQ".

Generate measurement from filters creates a new measurement whose response matches the current EQ filters. It has the same name as the current measurement prefixed by "Filters". If there are no active filters the result is a perfect impulse.

Modal Analysis

REW can analyse the low frequency part of the measured response to search for modal resonances. The search is controlled by the settings in the **Modal analysis** panel. To determine the modal characteristics a parametric analysis of a segment of the impulse response is carried out to identify the frequencies, amplitudes and rates of decay of the resonant features that make it up. Such an analysis is not constrained by the frequency resolution limits of an FFT, allowing precise values for each mode's parameters to be determined. However, the accuracy of the results depends on the signal-to-noise ratio of the measurement. The better the measurement, the better the results. To get the highest measurement quality for modal analysis set the sweep end frequency to match the highest frequency of interest, use the longest sweep and adjust levels so that the peaks of the captured signal are around -6 to -12 dB.
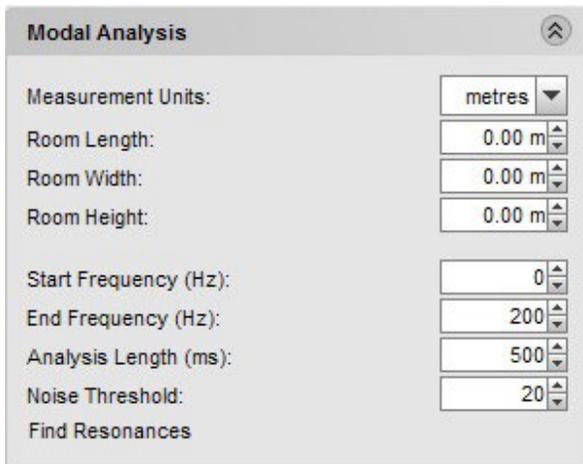
The modal analysis panel includes selections for the room dimensions. These are used to determine the room's theoretical modal frequencies up to 200 Hz, which may be plotted on the SPL & Phase, Group Delay, Spectral Decay, Waterfall and Spectrogram graphs. If any dimension is set to zero the corresponding modal frequencies will not be plotted, if all the dimensions are zero no modal frequencies will be plotted. The colours used for the modal frequencies are the same as those used in the Room Simulator.

The panel controls select the range to search for resonances (which will be restricted to the range of the measurement if smaller), the duration of the impulse response to analyse and a threshold for filtering out spurious resonances due to noise in the measurement. Best results are obtained by keeping the frequency span to around 100 - 200 Hz. The Analysis Length, 500 ms by default, may be reduced if the measurement is noisy or increased if the measurement has particularly low noise (noise floor of the impulse more than 60 dB below the peak).

Small alterations of the analysis length, 10-20 ms or so, can help establish whether the modal resonances identified are accurate - modes with consistent frequency, amplitude and decay time at differing analysis lengths indicate reliable data. When **Find Resonances** is clicked the analysis begins, it usually completes after a few seconds. The results are shown in the **Resonances** panel.

Resonances

Min T60 (ms):	300
Max T60 (ms):	3000
Min SPL (dB):	65

Clear Selection

Freq Hz	Peak dB	T60 ms
25.0	80.9	2,497
27.0	87.3	609
30.6	67.3	923
45.6	69.5	528
51.0	79.4	341
60.7	95.7	447
61.4	90.9	549
77.2	70.1	461
92.6	75.2	374
148.1	69.0	373
184.9	70.1	1,649

The Resonances panel includes controls to filter the results list according to the T60 decay times of the resonances and their amplitude. The list of resonances may be sorted by frequency, SPL ("Peak dB") or T60 decay time by clicking on the column headers in the table. Clicking on a resonance in the table will show a plot of its shape on the Filter Adjust graph, multiple resonances can be selected by clicking and dragging or using Ctrl+click or Shift+click. **Clear Selection** clears any selections made.
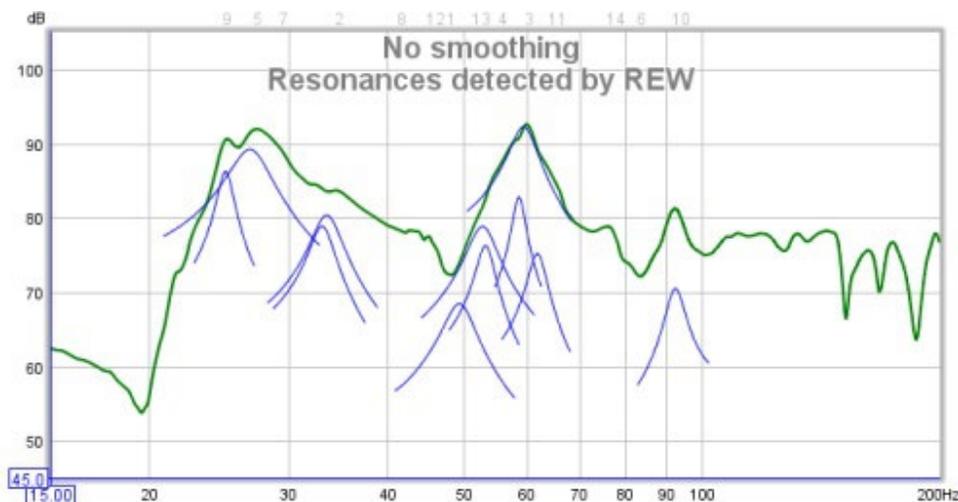
Filters which accurately counter specific resonances can be generated by selecting the "Modal" filter type and setting the Target T60 value to the T60 time determined by REW. Modal filters are normal parametric EQ filters whose Q or bandwidth is adjusted by REW as their gain is changed to ensure they target the specified T60 value as closely as the equaliser settings resolution permits.

Pole-Zero Plot

REW provides a Pole-Zero plot as an alternative way of viewing the results of the modal analysis. This may be an entirely unfamiliar way of viewing a response to many, but it does have some virtues when looking at resonances and filters. However, little would be lost by ignoring this section.
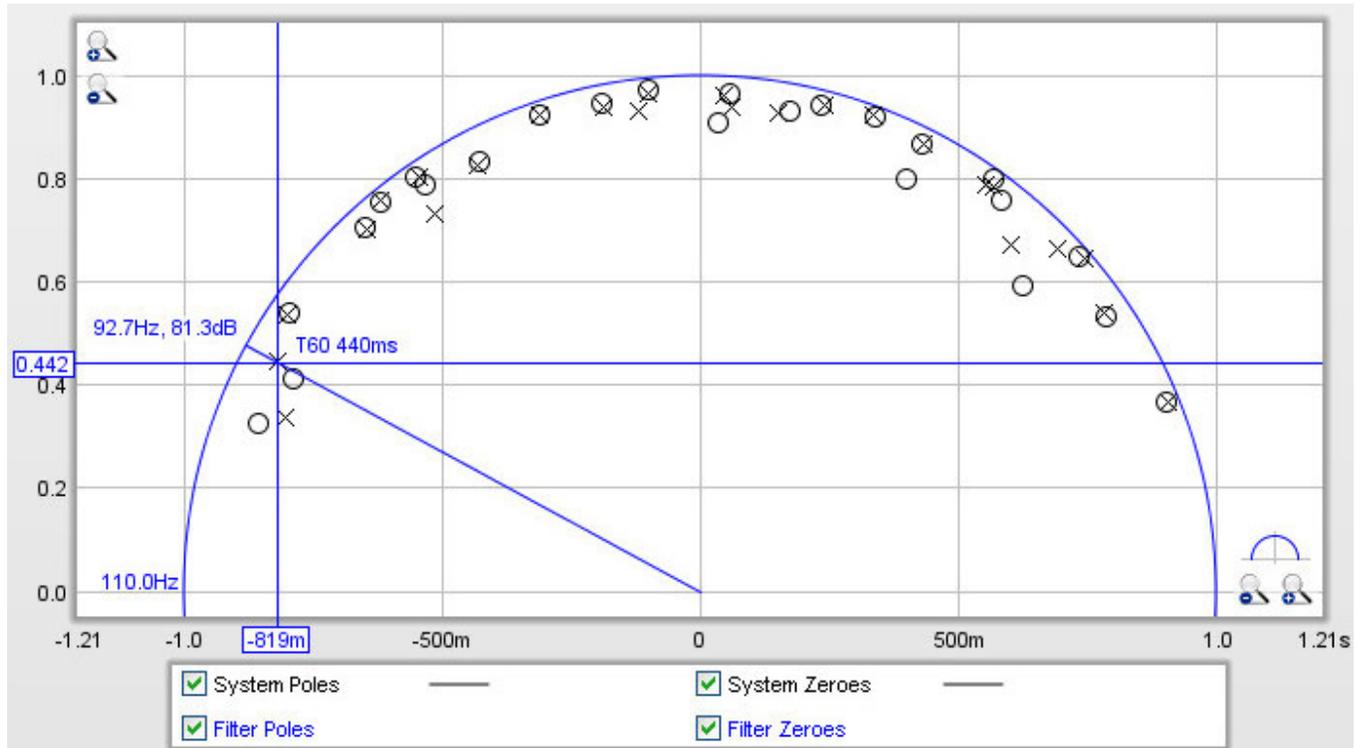
The pole-zero plot is a graph of complex numbers with the real part along the horizontal axis and the imaginary part along the vertical axis. There is a circle on the plot with a radius of one unit (referred to as the "unit circle") which corresponds in a way to the frequency axis of a frequency response. The plot shows results up to a frequency a little above the end of the modal analysis search, the upper frequency span of the plot is shown just to the left of the unit circle, near the point (-1, 0). As we move around the upper half of the unit circle the frequency increases from zero at the right side to the upper limit of the plot at the left. The lower half of the circle corresponds to negative frequencies, but for the signals we are looking at the bottom half is always a mirror image of the top half and can be ignored.

The plot shows poles, represented by crosses, and zeroes, represented by circles. Poles are places where the response becomes infinite, zeroes places where it becomes zero. The closer a pole gets to the unit circle, the more it pulls the frequency response upwards. Conversely, zeroes pull the response towards zero. Poles and zeroes at the same location cancel each other out completely, poles and zeroes close to one another partially counter each other's effects. If the plot has many pole/zero pairs that overlap they can be reduced by increasing the Noise Threshold setting. Poles outside the unit circle would correspond to an unstable system, none should appear there. Zeroes outside the circle would mean the response is not minimum phase, but the analysis may not start at the zero time of the impulse so this plot is not necessarily a good indicator of whether a response is minimum phase, for the correct method of determining that (using the excess group delay plot) refer to the Minimum Phase help topic.

Each modal resonance has a corresponding pole (actually a pair, the second is a mirror image below the axis). The frequency of the pole can be seen by drawing a line from the (0,0) point out through the pole to the point it reaches the unit circle, in the plot above the pole is at approximately 92.7Hz. REW shows the frequency value, and the SPL level at that frequency (81.3dB above). The closer a pole gets to the unit circle, the longer its T60 decay time. REW shows the T60 time corresponding to the cursor position, in the example above it is 440ms. If a resonance is selected in the Resonances panel its pole will be highlighted on the plot. The plot can be zoomed in on to get a closer view, clicking the button just above the x axis zoom buttons will reset the axis ranges to show the upper half of the unit circle.

Filters also have poles and zeroes, a parametric EQ filter has a pair of poles and a pair of zeroes (one pole and one zero above the axis, the other below). The locations of the filter's poles and zeroes vary as the filter's settings (frequency, Q/bandwidth and gain) are adjusted. If the settings of a filter are adjusted so that its zero is directly over the pole of a resonance, it completely counters the effect of that resonance in the time and frequency domains. Seeing how filter zero locations compare to response pole locations is where the pole-zero plot can be useful. In the case of the "Modal" filter type REW makes the adjustments that keep the filter's zero at a distance from the unit circle that matches the filter's target T60 time.

Filter poles and zeroes are shown in colour on the plot, corresponding to the colour used for that filter on the filters panel and the Filter Adjust plot. An example of a set of filters is shown below. Filters that cut (negative gain) have their zeroes closer to the unit circle than their poles, filters that boost have their poles closer to the unit circle than their zeroes.
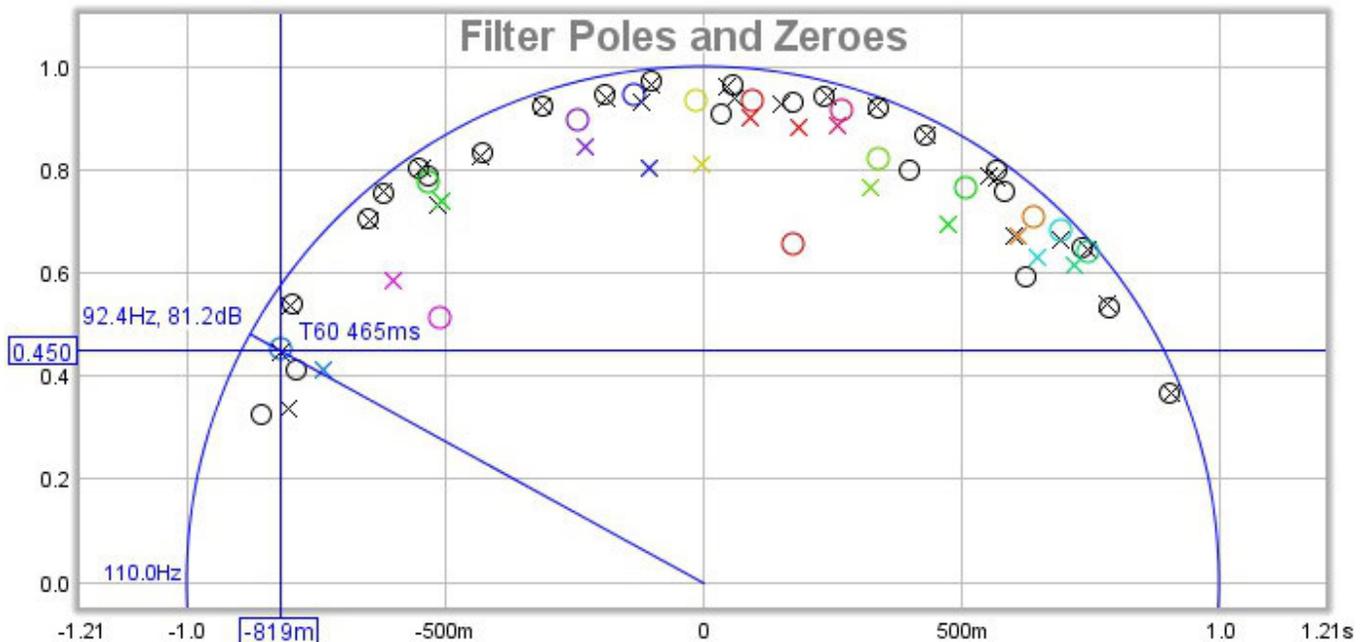
Pole-Zero Controls

Show cursor annotations controls whether REW draws a line from the origin through the cursor position to the unit circle and labels the frequency, response SPL and T60 values. If **Show 500ms T60 boundary** or **Show 1000ms T60 boundary** are selected REW will draw circles on the plot corresponding to those T60 times, any pole outside those circles has a T60 time greater than the circle value. If **Show Resonance Poles Only** is selected the pole-zero plot will only show the poles for the resonances displayed in the Resonances panel, otherwise it shows all poles found during the analysis.

[Help Index](#)

EQ Filters Panel

The EQ Filters panel is displayed by clicking its button at the top of the EQ window.

EQ Filters

The panel shows the filters settings for the current measurement. Buttons at the top of the panel allow the filter settings to be sorted, loaded, saved or deleted and the sort direction and key to be specified. At the right of the panel is a figure for the headroom the filter set requires, this will only be greater than zero if the filter set has overall gain. The equaliser type can be changed in the Equaliser selector at the right of the EQ window.

Each filter has:

- A check box to select/deselect it
- An identifying number and coloured line showing how the filter will be displayed when showing individual filter responses
- A "Control" setting which should be set to "Automatic" for filters REW is allowed to configure or "Manual" for filters it must leave unaltered.
- A "Type", which for TMREQ and Generic filters can be:
  - PK for a peaking (parametric) filter
  - LP for a 12dB/octave Low Pass filter (Q = 0.7071)
  - HP for a 12dB/octave High Pass filter (Q = 0.7071)
  - LS for a Low Shelf filter (S = 0.9)
  - HS for a High Shelf filter (S = 0.9)
  - NO for a notch filter
  - Modal for a Modal filter

  The Generic and DCX2496 also have shelving filters implemented per the DCX2496
  - LS 6dB for a 6dB/octave Low Shelf filter (S = 0.5)
  - HS 6dB for a 6dB/octave High Shelf filter (S = 0.5)
  - LS 12dB for a 12dB/octave Low Shelf filter (S = 1.0)
  - HS 12dB for a 12dB/octave High Shelf filter (S = 1.0)

  The Generic equaliser setting also has
  - LPQ, a 12dB/octave Low Pass filter with adjustable Q
  - HPQ, a 12dB/octave High Pass filter with adjustable Q

  For most other equalisers the only types available are PK and Modal, although the labelling for the PK filter varies. The miniDSP equaliser setting supports all the filter types that Generic supports.
- Centre Frequency/Corner Frequency, Gain and either Q, Bandwidth or Target T60 controls as appropriate for the filter type and selected equaliser. The filter bandwidth in Hz at the half-gain points is shown alongside the Q or Bandwidth control for PK filters.
- Displays of the 60dB decay time in milliseconds for a mode the current filter settings would match and the 60dB decay time of the filter itself, which is the decay which would remain after cancelling the decay of a modal resonance of the indicated modal decay. These correspond to the locations of the zeroes and poles of the filter.

The DSP1124P mode has an additional display showing frequency in the form in which is must be entered on that unit, i.e. as a one-third octave centre and a fine adjustment which ranges from -9 to +10 (63 -5 in the example below).

DSP1124P	Control	Type	Frequency		Gain	BW/60	Hz	Target T60	Mode T60	Filter T60
☑ 1 ——	Auto ▼	PA ▼	25.00 ⬍	25	-10 ⬍	4 ⬍	2.4		1,659	525

The Modal filter type is a peaking filter whose bandwidth or Q is adjusted by REW to match a Target T60 time, it is used to accurately counter a modal resonance whose T60 time is known. To match a specific T60 time the filter's bandwidth or Q must be altered as its gain or centre frequency change. REW chooses the bandwidth or Q setting supported by the selected equaliser that most closely matches the target T60.

Generic	Control	Type	Frequency	Gain	Q	Hz	Target T60	Mode T60	Filter T60
☑ 1 ——	Auto ▼	Modal ▼	24.90 ⬍	-9.8 ⬍	11.202 ⬍	2.2	1,739 ⬍	1,739	563

[Help Index](#)

Equaliser Selection

The Equaliser panel is used to select the type of equaliser whose responses REW is to model. Changing the equaliser type updates the filter panel, applying the settings appropriate to the selected equaliser. Filters already defined are retained where possible, but parameter values will be adjusted if necessary to comply with the ranges and resolutions of the chosen equaliser. The currently selected equaliser is shown in the panel title, with the sample rate at which it operates shown in brackets, and in the EQ Filters panel.

TMREQ

The TMREQ equaliser is a software EQ that runs at the sample rate of the measurement it is applied to. This setting offers the full range of filters and filter settings supported by TMREQ (peaking = parametric, low pass, high pass, low shelf, high shelf and notch). Filters can be transferred over a serial connection to and from AV32R or AV192R using the **Send filter settings to equaliser** and **Retrieve filter settings from equaliser** actions. For the Peaking filters the bandwidth in Hz between the half gain points is given by:

$$Bandwidth = centre\ frequency/Q$$

The TMREQ setting allows 8 filters. The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	20000	1 Hz
Gain	-15	+6	0.1 dB
Q	0.1	50	0.1

BFD Pro DSP1124P/FBQ1000

The DSP1124P equaliser setting supports the DSP1124P's parametric filters, allowing 12 filters. Note that this equaliser is now sold under as the FBQ1000. Filters can be transferred over a Midi connection using the **Send filter settings to equaliser** action. The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	20000	see below
Gain	-48	+16	1 dB
BW/60	1	120	1

The frequency control adjusts in the pseudo-1/60th octave steps DSP1124P supports (20 evenly spaced subdivisions of the ISO one-third octave intervals), with the one-third octave and fine adjustment values DSP1124P uses shown alongside the actual frequency in the EQ Filters Panel.

The "BW/60" control replicates the effect of the DSP1124P's bandwidth setting. This control sets the bandwidth of the filter between the half-gain points with:

$$\text{Bandwidth (Hz)} = \text{centre frequency}*(BW/60)*\sqrt{2}$$

For example, at a bandwidth setting of 60/60 a filter centred on 1kHz with a gain of -6dB will have a bandwidth of 1,414Hz between the points where its response crosses -3dB. This bandwidth remains constant as the filter's gain is adjusted (Note that the Behringer DSP1100 software package does **NOT** correctly reproduce the way the bandwidth control actually operates, its bandwidths are too small by a factor of sqrt(2)).

Defining filter bandwidth in this way is not uncommon (the TMREQ filters use a similar definition). The relationship between Q and BW for the DSP1124P is

$$Q = 60/[(BW/60)*\sqrt{2}]$$

so the bandwidth range of 1/60 to 120/60 gives a Q range from 42.4 to 0.35.

BFD Pro FBQ2496

The FBQ2496 equaliser setting supports the FBQ2496's parametric filters, allowing 20 filters. Filters can be transferred over a Midi connection using the **Send filter settings to equaliser** action but the transfer is not always reliable. The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	20000	1/60th octave
Gain	-36	+15	0.5 dB (1 dB below -15 dB)
Bandwidth	1/60	10	octaves, see below

The frequency control adjusts in approximately 1/60th octave steps (more precisely, 1/200th of a decade).

The bandwidth control adjusts in 1/60 of an octave steps from 1/60 to 5/60 of an octave, then goes through 1/10, 1/9, 1/8, 1/7, 1/6, 1/5, 1/4, 1/3, 1/2, 3/4, 1, 1.5, 2, 3, 4, 5, 6, 7, 8, 10 octaves.

The relationship between Q and BW in octaves for the FBQ2496 is

$$Q = sqrt(2)/BW$$

so the bandwidth range of 1/60 to 10 octaves gives a Q range from 84.85 to 0.14.

DCX2496

The DCX2496 equaliser setting supports parametric filters (labelled "BP" for Band Pass) and low and high shelving filters (with 6 and 12 dB/octave slopes). It also allows two crossover filters to be defined, LP or HP (Butterworth 6, 12, 18, 24, 48 dB/oct, Bessel 12, 24 dB/oct, Linkwitz-Riley 12, 24, 48 dB/oct). It allows up to 9 EQ filters per channel, depending on the other processing the unit is doing. The parametric filter bandwidth in Hz between the half gain points is given by:

$$Bandwidth = sqrt(gain)*centre\ frequency/Q$$

where gain is from the absolute value of the dB figure so always >= 1. The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	20000	106 steps per decade
Gain	-15	+15	0.1 dB
Q	0.1	10	20 steps per decade

The frequency control adjusts in steps of 1/106th of a decade. The Q control adjusts in steps of 1/20th of a decade, i.e. there are 20 Q values between 0.1 and 1 and another 20 between 1.0 and 10.

SMS-1

The SMS-1 equaliser setting supports parametric filters only, allowing 8 filters. The filter bandwidth in Hz between the half gain points is given by:

$$Bandwidth = centre\ frequency/Q$$

N.B. The SMS-1 filter shapes have not been verified against an actual unit.

The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	15	120	1 Hz
Gain	-13	+6	0.5 dB
Q	0.3	20.0	0.1

QSC DSP-30

The DSP-30 equaliser setting supports parametric filters only, allowing 20 filters. The filter bandwidth in Hz between the half gain points is given by:

$$\text{Bandwidth} = \text{centre frequency}/Q$$

N.B. The DSP-30 filter shapes have not been verified against an actual unit.

The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	20000	0.1 Hz
Gain	-120	+12	0.1 dB
Q	0.3	50	0.01

Crown USM-810

The USM-810 equaliser setting supports parametric filters only, allowing 10 filters. The filter bandwidth in Hz between the half gain points is given by:

$$\text{Bandwidth} = \text{sqrt(gain)*centre frequency}/Q$$

where gain is from the absolute value of the dB figure so always >= 1. **N.B. The USM-810 filter shapes have not been verified against an actual unit.**

The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	20000	1 Hz
Gain	-24	+24	0.1 dB
Q	0.1	35	0.01

miniDSP

The miniDSP equaliser setting supports the same filter types and resolutions as the Generic setting, but for 6 filters operating at 48 kHz. It is aimed at the miniDSP plug-in Advanced mode, which allows filters to be specified by their biquad coefficients. The **Save Filter Coefficients to File** action writes the filter coefficients to a file in a format suitable for use with the miniDSP software (note that the a1 and a2 coefficients are negated per the miniDSP format). An advantage of this is the very high filter frequency and Q resolution it allows, permitting exact targeting of modal resonances. The miniDSP plug-in has an **Import REW File** button on its Parametric EQ configuration screens to load the files.

miniDSP-96k

The miniDSP-96k equaliser setting supports the same filter types and resolutions as the Generic setting, but for 5 filters. It is aimed at miniDSP plug-ins that operate at 96 kHz. The **Save Filter Coefficients to File** action writes the filter coefficients to a file in a format suitable for use with the miniDSP software. The miniDSP plug-in has an **Import REW File** button on its Parametric EQ configuration screens to load the files.

miniDSP 2x4 HD

The miniDSP 2x4 HD equaliser setting supports the same filter types and resolutions as the Generic

setting, but for 10 filters at 96 kHz. The **Save Filter Coefficients to File** action writes the filter coefficients to a file in a format suitable for use with the miniDSP software. The miniDSP plug-in has an **Import REW File** button on its Parametric EQ configuration screens to load the files.

nanoAVR

The miniDSP nanoAVR equaliser setting supports the same filter types and resolutions as the Generic setting, but for 10 filters at 96 kHz. The **Save Filter Coefficients to File** action writes the filter coefficients to a file in a format suitable for use with the miniDSP software. The miniDSP plug-in has an **Import REW File** button on its Parametric EQ configuration screens to load the files.

C-DSP 8x12

The miniDSP C-DSP 8x12 equaliser setting supports the same filter types and resolutions as the Generic setting, but for 10 filters at 192 kHz. The **Save Filter Coefficients to File** action writes the filter coefficients to a file in a format suitable for use with the miniDSP software. The miniDSP plug-in has an **Import REW File** button on its Parametric EQ configuration screens to load the files.

waveFLEX DSP A8

The DSP A8 equaliser supports the same filter types and resolutions as the Generic setting, but for 5 filters operating at 96 kHz. The **Save Filter Coefficients to File** action writes the filter coefficients to a file in a format suitable for use with the waveFLEX software.

ADA PEQ

The ADA equaliser setting supports parametric filters only, allowing 12 filters. The filter bandwidth in Hz between the half gain points is given by:

$$\text{Bandwidth} = \text{centre frequency}/Q$$

N.B. The ADA filter shapes have not been verified against an actual unit.

The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	20000	1 Hz
Gain	-10	+10	0.5 dB
Q	0.1	10	0.1

Xilica XP2040

The XP2040 equaliser setting supports parametric filters only, allowing 16 filters. The filter bandwidth is specified in octaves, but the corresponding bandwidth in Hz is shown in the filter controls panel. The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	30000	1 Hz
Gain	-30	+15	0.25 dB

| | BW | 0.02 | 3.61 | 0.01 octaves |

Emotiva UMC-200

The UMC-200 equaliser setting supports parametric filters only, allowing 11 filters. The filter bandwidth in Hz between the half gain points is given by:

$$\text{Bandwidth} = \text{sqrt(gain)} * \text{centre frequency}/Q$$

where gain is from the absolute value of the dB figure so always >= 1. **N.B. The UMC-200 filter shapes have not been verified against an actual unit.**

The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	20000	256 steps
Gain	-15	+3	0.25 dB
Q	0.25	24	0.125

Emotiva XMC-1

The XMC-1 equaliser setting supports parametric filters only, allowing 11 filters. The **Save filter settings to file** action writes the filter settings to a file in a format suitable for use with the Emotiva software. The filter bandwidth in Hz between the half gain points is given by:

$$\text{Bandwidth} = \text{centre frequency}/Q$$

The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	20000	1 Hz
Gain	-64	+6	0.5 dB
Q	0.6	50	0.01

Emotiva RMC-1

The RMC-1 equaliser setting supports the same filter numbers and types as XMC-1 but offers further choices of channel when saving or loading .emo files. This setting may also be used with the RMC-1L and XMC-2.

Dual Core

The Dual Core equaliser setting supports parametric filters only, allowing 16 filters. The filter bandwidth in Hz between the half gain points is given by:

$$\text{Bandwidth} = \text{sqrt(gain)} * (\text{-3dB BW})$$

$$Q = \text{centre frequency}/(\text{sqrt(gain)} * (\text{-3dB BW}))$$

where gain is from the absolute value of the dB figure so always >= 1. The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	5	20000	1 Hz
Gain	-36	+12	0.1 dB
-3dB BW	1	10000	1

Rockford Fosgate 3Sixty.3

The 3Sixty.3 equaliser setting supports parametric filters only, allowing 20 filters. The filter bandwidth in Hz between the half gain points is given by:

$$\text{Bandwidth} = \text{sqrt(gain)*centre frequency/Q}$$

where gain is from the absolute value of the dB figure so always >= 1. The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	20000	0.1 Hz
Gain	-24	+24	0.1 dB
Q	0.1	6	0.01

rePhase

The rePhase equaliser setting supports parametric filters only. It allows up to 17 filters per channel. The **Save filter settings to file** action writes the filter settings to a file in a format suitable for use with the rePhase software. The parametric filter bandwidth in Hz between the half gain points is given by:

$$\text{Bandwidth} = \text{sqrt(gain)*centre frequency/Q}$$

where gain is from the absolute value of the dB figure so always >= 1. The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	10	24000	0.01 Hz
Gain	-96	+96	0.1 dB
Q	0.1	10	0.01

Outlaw Audio Model 976

The Model 976 equaliser setting supports 10 filters. The first filter is a low shelf, the last is a high shelf. The shelf filters have 12 dB/octave slope. The other 8 filters are peaking filters where the bandwidth in Hz between the half gain points is given by:

$$\text{Bandwidth} = \text{centre frequency/Q}$$

The filter adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	10	20000	1 Hz up to 200Hz, then 10 Hz up to 2 kHz, then 100 Hz
Gain	-20	+3	0.5dB
Q	0.5	14	Values are 0.5, 1.0, 1.4, 2.0, 2.4, 2.9, 3.6, 4.8, 7.2 and 14

StormAudio

The StormAudio equaliser setting supports 12 parametric filters and 2 crossover filters. The crossover filter types supported are Butterworth and Linkwitz-Riley up to 8th order (48 dB/octave). The **Save filter settings to file** action writes the filter settings to a file in a format suitable for use with the StormAudio software. The parametric filter bandwidth in Hz between the half gain points is given by:

$$Bandwidth = centre\ frequency/Q$$

The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	16	20000	1 Hz
Gain	-96	+48	0.25 dB
Q	0.25	32	0.25

Dutch & Dutch 8c

The Dutch & Dutch 8c Equaliser supports 24 filters, peaking and shelf which can be written to or read from the speakers over the local network. It also shows the speakers bass, treble and sub filter settings. On Windows a firewall prompt will appear to allow REW to access the network when the 8c equaliser is first selected. The peaking filter bandwidth in Hz between the half gain points is given by:

$$Bandwidth = centre\ frequency/Q$$

The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	15	20000	0.01 Hz below 100 Hz, 0.1 Hz below 1 kHz, 1 Hz above 1 kHz
Gain	-120	+3	0.1 dB
Q	0.1	50	0.001

REW scans the local network for speakers when the 8c equaliser setting is selected and, if the 8c is the default equaliser in the Equaliser preferences, on startup. There is also a **Scan network for speakers** action in the filter tasks pane that will trigger a rescan and show the results in a dialog. The **Send filter settings to speaker or group** action writes the filter settings to the selected speaker or group of speakers. The **Load filters from speaker or group** action retrieves the selected speaker or group filter settings. When initially retrieved the filters will be set to **Manual** control, to allow REW to use a filter for automatic target matching set the control to **Automatic**. Filters that were set to automatic when they were sent to the speaker will be set to automatic when they are retrieved.

If the 8c is the default equaliser in the Equaliser preferences the measurement dialog will offer an option to measure with playback from the speakers, which uses measurement sweep files which the speakers contain. When measuring with playback from the speakers REW will capture the filter settings from the speaker when it is measured, they can be viewed in the Info dialog. Those settings are retained in the measurement file and allow iterative adjustment of filter settings, taking into account the filters that were active when the measurement was made. If the settings have been captured the **Reset filters for current measurement** action in the filter tasks pane will offer a choice of resetting the filters to the settings they had when the measurement was made or to None.

When the 8c is selected as the equaliser the target settings will be set to:

- Speaker Type: Full Range
- LF Slope: 24dB/octave
- LF Cutoff (Hz): 20
- LF Rise Start (Hz): 100
- LF Rise End (Hz): 20
- LF Rise Slope (dB/octave): 1.0
- HF Fall Start (Hz): 1000
- HF Rise Slope (dB/octave): 0.0

and the target match range will be set to 20 .. 100 Hz, with a maximum match end frequency of 1000 Hz.

JL Audio TwK-88

The TwK-88 equaliser setting supports peaking filters only, allowing 10 filters. The filter bandwidth in Hz between the half gain points is given by:

$$Bandwidth = sqrt(gain)*centre\ frequency/Q$$

Where gain is less than one for cuts so cuts have narrower bandwidths for the same Q value as a boost. The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	20000	0.1 Hz
Gain	-12	+6	0.1 dB
Q	0.25	10	0.01

TimeZero2.6D/4.8D

The Elder Audio TimeZero2.6D/4.8D equaliser setting offers peaking, low shelf, high shelf, notch and all pass filters. For the Peaking filters the bandwidth in Hz between the half gain points is given by:

$$Bandwidth = centre\ frequency/Q$$

The TimeZero2.6D/4.8D setting allows 5 filters. The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	20000	0.1 Hz
Gain	-16	+12	0.5 dB

Q	0.31	19.4	0.01

PerListen Audio

The PerListen Audio equaliser setting supports parametric filters only, allowing 10 filters. The filter bandwidth in Hz between the half gain points is given by:

$$Bandwidth = centre\ frequency/Q$$

The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	10	200	0.1 Hz
Gain	-20	+3	0.1 dB
Q	0.3	20	0.01

AU N-Band EQ

The AU N-Band EQ is a software EQ that runs at the sample rate of the measurement it is applied to. This equaliser setting supports 16 filters. The parametric filter bandwidth in Hz between the half gain points is given by:

$$Bandwidth = centre\ frequency/Q$$

but bandwidths are entered in octaves.

The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	20	22000	0.01 Hz below 100 Hz, 0.1 Hz below 1 kHz, 1 Hz above 1 kHz
Gain	-96	+24	0.1 dB
BW	0.02	5	0.01 octaves

Generic

The Generic Equaliser is a software EQ that runs at the sample rate of the measurement it is applied to. It supports a full range of filters and filter settings (peaking = parametric, 1st and 2nd order low pass & high pass including adjustable Q 2nd order variants, low shelf and high shelf including adjustable Q variants, notch and all-pass) based on the Robert Bristow-Johnson 'Cookbook' equations. It also offers a pair of crossover filters with Bessel, Butterworth or Linkwitz-Riley characteristics up to 8th order.

The **Save Filter Coefficients to File** action writes the filter coefficients corresponding to the selected sample rate to a file in the selected format. Currently supported formats are miniDSP and SigmaStudio, note that the a1 and a2 coefficients are negated in the miniDSP format.

LPQ and HPQ filter types can be used to implement higher order filters as cascades of biquad filters. Here are the Q values for some common crossover types and orders, **for odd orders the last**

section is LP1/HP1 (a first order section), those are highlighted in blue:

2	Butterworth	0.707			
3	Butterworth	1.000	LP1/HP1		
4	Butterworth	1.307	0.541		
5	Butterworth	1.618	0.618	LP1/HP1	
6	Butterworth	1.932	0.707	0.518	
7	Butterworth	2.247	0.802	0.555	LP1/HP1
8	Butterworth	2.563	0.900	0.601	0.510
2	Linkwitz-Riley	0.500			
4	Linkwitz-Riley	0.707	0.707		
8	Linkwitz-Riley	1.307	0.541	1.307	0.541

For the Peaking filters the bandwidth in Hz between the half gain points is given by:

$$Bandwidth = centre\ frequency/Q$$

The Generic setting allows 20 parametric filters. The adjustment ranges are:

Parameter	Minimum	Maximum	Resolution
Frequency	10	22000	0.01 Hz below 100 Hz, 0.1 Hz below 1 kHz, 1 Hz above 1 kHz
Gain	-120	+30	0.1 dB
Q	0.1	50	0.001

[Help Index](#)

Room Simulator

The Room Simulator generates frequency responses for multiple sources at multiple locations in a rectangular room. It uses a frequency domain method based on the rigid boundary solution to the wave equation, modified for lossy boundaries. Results are equivalent to those obtained by the image source method in the time domain (Allen and Berkley 1978). Sources and listening positions can be altered by dragging on plan and elevation views of the room.

The Room Simulator window looks like this when first opened:
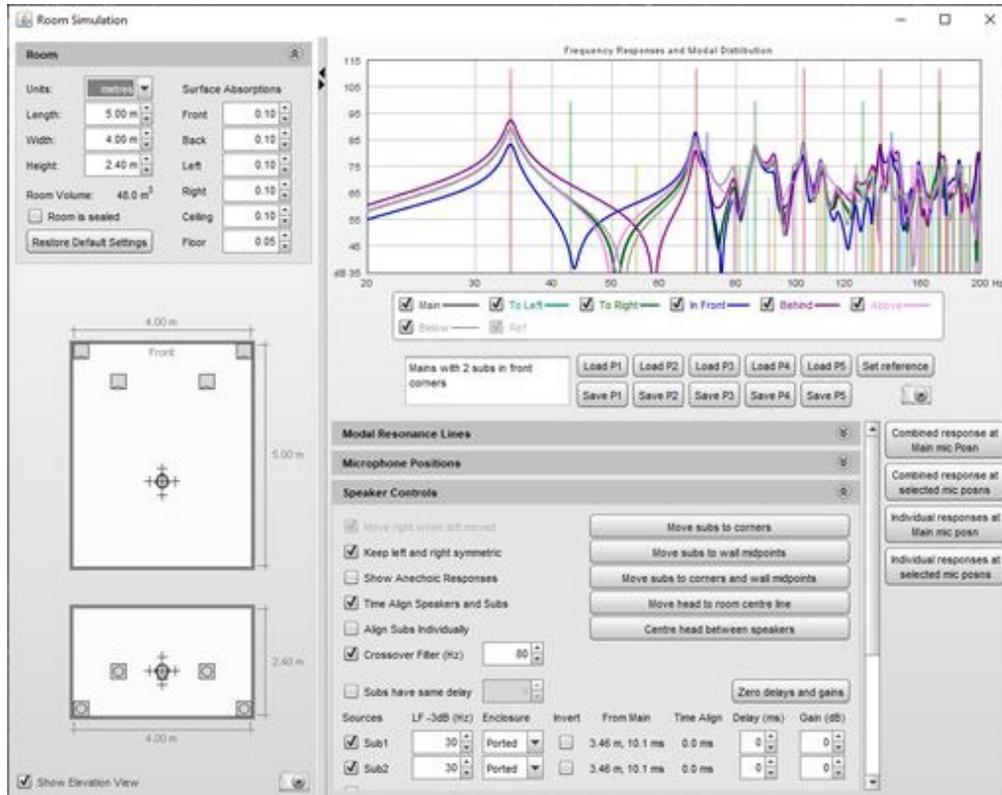
The left hand panel shows a view of the room with controls for room dimensions and the acoustic absorptions of the room's surfaces. The right hand side shows the frequency response at the main listening position and additional positions around it and has controls for which modal resonances are shown, the positions at which responses are to be calculated, the sources to be modelled and and how they are managed. The entire window can be resized and the divider between the left and right panels can be dragged to adjust the proportion allocated to each. The small triangles at the top of the divider allow either panel to be collapsed completely.

Room Panel

The dimensions and properties of the room are configured in the controls at the top of the room panel. The controls may be collapsed by clicking on the chevrons at the top right of the panel.

Dimensions may be displayed in metric or imperial units according to the selected **Units**. Regardless of the units selected, the dimension controls accept input in metric or imperial units, for example 2.5m, 250cm, 2500mm, 8.2ft, 8ft 2in, 8' 2", 8f2i, 8f2 and 98in are all valid entries. If an entry is a number without any units it is assumed to be in the selected measurement units.

If the room is well sealed select the **Room is Sealed** box, this increases the response boost at the lowest frequencies.

The surface absorptions define how sound is absorbed when it meets the surface. The absorptions are independent of angle or frequency. The higher the absorption figures, the more sound is absorbed at that surface and the more damped the room's modal resonances become.

Below the room panel controls are the views of the room, in plan and elevation.
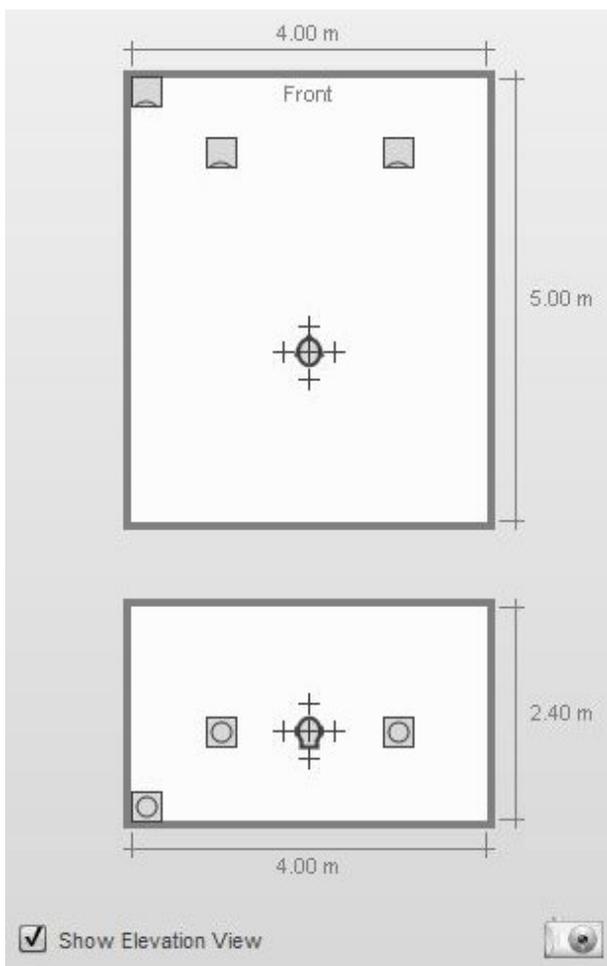
The main listening position is indicated by the head. Crosses around the head show the locations of any additional points selected for responses to be generated, in the image above the positions to left, right, in front, behind, above and below the main listening position have been selected. The elevation view can be hidden by un-ticking the **Show Elevation View** box at the bottom of the panel. An image of the room panel can be captured by clicking on the camera icon in the bottom right corner.

A source can be selected by moving the mouse cursor over it. The source will be highlighted and can be moved by left-clicking and dragging or by using the arrow keys, the arrow keys allow finer adjustment of position. If the shift key is held down while dragging, movement will be restricted to either horizontal or vertical only. Source outlines may overlap the room boundaries at large room sizes, the location of the acoustic centre is not affected. The source can be rotated by right-clicking or by pressing the R key (clockwise rotation) or L key (anticlockwise rotation). Note that rotating the source does not alter its response, all sources are treated as omnidirectional. The main listening position can similarly be moved using the mouse or, after highlighting it, the arrow keys, as can any of the additional listening positions. When a source or listening position is highlighted its location is shown:
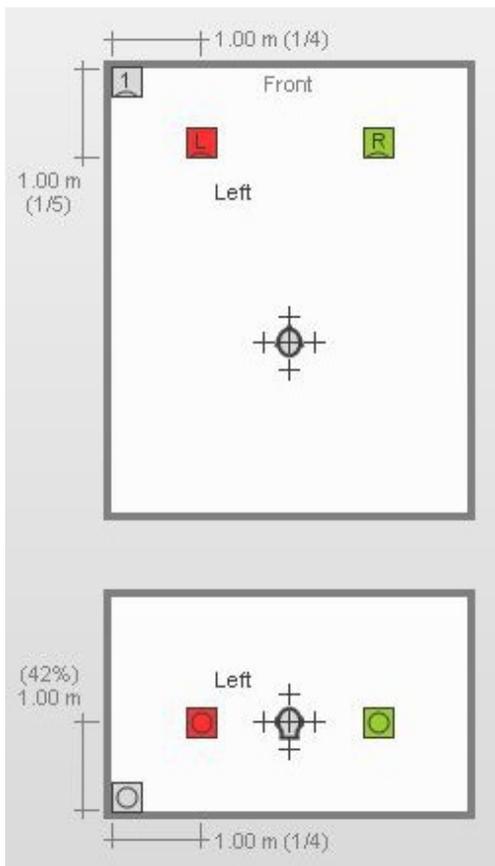
When a source is highlighted the dimensions shown are to the acoustic centre, which is located at the centre of the front face. Highlighting a source also shows its individual contribution to the combined response at the main listening position on the response graph.

Highlighting the head position when the simulation includes left and right speakers will show the angle between them at the head position.
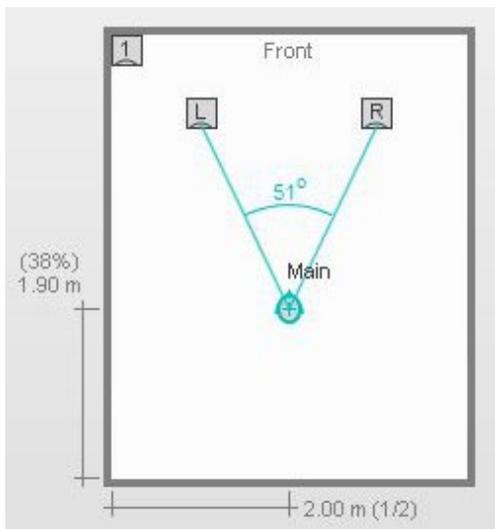
Response Panel

The frequency response at the main listening position and additional positions around it are shown on the response panel. An additional "Ref" trace can be set to retain the current main position response and act as a reference as changes are made. The modal distribution for the room is shown on the response panel using lines that are colour-coded according to the axes they include:
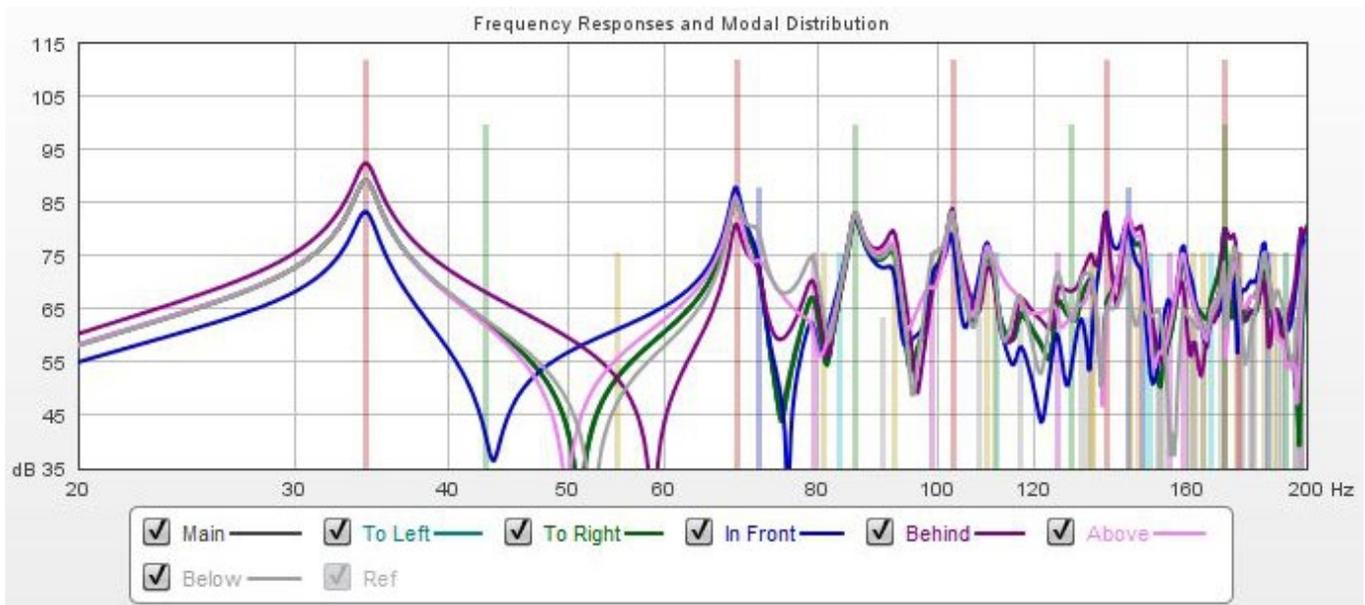
Presets

Five sets of room simulator configurations can be saved to and loaded from presets using the buttons below the graph legend. Each preset contains the room dimensions and absorption values, the source selections and positions and mic selections and positions. Presets are remembered for the next startup. Any comment entered is saved with the preset. To the right of the preset buttons are a button to set the reference trace to the current Main position response and a button to capture an image of the response graph, which includes any comment entered.
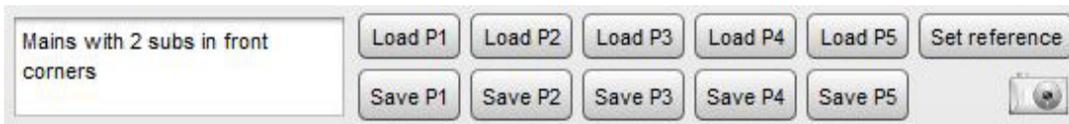
Controls

The **Modal Resonance Lines** controls identify the colours of the individual lines and allow their transparency to be adjusted. Any lines which are not selected will not appear on the graph.
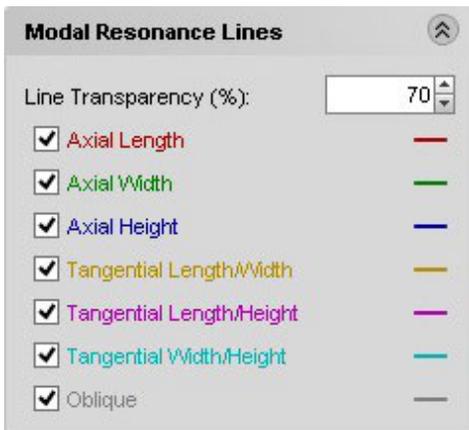
Colour	Mode
Red	Axial Length
Green	Axial Width
Blue	Axial Height
Orange	Tangential Length, Width
Magenta	Tangential Length, Height
Cyan	Tangential Width, Height
Grey	Oblique

The **Microphone Positions** controls set the distances for the additional listening positions from the main position. They can also be adjusted by dragging the crosses on the room view.
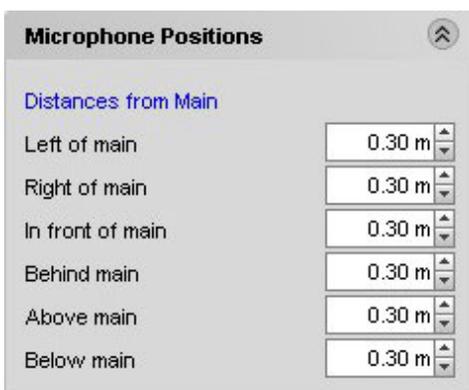
The **Speaker Controls** allow a number of sources to be selected, including up to 8 subwoofers. If **Subs have same delay** is selected all the subs use the delay figure next to the check box, otherwise each sub uses its

own delay setting. The low frequency extension of each source can be configured independently - this is the frequency at which the source begins to roll off, it is not the bass management frequency, which is set using the **Crossover Filter** control. The room responses shown are the sums of the contributions of all the selected sources.
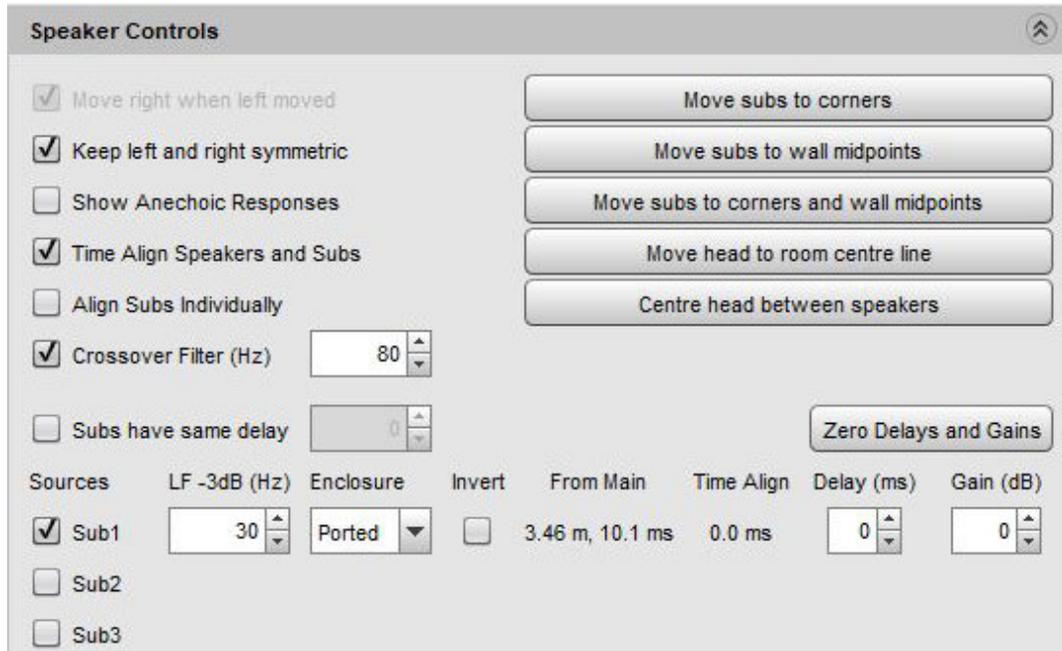
Subwoofers can be relocated to the corners or wall midpoints by pressing the appropriate buttons, the head can be moved to be centred across the width of the room or centred between the speakers.

If **Move right when left moved** is selected moving the left speaker will also move the right speaker, keeping the same distance between the speakers. The right speaker can be moved independently. If **Keep left and right symmetric** is selected moving either speaker will move the other, keeping them at the same height and distance from the front wall and equal distances from the side walls.

If **Show Anechoic Responses** is selected the response plots remove the room contributions, leaving only the responses of the sources themselves. Sources can be time aligned at the main listening position and, if multiple subs are being simulated, the effect of time aligning each sub individually can be observed (note that if a symmetric placement of multiple subs is being used to minimise modal excitation the subs should **not** be individually aligned). The distances and times of flight to each source are shown and the effect of adding a time delay or gain adjustment to each source can be observed by altering the relevant controls.

Note that the simulation automatically level aligns sources to the main listening position, but that all the simulated subwoofers generate signals at the same level - this is necessary if they are arranged symmetrically to minimise modal excitation and reduce seat-to-seat variation.

Generated measurements

On the right side of the panel there are buttons to generate measurements from the simulated responses.
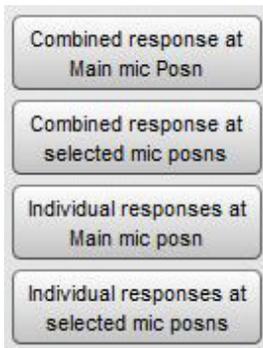
The **Combined response at Main mic posn** button produces a measurement showing the simulated response at the Main mic position when all the currently selected sources are active. **Combined response at selected mic posns** produces a measurement showing the average of the simulated responses at the currently selected mic positions when all the currently selected sources are active, with an underlay showing the overall span of the individual responses.

Individual responses at Main mic posn produces a set of measurements, one for each of the currently selected sources, showing the simulated response at the Main mic position from that source. **Individual responses at selected mic posns** produces a set of measurements, one for each of the currently selected sources, showing the average of the simulated responses at the currently selected mic positions when that source is active, with an underlay showing the overall span of the individual responses.

The measurements generated include the room dimensions and speaker and mic positions, a plan view of the room configuration used to produce the measurement can be shown on the SPL & Phase graph by selecting **Show room panel** in the graph controls. The view has the same highlighting features as the room simulator view, except positions cannot be altered.
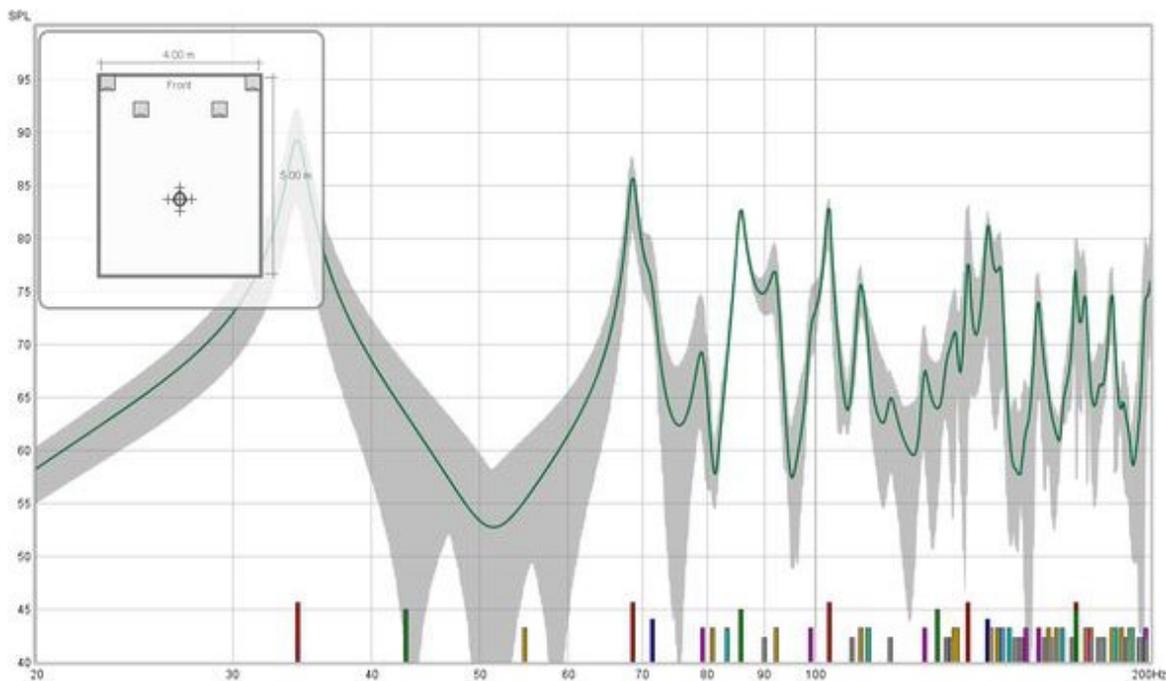
Help Index

Importing Measurement Data

REW can import frequency response and impedance measurement data from other applications, including the ETF5 measurement system (http://www.acoustisoft.com). Most ETF5 export formats are supported, along with generic comma, space or TAB-delimited text files. Note that when making measurements with ETF5 it is best to use the calibrated SPL option to allow level comparisons between channels.

ETF *.pcm Format Impulse Responses

Full range ETF measurements allow the impulse response to be exported in a .pcm raw data format, using the *File → Write Impulse As *.pcm* option. Use REW's **File → Import Impulse Response** command to load these files. Importing the impulse response allows REW's Spectral Decay and Waterfall plots to be generated from the data (this is not possible when importing frequency response text files). REW loads the first 128k samples from the file (approx 2.73s at 48k sampling).

WAV or AIFF Format Impulse Responses

The **File → Import Impulse Response** command can also be used to load impulse responses that have been saved as .wav or .aiff format. REW loads the first 256k samples from the file (approx 5.46s at 48k sampling), if the file contents are shorter than 1 second the response is padded out to 1 second with zeroes.

Frequency Response and Impedance Text Files

The **File → Import Frequency Response** and **File → Import Impedance Measurement** commands accept text files with extension .txt, .frd, .dat or .zma. If the extension is .zma the data is always treated as impedance, otherwise it is treated as SPL if loaded via **Import Frequency Response** or impedance if loaded via **Import Impedance Measurement**. The following formats are accepted:

- Generic comma, TAB, space or semicolon delimited files
  - Data must be presented as freq, magnitude (SPL or impedance) and (optionally) phase in degrees, one set of values per line
  - Samples can be at arbitrary frequency spacing, but each line must have a higher frequency than the one before and there must be at least 5 data entries
  - Only lines which begin with a number are imported, others are ignored
  - In comma-delimited files there must be at least one space after the comma
  - Spaces before values are ignored
  - To add a comment on a data line put a tab, comma or space after the last value

- ETF5 Export Data option for the Low Frequency Room Response window
- ETF5 Export Bode Response option
- ETF5 Export Data option for the Logarithmic Frequency Response window
  - Cubic spline interpolation is used between sample points

Comma-delimited File Format

Here is an example of a valid format for comma-delimited SPL data without phase:

```
SPL measurements acquired by REW V3.08
```

```
            Source: D:\REW\test files\testfile.txt
            Format: Comma delimited data
            Dated: 05-Mar-2005 17:53:56
            Channel: Left, Bass limited 80Hz


            20.0, 65.01
            21.0, 65.77
            22.0, 67.50
            23.0, 67.93
            24.0, 68.22
            25.0, 67.88
            26.0, 67.92
            27.0, 68.31, this line has a comment
            28.0, 69.14
            29.0, 69.16
            30.0, 69.29
```

If comma is used as the decimal delimiter in your locale it is best for clarity to use TAB, space or semicolon as the separator.

Space-delimited File Format

Here is some impedance data in space-delimited format, with phase

```
            * Measurement data saved by REW V5.00
            * Source: Line (ESI MAYA44 Audio), no input selected, Right channel,
volume: no control
            * Format:    1M Log Swept Sine, 1 sweep at -30,0 dBFS
            * Dated: 31-Dec-2010 11:26:49
            * Sense Resistor: 100.0
            * Lead resistance: 0.000
            * Calibration factor: 1.0028
            * Note: horizontal
            * Measurement: SPH170 horz
            * Frequency Step: 0.36621094 Hz
            * Start Frequency: 1.8310547 Hz
            *
            * Freq(Hz) Z(Ohms) Phase(degrees)
            1.831 6.423 5.392
            2.197 6.444 6.426
            2.563 6.481 7.302
            2.930 6.522 8.049
            3.296 6.558 8.714
            3.662 6.586 9.368
            4.028 6.609 10.076
            4.395 6.632 10.864
            4.761 6.664 11.706
            5.127 6.705 12.549
            5.493 6.753 13.346
            5.859 6.803 14.082
            6.226 6.851 14.777
            6.592 6.896 15.467
            6.958 6.939 16.180
            7.324 6.985 16.919
            7.690 7.034 17.669
            8.057 7.089 18.409
            8.423 7.147 19.126
            8.789 7.207 19.822
            9.155 7.269 20.505
            9.521 7.332 21.181
            9.888 7.398 21.851
```

Communicating with AV32R DP or AV192R

AV32R Dual Processor and AV192R support serial communication via a TAGtronic Programming Cable. The serial format is 115.2kBaud, 8 data bits, no parity, 1 stop bit. The Programming Cable converts RS232 levels to the RS485 levels of the TAGtronic bus. Only connect the black plug (to the "out" socket) for communication with REW.

Owners of AV192R with the front panel inputs option can connect to the unit's front panel programming connector using the RS232-jack plug lead.

The serial interface allows REW to read loudspeaker configuration and filter settings from the unit and to send filter settings to the unit.

REW only supports serial communication on Windows platforms. The COM port is set via the Comms panel in the Preferences dialog.

Communicating with a TAG McLaren AV Processor

Connect to the AV processor using the programming lead. If using the rear panel connections, plug the black lead into the "out" socket. Do not connect the red lead.

Backup TAG McLaren AV Processor Settings

Before proceeding further it is advisable to back up the settings in your AV32R DP or AV192R using the TMA User Settings Backup utility.

Read and Save Current Filter Settings

Use the Retrieve Channel Filter Settings from Unit (Ctrl+F) entry in the Equaliser menu to retrieve the current TMREQ filter settings for a channel from the processor then save them as a .req file

using the <u>Save Filters</u> entry in the File menu.

Selecting the Test Signal Input on AV32R DP or AV192R

Select the AV processor input to which the soundcard's output has been connected then go to the TMREQ menu. If the channel you wish to measure already has filters defined, they can be disabled temporarily by pressing the **eject** key on the remote control (the **record** key is used to re-enable them).

Go into one of the filter menus for the channel and set the `Test Signal` to `Current R`. Set `Repeat Sig.` to `No` (this is only used when you want to measure the effect of running two speakers at the same time). Set `Bass Redir.` to `No`, which prevents the subwoofer being activated when measuring a Bass Limited speaker - the signal which would normally be redirected to the subwooofer is discarded. After making corrections, re-measuring with `Bass Redir.` set to `Yes` will allow the integration between the speaker and the subwoofer to be checked.

<u>Help Index</u>

Communicating with the BFD Pro DSP1124P and FBQ2496

The BFD Pro DSP1124P and FBQ2496 models support Midi communication for setting up their filters and operating modes. To communicate with the units over Midi a Midi Interface is required, for example the Edirol UM-1X USB-Midi interface or the M-Audio UNO.

Connect the plug labelled "OUT" on the Midi interface to the socket labelled "IN" on the BFD (on the UM-1X the adaptor OUT plug has the text "Connect to Midi IN" moulded into it). It is not necessary to connect to the Midi OUT of the BFD.

The Midi Output port is selected via the Comms panel in the Preferences dialog, REW will not be able to communicate over Midi until the port has been selected. The selection is remembered for the next power-up.

Midi communication is supported on Windows platforms. Linux platforms will require Tritonus (www.tritonus.org) to support Midi comms. macOS platforms with JRE V6 or later installed should support Midi.

Configuring DSP1124P for Midi Communications

In the DSP1124P default configuration Midi comms are disabled. To enable the Midi features used by REW, the unit's Midi menus need to be set up as follows (all buttons referred to are on the DSP1124P front panel):

1. Press the IN/OUT and STORE buttons together to access the Midi menus, the LEDs in both buttons start flashing and the display changes to show:

2. This is the Midi channel menu (indicated by the "c" in the right hand digit), when the channel shows "-" midi is off. Use the jog wheel to change the channel to 1:

3. Press the IN/OUT button twice to change to the Controller menu ("C" in right hand digit) and use the jog wheel to select mode 3:

4. Press the IN/OUT button again to change to the Program menu ("P" in right hand digit) and use the jog wheel to select mode 3:

5. Press the IN/OUT button again to change to the Store Enable menu ("S" in right hand digit) and use the jog wheel to set the value to 1:

     This enables REW to save settings to the DSP1124P's presets

6. Press the IN/OUT button 2 more times to exit the Midi menus

The DSP1124P is now configured for Midi communications. This configuration, with the exception of the Store Enable setting, is remembered for the next power-up and does not need to be entered again.

Configuring FBQ2496 for Midi Communications

Midi comms are enabled by default in the FBQ2496. If Midi comms has been turned off, or the channel has been set to something other than 1, proceed as follows (all buttons referred to are on the front panel):

1. Make sure the unit is NOT in PEQ mode (i.e. the LED in the PEQ button must be off, if it is on press the PEQ button to turn it off). Press the BANDWIDTH and BYPASS buttons together to access the Midi menus, the LEDs in both buttons start flashing as does the MIDI LED below the numeric display. The display itself shows the Midi on/off status, if it shows **OFF** turn the knob until it changes to **on**. Then press the BANDWIDTH button to change to the Midi channel menu, shown in the display by **C** followed by the channel number. Turn the knob to select channel 1.
2. Press any button except BANDWIDTH or BYPASS to exit the Midi menus.

The FBQ2496 is now configured for Midi communications. This configuration is remembered for the next power-up and does not need to be entered again.

Notes

1. Store Enable is turned off by the DSP1124P when it powers up, REW will prompt you to turn on Store Enable for each measurement session when using DSP1124P. If you do not turn on Store Enable REW will not be able to save filter settings to presets - after downloading filters to the DSP1124P the red LED in the STORE button will be flashing as a warning that changes have been made but not stored. You can manually save to presets by pressing the STORE button, using the jog wheel to select the preset to store to, then pressing the STORE button again (just press the button twice if you are already on the preset you want to use).

2. The IN/OUT button LED on DSP1124P flickers during Midi communications, on FBQ2496 the MIDI LED flickers.

3. When filters are downloaded to FBQ2496 REW will configure the unit to have 20 parametric filters on the channel being downloaded and, after the download, will turn off the bypass (if it is on).

4. Downloading a set of filters takes about 1 second per filter.

[Help Index](#)

Soundcard Preferences

The Soundcard Preferences panel is used to configure the audio input and output used for measurement, calibrate the audio interface and establish the correct levels for making measurements.
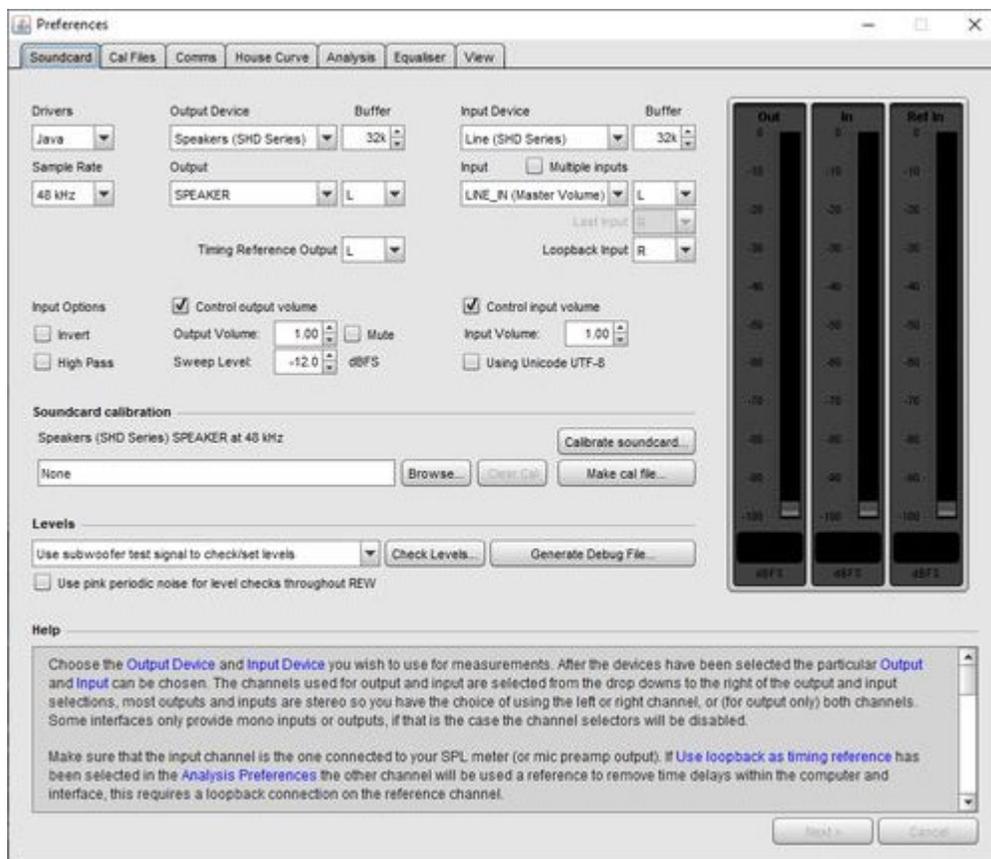
The various controls on the panel are as follows:

Drivers

On Windows platforms there is a choice of Java or ASIO drivers for the audio interface. The Java drivers support 44.1, 48, 88.2, 96 and 192 kHz sample rates and 16-bit data. On macOS and Linux 32-bit or 24-bit data is used if the interface offers it. Java drivers permit the input and output to be on different devices and allow volume control from REW.

The ASIO drivers support up to 1536 kHz and a variety of formats depending on the driver. ASIO drivers support one ASIO device which must be used for both input and output and REW has no control over levels. Pseudo-ASIO drivers such as ASIO4All create an ASIO wrapper around the WDM drivers for devices, allowing input and output through different devices.

Sample Rate

With Java drivers the sample rate may be set to 44.1, 48, 88.2, 96 or 192 kHz, the default is 48kHz. To prevent resampling in the OS make sure the audio interface is configured to operate at the sample rate selected in REW. With ASIO drivers the choice of sample rates offered will reflect those the interface supports, with a maximum sample rate of 1536 kHz. Note that the lists of input and output devices only include those devices that report they support the selected sample rate, if your device

does not appear in the lists try changing the sample rate.

Inputs and Outputs (Java drivers)

The input and output device lists show the physical devices that Java has found that report they support the selected sample rate, along with some OS virtual devices. The lists of inputs and outputs are specific to the selected input device and output device. The **Default Device** settings tell REW to request the defaults that have been set in your OS (in the Sounds and Audio Devices control under Windows or the Audio and Midi Setup utility under macOS). When the default devices have been selected REW leaves all control of the audio inputs and outputs and their associated volume controls to you, use the controls provided by your interface's mixer or OS controls to set levels and select inputs and outputs as required. **Note** that when using a USB mic with a cal file that contains a sensitivity figure REW needs to read the input volume setting to correctly show SPL, to allow that the input device and input for the mic must be selected (they must not be left as "Default Device").

Output Channel (Java drivers)

REW can place its test signals on either or both of the output channels. If there are more than 2 available output channels the first and second can be driven at the same time - which channels those are can be configured using the **Output Channel Mapping** dialog (see below). The output channel selection can also be made directly on the measurement panel or signal generator.

Timing Reference Output Channel (Java drivers)

When using a timing reference the output channel will be the channel selected here. The selection can also be made directly on the measurement panel.

Output Mapping (Java drivers)

On some platforms Java supports multichannel output - currently Windows does not offer this. If the platform supports multichannel and the current output has more than 2 channels an output channel mapping button is shown:
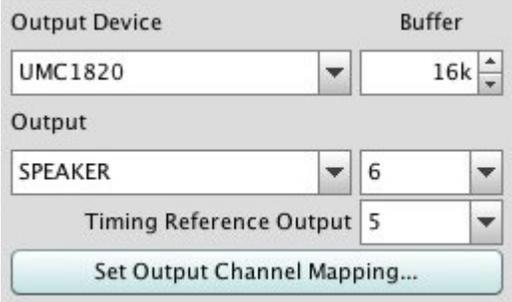
The button brings up a dialog to select up to 8 channels for use while measuring:

to specific dB increments which vary according to the device and its volume range. Use Audio MIDI Setup to control volume settings on macOS.

Sweep Level

The Sweep Level control sets the RMS level at which REW will generate its measurement sweep, relative to digital full scale. The highest level possible is -3 dBFS, unless the View preference Full scale sine rms is 0 dBFS has been selected, in which case the maximum is 0 dBFS. Using the maximum value places the peaks of the signal at digital full scale. A typical setting is -12 dBFS (the default). This selection can also be made directly on the measurement panel.

Output Buffer, Input Buffer (Java drivers)

The Output Buffer and Input Buffer controls set the size of the buffers used when accessing the interface. The default settings are 32k (meaning the buffer sizes are 32,768 pairs of audio samples). If you experience occasional glitches or interruptions in the signal generator output try increasing the replay buffer size, but note that there are other possible causes of this, such as interference from wireless cards. Similarly if the captured audio signals (as shown in the Scope graph panel) have occasional dropouts try increasing the record buffer. Using larger buffers will increase latency (delays when starting and stopping replay and recording) but should otherwise not be detrimental. If you are not experiencing any problems with audio input or output you may wish to reduce the buffer sizes to minimise latency.

Calibration Panel

The controls in the Calibration panel are used to calibrate the interface.

The **Browse...** button is used to select a calibration file, a plain text file which by default has the extension .cal, though other extensions are also accepted. The file format is detailed below. **Clear Cal** clears the calibration data structures, all subsequent measurements will not have any interface calibration corrections applied to them and REW will not load any previously specified interface calibration file on the next startup. **Calibrate...** starts a process of measuring the interface response via an external loopback connection. **Make cal file...** is used to save a measurement as a calibration file - this should only be used with the results of a loopback measurement, and then only after checking that the measurement is valid. The measurement data is saved as a text file, with the SPL values offset to give 0dB at 1kHz. The file is automatically loaded on startup and applied to subsequent measurements.

Levels Panel

The controls in the Levels panel are used to set the output and input levels for measurement. Levels can be set using either a subwoofer or one of the main speakers, this is selected in the drop-down box in the panel. The **Check Levels...** button starts a process of establishing and verifying the levels. The **Generate Debug File...** button generates a text file with information about all the audio devices and controls that Java has been able to identify. If there are problems configuring the interface for use with REW provide a copy of this file along with a description of the problem. **Use pink periodic noise for level checks throughout REW** controls whether REW uses random pink noise or periodic pink noise when checking levels. Periodic noise provides much more stable readings at low frequencies but sounds quite different to random noise.

Example Input and Output Settings

Here are some example settings, firstly using Java drivers, a PC's built-in audio interface and a USB mic. REW has been set to control the levels and the Right channel is being used for input.

Here are some settings using ASIO drivers, in this case ASIO4All (which provides an ASIO wrapper around Windows audio drivers). Note that it is not necessary to select a timing reference output if a timing reference is not being used. The **ASIO Control Panel** button launches the ASIO control panel for the interface.

FlexASIO

If the FlexASIO ASIO driver is being used the **ASIO Control Panel** button will show a dialog that provides a graphical interface to generate and update the FlexASIO configuration file, FlexASIO.toml.

REW's handling of the configuration file only supports the **backend** and **bufferSizeSamples** options. The **[input]** and **[output]** sections will only contain the selected input and output device names, they will not contain any other options that FlexASIO supports. The configuration file is updated when the **Close panel and update config file** button is pressed, the FlexASIO driver is then reloaded to update the available inputs and outputs. If the dialog is closed without using the **Close panel and update config file** button any changes will be discarded.

Soundcard Calibration File Format

The calibration file is a plain text file which by default has the extension .cal, though other extensions are also accepted. It should contain the actual gain (and optionally phase) response of the interface at the frequencies given, these will then be subtracted from subsequent measurements. The values in the calibration file can be separated by spaces, tabs or commas.

- Each line of calibration data must have a frequency value and a gain value, a phase value is optional
- Frequency is in Hz, gain in dB, phase in degrees
- The cal points can be at arbitrary frequency spacing, but each line must have a higher frequency than the one before and there must be at least 2 freq, gain data pairs
- Only lines which begin with a number are loaded, others are ignored
- In comma-delimited files there must be at least one space after the comma
- Spaces before values are ignored
- The sample rate at which the data was generated can be indicated by having a line which starts "Sample Rate:" (without the quotation marks) followed by the sample rate in Hz. REW checks for this when loading a file and will warn if the rate does not match the current interface setting - calibration data generated at a different sample rate will not provide accurate correction.

Here is an example section of a valid file format:

```
* Soundcard Calibration data saved by Room EQ Wizard V5.00
* Source: EDIROL UA-1A, Digital Audio Interface, Right channel, volume: no
control
* Format:  256k Log Swept Sine, 1 sweep
* Dated: 21-Nov-2010 21:47:56
* Sample Rate: 44100
*
2.019 -1.424 53.471
2.219 -1.238 49.420
2.419 -1.062 45.118
2.619 -0.929 41.888
2.819 -0.823 39.056
3.019 -0.740 36.590
3.219 -0.668 34.409
3.419 -0.607 32.468
3.619 -0.557 30.736
3.819 -0.513 29.177
4.019 -0.475 27.777
4.219 -0.443 26.495
```

After a calibration file has been loaded it will be applied to all subsequent measurements. Loading the calibration file does NOT affect any data already measured and does not affect any measurement data that is imported. The graph display is updated to show the calibration curve, offset to lie at the current Target level.

Linear interpolation is used between calibration points. Outside the range of the calibration data the behaviour depends on whether C weighting compensation has been selected. If C weighting compensation

is selected, C weighting curve figures will be used for frequencies above or below the range of frequencies in the calibration data. If not, the calibration values for the lowest frequency in the file will also be applied for all lower frequencies and the calibration values for the highest frequency in the file will be applied for all higher frequencies.

The calibration file name and path are remembered for the next startup, the file will be loaded automatically when REW is started. A message confirming loading of the file is given.

To stop calibration data being applied, use the **Clear Cal...** button. **Useful tip:** To apply or remove a soundcard calibration file after a measurement has been taken, simply load or clear the cal data as required and press the **Apply Windows** button in the IR Windows panel to recalculate the frequency response.

[Help Index](#)

Cal files Preferences

The Cal files Preferences list all the outputs and inputs that have previously been selected in REW with any calibration files that have been specified for them. Choosing an output device and output or input device and input will cause REW to automatically load the associated calibration file. Entries can be deleted if they are no longer needed. Individual calibration files can be added using the **Browse** buttons or removed using the **Clear Cal** buttons. The calibration data will be applied to all new measurements taken after it has been loaded and will be shown on the SPL & Phase graph for the measurements. The currently active entries are highlighted. The calibration files that will be used for a new measurement can be seen by using the **Cal files** button on the Measure dialog.
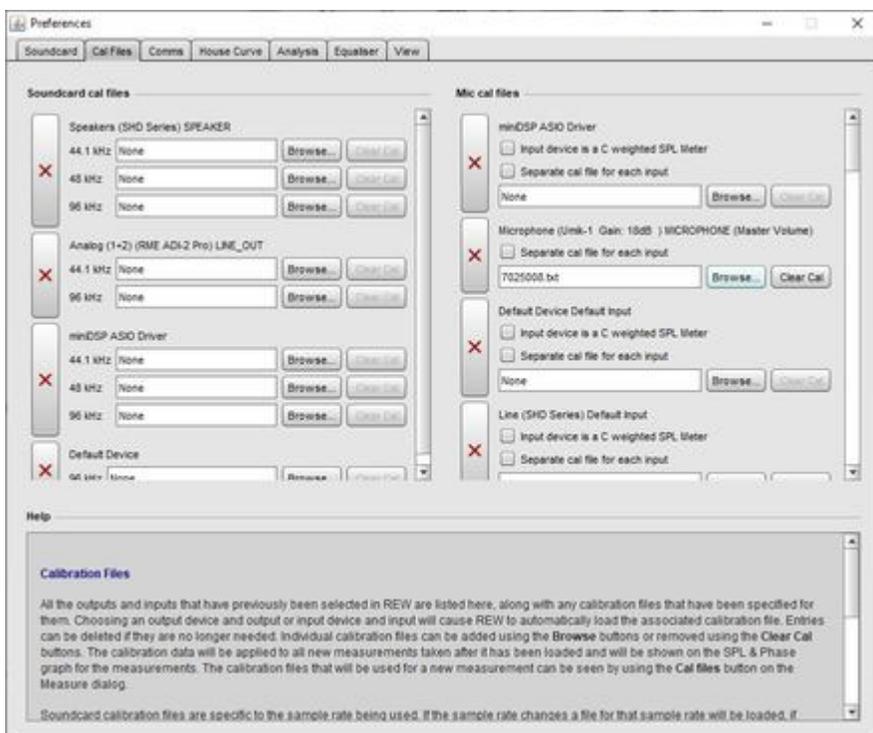
Soundcard calibration files are specific to the sample rate being used. If the sample rate changes a file for that sample rate will be loaded, if available.

Mic calibration has a C weighting option. Select Input device is a C Weighted SPL Meter if you are using a C weighted SPL meter as the input to REW, measurements will then be corrected to remove the low and high frequency roll-offs of the C weighting characteristic. If a cal file is also loaded the correction will only be applied outside the frequency range covered by the cal file.

Mic cal files can apply to all inputs on a device or, if Separate cal file for each input is selected individual cal files can be specified for the separate inputs of the device. For example, the left and right inputs of a miniDSP EARS headphone measurement system have separate cal files. Note that most USB microphones are mono devices and provide the same data on both left and right channels, a single cal file for the device is appropriate for those.

To apply or remove a calibration file for an existing measurement, use the **Change Cal...** button in the measurement panel.

To calibrate the REW SPL meter reading against an external meter use the **Calibrate** button in the SPL Meter controls.

Calibration File Format

The calibration file is a plain text file which by default has the extension .cal, though other extensions are also accepted. It should contain the actual gain (and optionally phase) response of the meter or microphone at the frequencies given, these will then be subtracted from subsequent measurements. The values in the calibration file can be separated by spaces, tabs or commas. Typically the values are relative to the level at some reference frequency, e.g. 1kHz, so the gain value there is 0.0.

- Each line of calibration data must have a frequency value and a gain value, a phase value is optional
- Frequency is in Hz, gain in dB, phase in degrees
- The cal points can be at arbitrary frequency spacing, but each line must have a higher frequency than the one before and there must be at least 2 freq, gain data pairs
- Only lines which begin with a number are loaded, others are ignored
- In comma-delimited files there must be at least one space after the comma
- Spaces before values are

Here is an example section of a valid file format:

```
SPL Meter Calibration data

20 -15.38
50 -3.69
100 -1.34
200 -0.62
500 -0.26
1000 0.0
2000 1.80
5000 3.95
10000 -0.71
20000 -6.28
```

Mic sensitivity calibration data

USB microphone cal files will typically have a line at the start of the file that defines their sensitivity, of the form:

```
"Sens Factor =1.2345dB, SERNO: XXXXXXX"
```

The sens factor figure is usually the input dBFS reading the mic will produce when driven by a 94 dB calibrator with the input volume set to maximum.

Analog microphone cal files can also include sensitivity information. Add a line at the start of the file of the form:

```
Sensitivity -12.34 dBFS
```

where "-12.34" is replaced by the REW dBFS input level when the microphone is driven by 94 dB SPL. When that cal file is used REW will use the sensitivity data to calculate the SPL calibration offset. Note, however, that this calibration will **only be valid when using the same mic interface gain and input volume that were used when it was measured**.
Help Index

Comms Preferences

The Comms Preferences panel is used to choose the Midi and RS232 interfaces for communications with an equaliser.
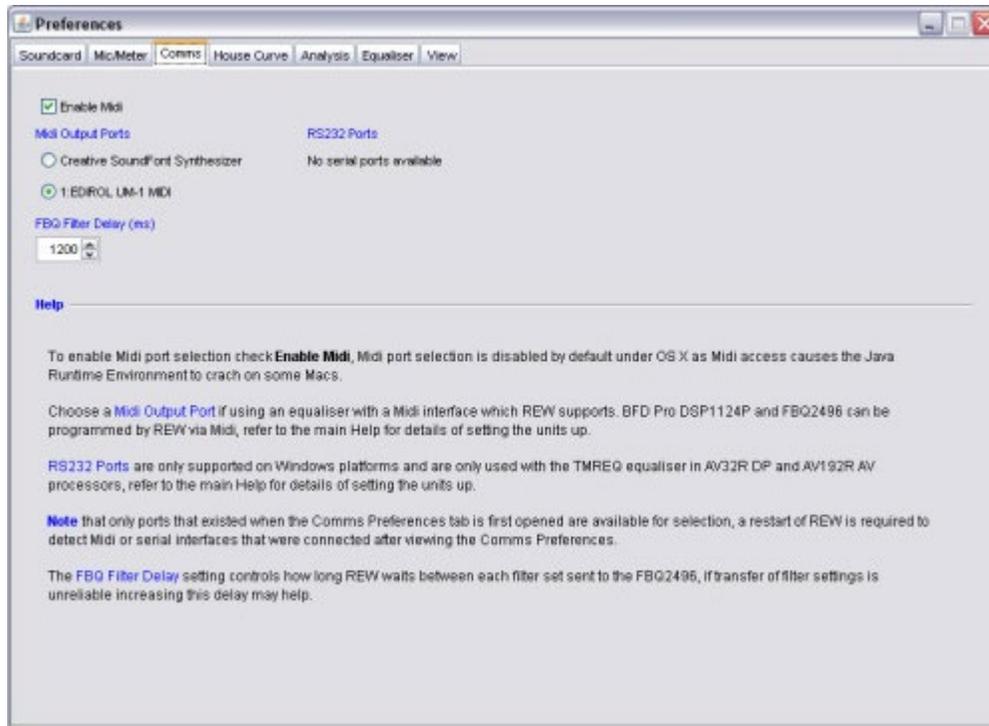
To enable Midi port selection check **Enable Midi**, Midi port selection is disabled by default under macOS as Midi access causes the Java Runtime Environment to crash on some Macs.

Choose a Midi Output Port if using an equaliser with a Midi interface which REW supports. BFD Pro DSP1124P and FBQ2496 can be programmed by REW via Midi, refer to the BFD Comms help for details of setting the units up.

RS232 Ports are only supported on Windows and are only used with the TMREQ equaliser in AV32R DP and AV192R AV processors. Refer to the AVP Comms help for details of setting those units up.

Note that only ports that existed when the Comms Preferences tab is first opened are available for selection, a restart of REW is required to detect Midi or serial interfaces that were connected after viewing the Comms Preferences.

The FBQ Filter Delay setting controls how long REW waits between each filter set sent to the FBQ2496, if transfer of filter settings is unreliable increasing this delay may help.

Help Index

House Curve Preferences

A "house curve" is a target response which differs from that defined by the bass management settings of the speakers. The House Curve Preferences allow a file containing house curve data to be loaded, or a curve that has been loaded to be removed. The selection is remembered for the next startup.
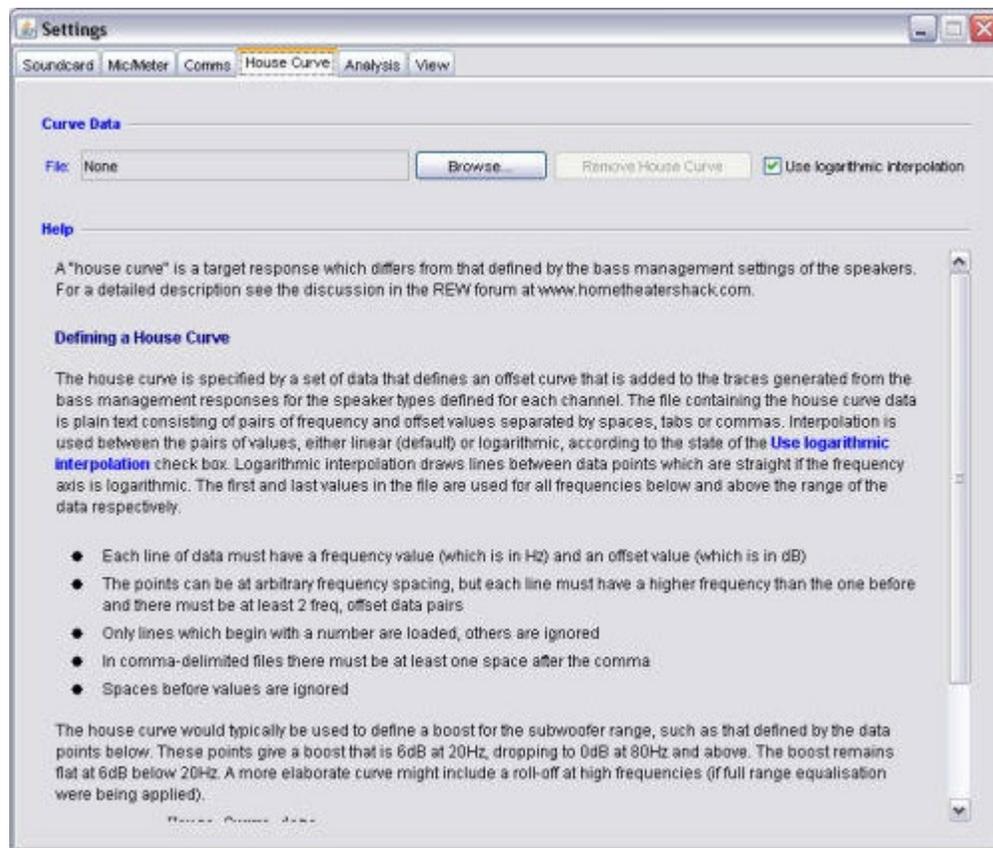
Defining a House Curve

The house curve is specified by a set of data that defines an offset curve that is added to the traces generated from the bass management responses for the speaker types defined for each channel. The file containing the house curve data is plain text consisting of pairs of frequency and offset values separated by spaces, tabs or commas. Interpolation is used between the pairs of values, either linear (default) or logarithmic, according to the state of the Use logarithmic interpolation check box. Logarithmic interpolation draws lines between data points which are straight if the frequency axis is logarithmic. The first and last values in the file are used for all frequencies below and above the range of the data respectively.

- Each line of data must have a frequency value (which is in Hz) and an offset value (which is in dB)
- The points can be at arbitrary frequency spacing, but each line must have a higher frequency than the one before and there must be at least 2 freq, offset data pairs
- Only lines which begin with a number are loaded, others are ignored
- In comma-delimited files there must be at least one space after the comma
- Spaces before values are ignored

The house curve would typically be used to define a boost for the subwoofer range, such as that defined by the data points below. These points give a boost that is 6dB at 20Hz, dropping to 0dB at 80Hz and above. The boost remains flat at 6dB below 20Hz. A more elaborate curve might include a roll-off at high frequencies (if full range equalisation were being applied).

```
House Curve data

20 6.0
80 0.0
```

When a house curve has been loaded the ᴴᴼᵁˢᴱ symbol is displayed next to the Target trace value in the Filter Adjust graph.

[Help Index](#)

Analysis Preferences

The Analysis Preferences alter the way REW carries out some of its calculations.
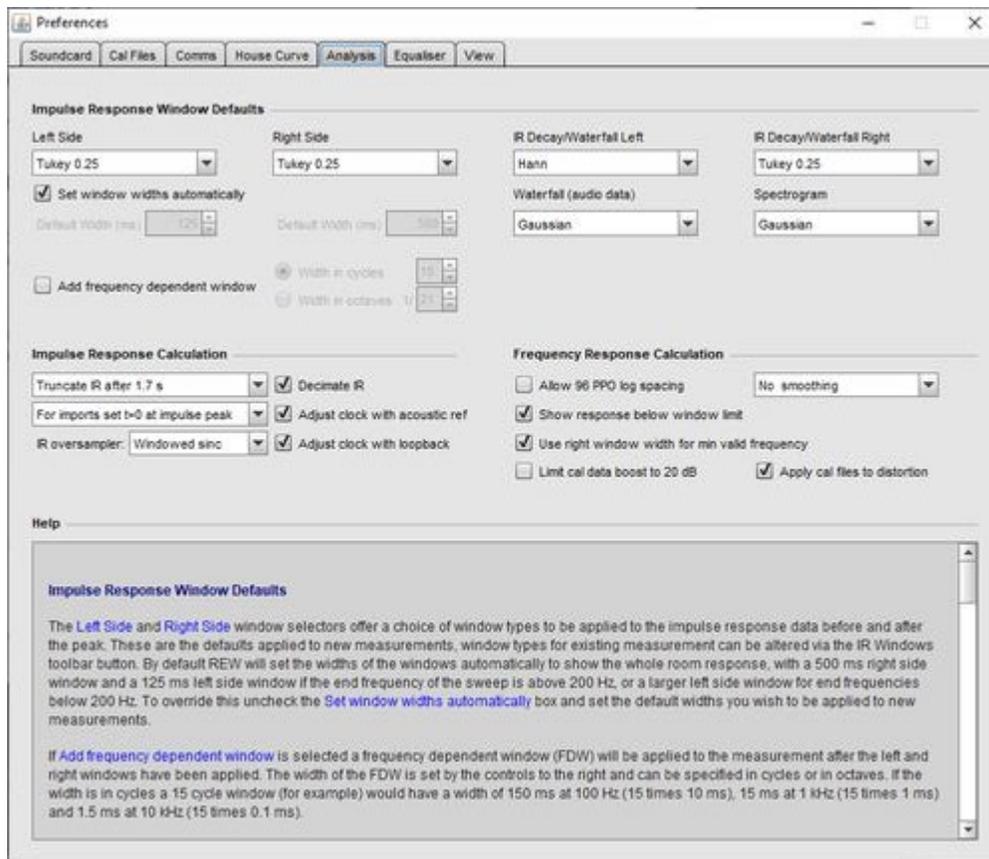
Impulse Response Window Defaults

The Left Side and Right Side window selectors offer a choice of window types to be applied to the impulse response data before and after the peak. These are the defaults applied to new measurements, window types for existing measurement can be altered via the IR Windows toolbar button. By default REW will set the widths of the windows automatically to show the whole room response, with a 500 ms right side window and a 125 ms left side window if the end frequency of the sweep is above 200 Hz, or a larger left side window for end frequencies below 200 Hz. To override this uncheck the Set window widths automatically box and set the default widths you wish to be applied to new measurements.

If Add frequency dependent window is selected a frequency dependent window (FDW) will be applied to the measurement after the left and right windows have been applied. The width of the FDW is set by the controls to the right and can be specified in cycles or in octaves. If the width is in cycles a 15 cycle window (for example) would have a width of 150 ms at 100 Hz (15 times 10 ms), 15 ms at 1 kHz (15 times 1 ms) and 1.5 ms at 10 kHz (15 times 0.1 ms).

The IR Decay/Waterfall Left and IR Decay/Waterfall Right window selections are applied to the impulse response data when generating Spectral Decay and Waterfall plots.

The Waterfall (audio data) window selection is applied when generating waterfall plots for imported audio data. This is the default, the window type used for any particular measurement can be altered via the selector in the graph controls.

The Spectrogram window selection is applied to the response data when generating Fourier Spectrogram plots. This is the default, the window type used for any particular measurement can be altered via the selector in the graph controls.

Impulse Response Calculation

After REW has made a measurement it can truncate the derived impulse response to preserve the important information while minimising the storage required for the measurement file. A 1 second period is retained before the peak, and by default a 1.7 second period is retained after the peak (this varies a little depending on the sample rate, at 44.1k (or multiples) it is approx 2 seconds, at 48k 1.7 seconds). In some cases it may be useful to retain more of the impulse response, such as measurements in very large spaces which have very long impulse responses. REW provides options to truncate the response after approx 4.4 seconds, or 9.9 seconds, or to retain the entire impulse response. It may also be necessary to use a longer sweep (with a 256k sweep there is about 6 seconds of impulse response data available after the peak). Note that retaining the entire impulse response will produce much larger measurement files, especially if long measurement sweeps are used. If the entire response is retained the peak will be centred within the response.

When impulse responses are imported the t=0 position can be set to either the first sample in the imported data or the location of the peak of the impulse response.

When REW generates an impulse response it is oversampled to provide data for graph rendering. The method used to generate that oversampled data is set by the IR oversampler setting. **Windowed sinc** uses a bandlimited interpolator, **FFT time shift** uses fractional delays implemented via the time shift property of the FFT, **None** disables IR oversampling (in which case Hermite interpolation is used in the graph plot generation between the original sample points). **Windowed sinc** is the default setting. This setting does not affect IR imports, which are not oversampled.

The Decimate IR selection controls whether REW reduces the sample rate of the impulse response to correspond to the range of frequencies in the measurement. Selecting this option greatly reduces the impulse response size for low frequency measurements and speeds up processing of the data.

The Adjust clock with acoustic ref selection controls whether REW compensates for clock rate differences between input and output when using the acoustic timing reference. If this option is selected an additional timing reference signal is played at at the end of the sweep and the time between the timing reference signals is used to calculate any clock adjustment required to match the input and output device clock rates.

The Adjust clock with loopback selection controls whether REW compensates for clock rate differences between input and output when using a loopback as a timing reference. If this option is selected the loopback signal is analysed to calculate any clock adjustment required to match the input and output device clock rates. To apply this correction the sweep must start below 1 kHz and end above 10 kHz.

If Loopback delay reference is IR peak is selected the system delay figure for loopback measurements is the relative times of the measured and loopback IR peaks. If it is not selected the system delay figure will use the estimated IR delay relative to the loopback when the peak is close to the IR start or the estimated IR start time relative to the loopback when the peak is later (limited bandwidth devices). The system delay figure for the acoustic timing reference always uses the estimated IR delay/estimated IR start.

Frequency Response Calculation

The Allow 96 PPO log spacing selection controls whether REW is permitted to convert frequency responses from linearly spaced data to logarithmically spaced data at 96 points per octave. The FFT that calculates the frequency responses produces data that is linearly spaced in frequency, i.e. there is a constant frequency step from each value to the next. For the high frequency parts of responses this means there are a very large number of points, using a lot of memory but not contributing anything useful to the displayed data. When this option is selected (it is on by default) REW will automatically convert frequency responses to more efficient logarithmic spacing with 96 data points in each octave of the response if this will reduce memory

usage (which is usually the case for sweeps that end above 300 Hz or so). As part of the conversion process REW first applies a 1/48th octave smoothing filter to the data to remove any high frequency combing from the response. The conversion takes place on any new measurement or when an IR Window is applied. Whether a measurement is log spaced or linearly spaced can be seen by bringing up the measurement info window by clicking the info button in the toolbar.

Note that conversion to 96 PPO is inhibited if the impulse peak is far from the impulse zero time, where "far" means the peak is offset from zero by a time that corresponds to more than 90 degrees of phase shift between samples at the measurement end frequency. This is to prevent aliased phase data at high frequencies which would lead to incorrect group delay figures.

The Show response below window limit selection controls whether REW displays the frequency and phase responses at frequencies lower than those valid for the current impulse response window width. For example, if the window width were 10 ms frequencies below 100 Hz would not be valid and are normally not displayed. There are circumstances in which it may be helpful to see that data, which this option allows, but the responses are drawn dashed to indicate they lie below the window cutoff. The Use right window width for min valid frequency selection controls whether the total or right hand window width is used for the minimum frequency calculation, the right hand window is the determining factor when the window reference time is at or near the IR peak, which is usually the case.

The default smoothing to apply to new SPL measurements can be selected from the drop-down box next to the log spacing check box.

If Limit cal data boost to 20 dB is selected REW will limit the total gain applied to compensate for calibration data attenuation to 20 dB. This prevents excessive boosting of the noise floor in areas where the combined mic/meter and soundcard response is more than 20 dB down. This setting affects the frequency response, RTA trace and SPL meter readings and is also applied when carrying out trace arithmetic.

If Apply cal files to distortion is selected REW will applying the cal files to distortion results. Applying the files provides more accurate results in regions where the fundamental or harmonics are affected by interface roll-offs but also boosts the noise floor in those regions.

Help Index

Equaliser Preferences

The Equaliser Preferences alter the way REW carries out its EQ filter calculations.
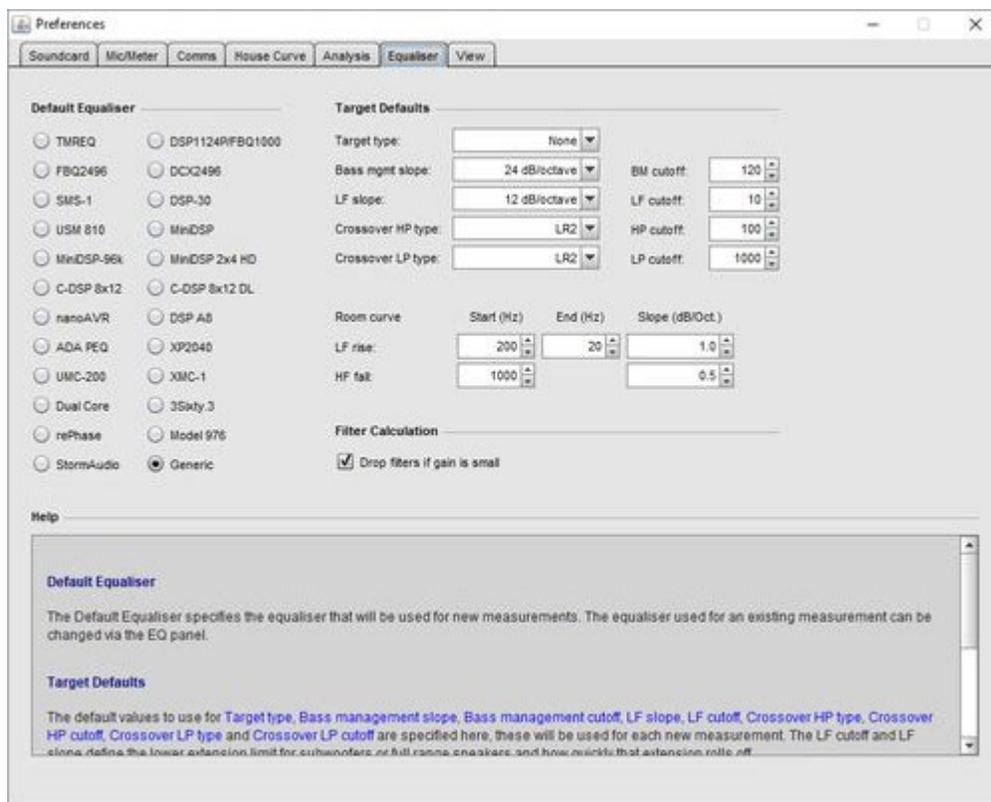
Default Equaliser

The Default Equaliser specifies the equaliser that will be used for new measurements. The equaliser used for an existing measurement can be changed via the EQ panel.

Target Defaults

The default values to use for Target type, Bass management slope, Bass management cutoff, LF slope, LF cutoff, Crossover HP type, Crossover HP cutoff, Crossover LP type and Crossover LP cutoff are specified here, these will be used for each new measurement. The LF cutoff and LF slope define the lower extension limit for subwoofers or full range speakers and how quickly that extension rolls off.

The Room curve LF rise and HF fall settings define the defaults that will be used for a room curve effect superimposed on the target shape for new measurements. The target curve will rise below the LF rise start frequency at the slope selected until reaching the LF rise end frequency. A rise at low frequencies may be subjectively preferred. Similarly, the target curve will fall above the HF fall start frequency at the slope selected. Falling HF response is a normal characteristic of in-room measurements at the listening position.

Filter Calculation

If Drop filters if gain is small is selected any Automatic filters which have gain of magnitude less than half the flatness target at the end of the optimisation process will be freed up (their Type will be set to "None")

View Preferences

The View Preferences control the appearance of REW
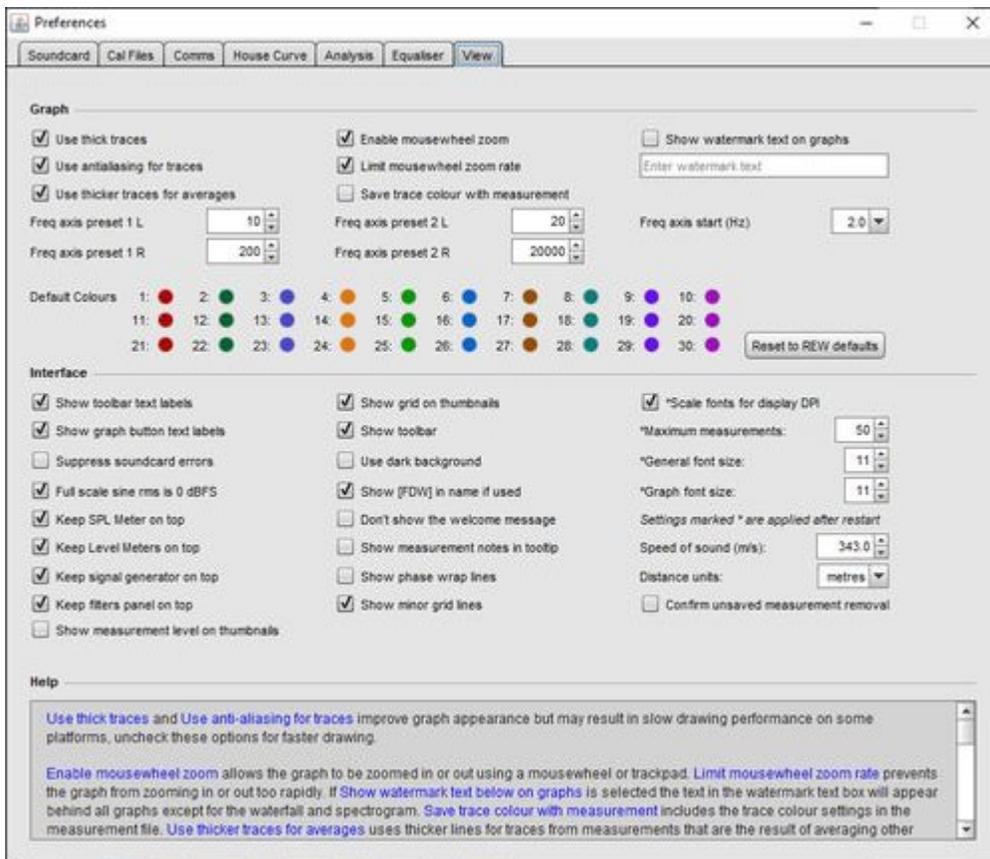
Use thick traces and Use anti-aliasing for traces improve graph appearance but may result in slow drawing performance on some platforms, uncheck these options for faster drawing. Use thicker traces for averages uses thicker lines for traces from measurements that are the result of averaging other measurements when plotting them on overlays or in the measurement thumbnail.

Enable mousewheel zoom allows the graph to be zoomed in or out using a mousewheel or trackpad. Limit mousewheel zoom rate prevents the graph from zooming in or out too rapidly. Save trace colour with measurement includes the trace colour settings in the measurement file. If Show watermark text below on graphs is selected the text in the watermark text box will appear behind all graphs except for the waterfall and spectrogram. Use dashed traces on overlays uses different line styles for overlay graphs to help distinguish the traces.

Freq axis preset 1 and 2 are buttons that appear on the graph to allow the frequency axis to be quickly changed between the two preset ranges defined here. Freq axis start (Hz) sets the lowest frequency that can be shown on graphs.

The default trace colours can be set up from here, or reset to the REW defaults. The r,g,b values for the trace colour are shown when hovering over the colour buttons. These default colours are remembered for the next time REW starts. Changing the default colours will not alter the trace colours of any measurements already loaded. To change the colours of existing measurements use the trace colour button on the measurement panel or click the **Set All Trace Colours To Defaults** button on the colour chooser panel.

Interface

The text under the toolbar buttons can be hidden by unchecking Show toolbar text labels. The text under the graph buttons can be hidden by unchecking Show graph button text labels. The grid on the measurement thumbnails can be turned off by unchecking Show grid on thumbnails.

Show toolbar controls whether the toolbar appears below the menu, hiding it makes more room for graphs on low resolution screens.

Maximum measurements is the maximum number of measurements that can be loaded in REW. Changes to this value do not take effect until REW is next started. Note that with a high maximum number of measurements REW could run out of memory, if that occurs it is possible to increase the amount of memory REW requests on startup during installation or by editing the roomeqwizard.vmoptions file in the REW program directory using the -Xmx option, REW uses -Xmx1024m by default. Admin privileges are required to change that file: copy it to a user-accessible directory, make the edits, then copy it back using Admin privileges.

Scale fonts for display DPI controls whether the font sizes are adjusted to allow for the screen resolution in dots-per-inch. Fonts will only be scaled if the resolution is more than 96 DPI. This option is not present on macOS (where scaling is automatic). It will have no effect on Windows unless the **High DPI scaling override** property is enabled with scaling set to **Application**. That property can be changed by right-clicking on the REW shortcut, selecting Properties, then the Compatibility tab, then clicking the **Change high DPI settings** button.

General font size is the base font size (in points) for REW. Graph font size is the font size used for the graph axes. Those font size changes do not take effect until REW is next started.

Max RTA inputs and Max level meters input set the maximum number of inputs that will be shown when using multi-input capture (Pro upgrade). Changes to these settings do not take effect until REW is next started.

Suppress soundcard errors stops REW showing error messages or warnings if it cannot access a soundcard on startup, this may be useful if REW is being used to analyse measurements on a computer without the audio interfaces needed for making measurements.

Use dark background changes the colour scheme from light to dark. Restart REW to update the colours for all components.

Full scale sine rms is 0 dBFS determines whether the rms level of a full scale sine wave is show as 0 dBFS or -3.01 dBFS. If this option is selected the rms level may be greater than 0 dBFS for some signal types, for example, a full scale square wave would have an rms level of +3.01 dBFS.

Show [FDW] in name if used appends [FDW] to the measurement name if a frequency dependent window has been applied.

Distance units selects whether distances are shown in metres or feet and inches.

If Confirm unsaved measurement removal is checked REW will ask for confirmation if you try to remove a measurement that has not been saved since it was last modified. This setting can also be changed from the removal confirmation dialog.

Speed of sound selects the value used for the speed of sound when converting times to distances, calculating modal resonances etc. The default value is 343.0 m/s (approximate speed of sound in dry air at 20 degrees C).

Use macOS file dialogs changes the file selection dialog to a format more like the native macOS file chooser, but note that the macOS chooser does not have the preview panel which REW uses to show the contents of mdat, req, calibration data and image files. There have been reports of the native file chooser sometimes not appearing, if that happens try switching back to the Java file chooser.

Don't show the welcome message controls whether REW shows a welcome message when it starts up.

Keep SPL Meter on top, Keep Level Meters on top and Keep signal generator on top control whether those windows stay on top of other windows.

If Show phase wrap lines is selected the wrap points of phase traces are shown with dashed lines.

If Show measurement level on thumbnails is selected the signal level used for a measurement will be shown on the thumbnail for the measurement.

If Animate measurement list is selected changes in the measurements list will be animated.

Using AZERTY keyboard is an option offered on macOS to work around a Java bug (JDK-8019498) causing incorrect shortcut key mappings.

Help Index

Keyboard Shortcuts

N.B. On macOS the shortcuts use *cmd* instead of *Ctrl* and *option* instead of *alt*

Grouped by function	
Keys	**Function**
F1	Show Help
F2	Show Help on Selected Item
Ctrl+M	Measure Response
Ctrl+N	Show Info Panel
Ctrl+S	Save Measurement
Ctrl+O	Open Measurement File
Ctrl+Shift+S	Save All Measurements
Ctrl+Alt+O	Open Filter file
Ctrl+Alt+S	Save Filters
Ctrl+Alt+R	Export RT60 data
Ctrl+I	Import Measured Data
Ctrl+Alt+I	Import Impedance Measurement
Ctrl+Shift+I	Import Impulse Response
Ctrl+Shift+U	Import Audio Data
Ctrl+Shift+N	Import Sweep Recording
Ctrl+Backspace	Delete Current Measurement
Ctrl+Shift+Backspace	Delete All Measurements
Ctrl+Shift+K	Show Room Simulator
Ctrl+Shift+L	Show Level Meters
Ctrl+Shift+M	Show SPL Meter
Ctrl+Shift+O	Show SPL Logger
Ctrl+Shift+P	Show TS Parameters window
Ctrl+Shift+Q	Show EQ panel
Ctrl+Shift+R	Show Signal generator
Ctrl+Shift+C	Show Scope
Ctrl+Shift+T	Show RTA
Ctrl+Shift+V	Show Overlays Window
Ctrl+Shift+W	Show IR Windows panel
Ctrl+Alt+M	Show Generate minimum phase dialog
Ctrl+Shift+E	Show Preferences
Ctrl+F	Retrieve Filter Settings from Unit

Ctrl+Shift+F	Send Filter Settings to Unit
Ctrl+Delete	Reset Filters for Current Measurement
Ctrl+LEFT	Select previous graph group
Ctrl+RIGHT	Select next graph group
Ctrl+1	Select first graph group
Ctrl+2	Select second graph group
Ctrl+3	Select third graph group
Ctrl+4	Select fourth graph group
Ctrl+5	Select fifth graph group
Ctrl+6	Select sixth graph group
Ctrl+7	Select seventh graph group
Ctrl+8	Select eighth graph group
Ctrl+9	Select ninth graph group
Ctrl+Shift+G	Show Grid
Ctrl+Shift+A	Frequency Axis Log/Linear
Ctrl+Z	Undo Zoom
Ctrl+Alt+F	Fit graph to data
Ctrl+Shift+1	Apply 1/1 octave smoothing
Ctrl+Shift+2	Apply 1/2 octave smoothing
Ctrl+Shift+3	Apply 1/3 octave smoothing
Ctrl+Shift+6	Apply 1/6 octave smoothing
Ctrl+Shift+7	Apply 1/12 octave smoothing
Ctrl+Shift+8	Apply 1/24 octave smoothing
Ctrl+Shift+9	Apply 1/28 octave smoothing
Ctrl+Shift+X	Apply variable smoothing
Ctrl+Shift+Y	Apply psychoacoustic smoothing
Ctrl+Shift+Z	Apply ERB smoothing
Ctrl+0	Remove smoothing
Ctrl+Shift+J	Save Graph As JPEG
LEFT, RIGHT, UP, DOWN	Move graph cursor
Shift + LEFT, RIGHT, UP, DOWN	Move graph
x	Zoom out horizontally
Shift+x	Zoom in horizontally
y	Zoom out vertically
Shift+y	Zoom in vertically
Alt+1	Select measurement 1
Alt+2	Select measurement 2

Alt+3	Select measurement 3
Alt+4	Select measurement 4
Alt+5	Select measurement 5
Alt+6	Select measurement 6
Alt+7	Select measurement 7
Alt+8	Select measurement 8
Alt+9	Select measurement 9
Alt+UP	Select previous measurement
Alt+DOWN	Select next measurement
Alt+A	Average the responses in All SPL graph
Alt+B	Save both in RTA window
Alt+C	Show the graph capture dialog
Alt+D	Toggle Distortion button in RTA window
Alt+G	Generate Decay or Waterfall plots
Alt+N	Open stepped sine dialog in RTA window
Alt+O	Open a WAV file for processing in RTA window
Alt+P	Save peak in RTA window
Alt+R	Reset Averaging in RTA window
Alt+S	Save current in RTA window
Alt+T	Start/stop the RTA
Alt+Z	Set t=0 at cursor in Impulse graph

N.B. On macOS the shortcuts use *cmd* instead of *Ctrl* and *option* instead of *alt*

Listed in key order	
Keys	**Function**
Alt+1	Select measurement 1
Alt+2	Select measurement 2
Alt+3	Select measurement 3
Alt+4	Select measurement 4
Alt+5	Select measurement 5
Alt+6	Select measurement 6
Alt+7	Select measurement 7
Alt+8	Select measurement 8
Alt+9	Select measurement 9
Alt+UP	Select previous measurement

Alt+DOWN	Select next measurement
Alt+A	Average the responses in All SPL graph
Alt+B	Save both in RTA window
Alt+C	Show the graph capture dialog
Alt+D	Toggle Distortion button in RTA window
Alt+G	Generate Decay or Waterfall plots
Alt+N	Open stepped sine dialog in RTA window
Alt+O	Open a WAV file for processing in RTA window
Alt+P	Save peak in RTA window
Alt+R	Reset Averaging in RTA window
Alt+S	Save current in RTA window
Alt+T	Start/stop the RTA
Alt+Z	Set t=0 at cursor in Impulse graph
Ctrl+LEFT	Select previous graph group
Ctrl+RIGHT	Select next graph group
Ctrl+0	Remove smoothing
Ctrl+1	Select first graph group
Ctrl+2	Select second graph group
Ctrl+3	Select third graph group
Ctrl+4	Select fourth graph group
Ctrl+5	Select fifth graph group
Ctrl+6	Select sixth graph group
Ctrl+7	Select seventh graph group
Ctrl+8	Select eighth graph group
Ctrl+9	Select ninth graph group
Ctrl+F	Retrieve Filter Settings from Unit
Ctrl+I	Import Measured Data
Ctrl+M	Measure Response
Ctrl+N	Show Info Panel
Ctrl+O	Open Measurement
Ctrl+S	Save Measurement
Ctrl+Z	Undo Zoom
Ctrl+Alt+F	Fit graph to data
Ctrl+Alt+I	Import Impedance Measurement
Ctrl+Alt+M	Show Generate minimum phase dialog
Ctrl+Alt+O	Open Filter file
Ctrl+Alt+R	Export RT60 data

Ctrl+Alt+S	Save Filters
Ctrl+Backspace	Delete Current Measurement
Ctrl+Delete	Reset Filters for Current Measurement
Ctrl+Shift+1	Apply 1/1 octave smoothing
Ctrl+Shift+2	Apply 1/2 octave smoothing
Ctrl+Shift+3	Apply 1/3 octave smoothing
Ctrl+Shift+6	Apply 1/6 octave smoothing
Ctrl+Shift+7	Apply 1/12 octave smoothing
Ctrl+Shift+8	Apply 1/24 octave smoothing
Ctrl+Shift+9	Apply 1/28 octave smoothing
Ctrl+Shift+A	Frequency Axis Log/Linear
Ctrl+Shift+C	Show Scope
Ctrl+Shift+E	Show Preferences
Ctrl+Shift+F	Send Filter Settings to Unit
Ctrl+Shift+G	Show Grid
Ctrl+Shift+I	Import .pcm Impulse Response
Ctrl+Shift+J	Save Graph As JPEG
Ctrl+Shift+K	Show Room Simulator
Ctrl+Shift+L	Show Level Meters
Ctrl+Shift+M	Show SPL Meter
Ctrl+Shift+N	Import Sweep Recording
Ctrl+Shift+O	Show SPL Logger
Ctrl+Shift+P	Show TS Parameters window
Ctrl+Shift+Q	Show EQ panel
Ctrl+Shift+R	Show Signal generator
Ctrl+Shift+S	Save All Measurements
Ctrl+Shift+T	Show RTA
Ctrl+Shift+U	Import Audio Data
Ctrl+Shift+V	Show Overlays Window
Ctrl+Shift+W	Show IR Windows panel
Ctrl+Shift+X	Apply variable smoothing
Ctrl+Shift+Y	Apply psychoacoustic smoothing
Ctrl+Shift+Z	Apply ERB smoothing
Ctrl+Shift+Backspace	Delete All Measurements
LEFT, RIGHT, UP, DOWN	Move graph cursor
Shift + LEFT, RIGHT, UP, DOWN	Move graph
x	Zoom out horizontally

Shift+x	Zoom in horizontally
y	Zoom out vertically
Shift+y	Zoom in vertically

[Help Index](#)

File Menu

Save measurement `Ctrl+S`

Save the current measurement in a binary format with the extension ".mdat". The path to the file is remembered for the next time the dialogue appears.

Open measurement `Ctrl+O`

Load measurements from an mdat file. The path to the file is remembered for the next time the dialogue appears. Mdat files can also be opened by dragging them onto the REW main window.

Save all measurements `Ctrl+Shift+S`

Save the data for all measurements in a single file with the extension ".mdat". A note can be entered that is displayed when the data set is next loaded. The path to the file is remembered for the next time the dialogue appears.

Remove measurements loaded from...

Remove all the measurements that were loaded from the file selected from the list.

Save filters `Ctrl+Alt+S`

Saves the EQ filter settings for the current measurement in a binary format with the extension ".req". A note can be entered that is displayed when the filters are next loaded. The path to the file is remembered for the next time the dialogue appears.

Open filters `Ctrl+Alt+O`

Load a set of EQ filters for the current measurement from a .req file. Any note saved with the settings will be displayed. If the .req file has more than one set of filters a selection box is offered to choose the set to load the filters from. Req files can also be opened by dragging them onto the EQ window or the EQ filters window.

Import frequency response `Ctrl+I`

Import frequency response data (SPL or, if the file extension is .zma, impedance) from other applications or load FR data which has been saved in a text format. See [Importing Measurement Data](#) for details. The path to the file is remembered for the next time the dialogue appears.

Import impedance measurement `Ctrl+Alt+I`

Import impedance measurement data from other applications which has been saved in a text format. See [Importing Measurement Data](#) for details. The path to the file is remembered for the next time the dialogue appears.

Import impulse response `Ctrl+Shift+I`

Import an impulse response from a WAV file, an AIFF file, a TXT file or a .pcm file (raw data) as produced by ETF from Full Range measurements using the *File → Write Impulse As *.pcm option*.

The sample rate must be specified manually for .pcm files as it is not contained in the file (ETF shows the sample rate in the pop-up which appears after exporting the response). TXT files can be REW IR text exports, COMSOL model files or IR files in a 2-column format with time in ms and a "%" comment line at the start of the file. WAV and AIFF files can be mono, stereo or multi-channel, the channel to be imported is chosen when loading the file. Files with 16, 24 or 32 bit signed PCM samples and 32-bit or 64-bit float samples are supported. Only the first 1,024,576 samples are imported. To adjust the SPL offset for a desired peak SPL figure in the frequency response use the Scale Response controls in the Impulse graph group. There is a setting in the Analysis Preferences to control whether the t=0 time for the imported response is placed at the first data sample or at the peak of the impulse response. The path to the file is remembered for the next time the dialogue appears. WAV, AIFF and .pcm files can also be opened by dragging them onto the main REW window. **N.B. Mic/Meter calibration, soundcard calibration and C weighting compensation are not applied to imported impulse responses.** Use the **Change cal** button on the measurement panel to apply calibration data after import.

Import audio data `Ctrl+Shift+U`

Import audio data from a WAV or AIFF file. Files with 16, 24 or 32 bit signed PCM samples and 32-bit or 64-bit float samples are supported. Only the first 60s (approx) is imported. The path to the file is remembered for the next time the dialogue appears.
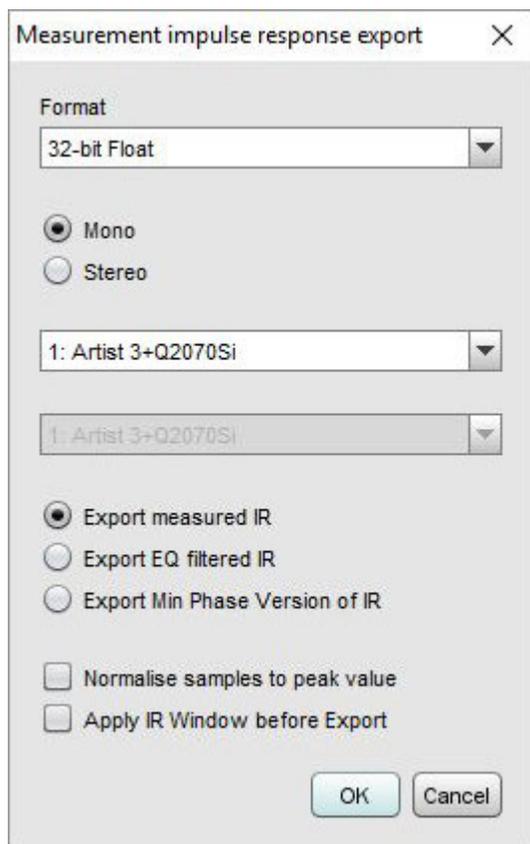
Import sweep recordings `Ctrl+Shift+N`

Import sweep stimulus and recorded responses from WAV or AIFF files to make offline measurements. Files with 16, 24 or 32 bit signed PCM samples and 32-bit or 64-bit float samples are supported. The sweep stimulus file must be produced by saving a measurement sweep from the REW signal generator with V5.19 beta 7 or later. The path to the files is remembered for the next time the dialogue appears.

Export impulse response as WAV

Export the impulse response for the selected measurement(2) in WAV format, written as mono or stereo data. The exported file will have a minimum duration of 1 second. The dialog provides options to choose whether to export the measured IR, the IR after any EQ filters have been applied to it or a minimum phase version of the IR. The response can optionally be normalised so that the peak value is unity (0 dBFS). There is also an option to apply the current impulse response window settings to the response before exporting it. The sample format can be chosen as 16, 24 or 32-bit signed PCM or 32-bit Float. Float is recommended if the application using the data can accept this, particularly if the response is not normalised. The peak of the exported response occurs 1 second after the start unless the IR window has been applied, in which case the export starts at the first windowed sample and the peak location depends on the IR window settings. That can be used to preserve relative timing for exported impulse responses that have been captured with a timing reference: define IR windows for the responses that have the same window ref time and the same left window width - use rectangular windows if the data should not be altered by the window. The exported IR for each impulse will then start at the start of the left window, which will be at the same absolute time for each, so their relative timing will be preserved.

Export all impulse responses as WAV

Export the impulse responses for all current measurements in WAV format. The exported file will have a minimum duration of 1 second. The dialog provides options to choose whether to export the measured IR, the IR after any EQ filters have been applied to it or a minimum phase version of the IR. The response can optionally be normalised so that the peak value is unity (0 dBFS). There is also an option to apply the current impulse response window settings to the response before exporting it. The sample format can be chosen as 16, 24 or 32-bit signed PCM or 32-bit Float. Float is recommended if the application using the data can accept this, particularly if the response is not normalised. The peak of the exported response occurs 1 second after the start unless the IR window has been applied, in which case the export starts at the first windowed sample and the peak location depends on the IR window settings. That can be used to preserve relative timing for exported impulse responses that have been captured with a timing reference: define IR windows for the responses that have the same window ref time and the same left window width - use rectangular windows if the data should not be altered by the window. The exported IR for each impulse will then start at the start of the left window, which will be at the same absolute time for each, so their relative timing will be preserved.

Export filters impulse response as WAV

Export the impulse response of the filters for the current measurement in WAV format, written as mono or stereo data, with the impulse response starting at the first sample in the file. The response length is 128k samples (131,072). The dialog provides options to choose the sample rate, sample format and to select whether or not to normalise the response so that the peak value is unity (0 dBFS). 32-bit Float is recommended if the application using the data can accept this, particularly if the response is not normalised. When exporting as stereo the first response selected is placed in the left

channel, the second in the right.

Export filter settings as text

Export EQ filter settings for the current measurement in a plain text format. The file includes the speaker setting and target level for the measurement. This file is a convenience format only, REW cannot load filter settings from text files - use the .req format to save and reload settings. The path to the file is remembered for the next time the dialogue appears. An example of the format is shown below.
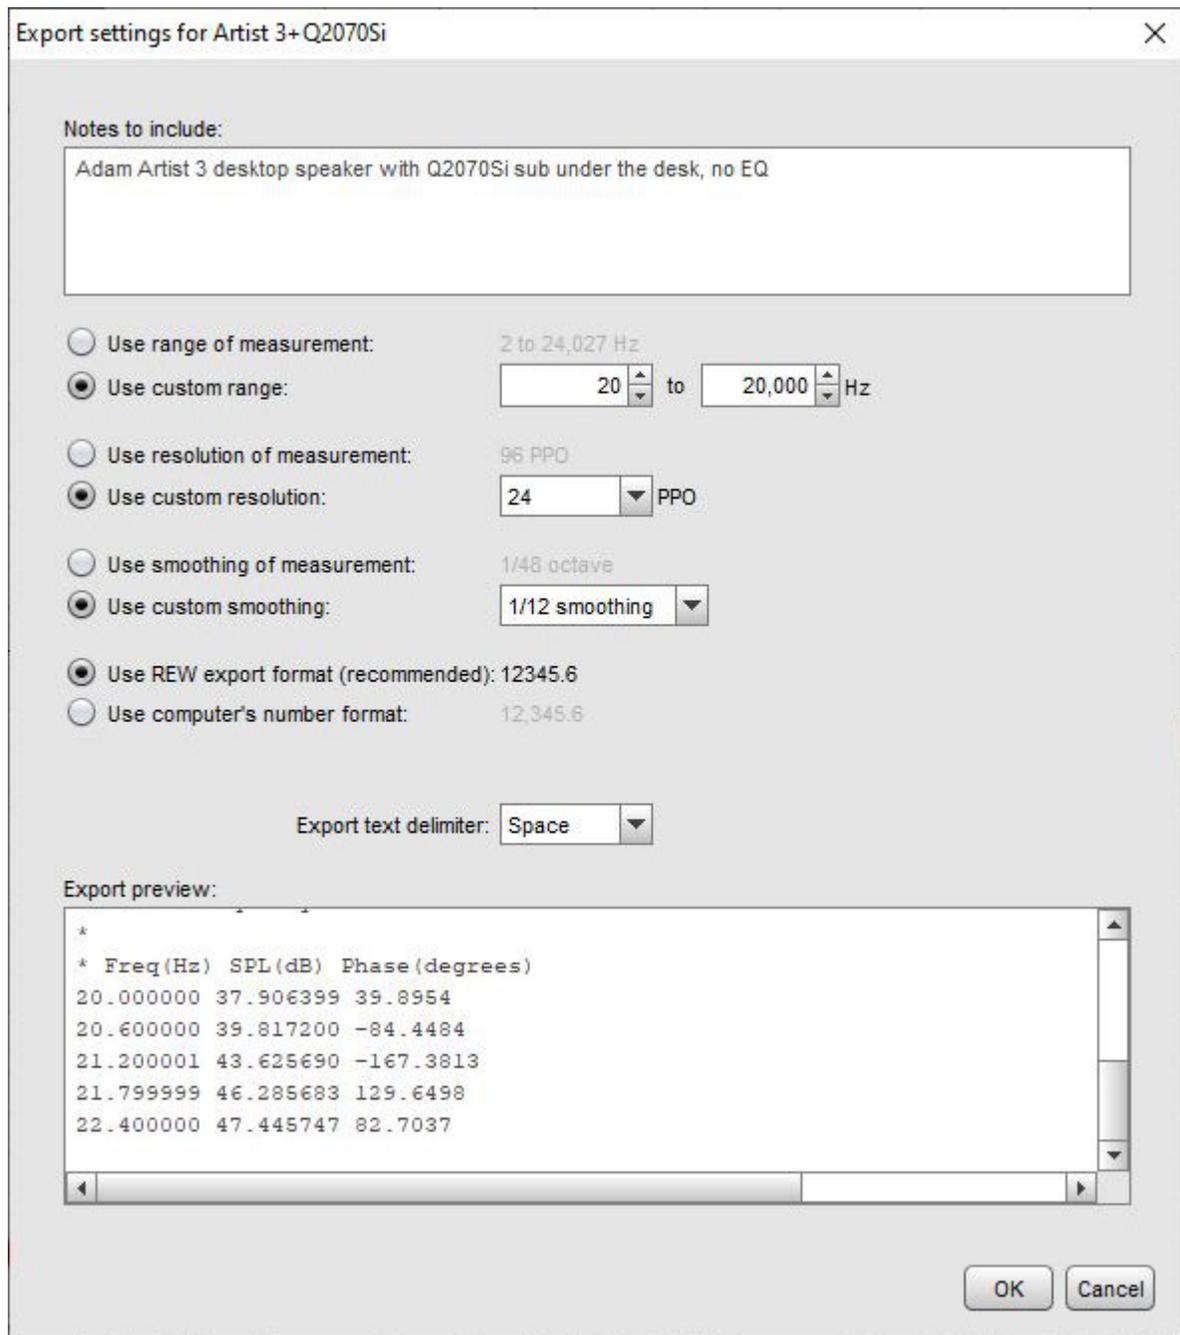
```
Filter Settings file

Room EQ V4.00
Dated: 07-Jan-2007 17:20:32

Notes:Example filter settings

Equaliser: DSP1124P
sampledata.txt
Bass limited 80Hz 12dB/Octave
Target level: 75.0dB
Filter 1: ON   PA    Fc   129.1Hz (  125 +2 )   Gain -18.5dB   BW/60   4.0
Filter 2: ON   PA    Fc    36.8Hz (   40 -7 )   Gain -15.5dB   BW/60 10.0
Filter 3: ON   PA    Fc    99.1Hz (  100 -1 )   Gain  -3.5dB   BW/60   1.0
Filter 4: ON   None
Filter 5: ON   None
Filter 6: ON   None
Filter 7: ON   None
Filter 8: ON   None
Filter 9: ON   None
Filter 10: ON  None
Filter 11: ON  None
Filter 12: ON  None
```

Export measurement as text

Export the measured data for the current measurement as a text file. The file includes the measurement settings. These text files can be re-loaded using Import Measured Data. The path to the file is remembered for the next time the dialogue appears. When the menu item is selected a settings dialog appears that allows a note to be entered and features of the export to be controlled.

Export settings for Artist 3+Q2070Si   ✕

Notes to include:

Adam Artist 3 desktop speaker with Q2070Si sub under the desk, no EQ

○ Use range of measurement: 2 to 24,027 Hz
● Use custom range: 20 ▲▼ to   20,000 ▲▼ Hz

○ Use resolution of measurement: 96 PPO
● Use custom resolution: 24   ▼ PPO

○ Use smoothing of measurement: 1/48 octave
● Use custom smoothing: 1/12 smoothing ▼

● Use REW export format (recommended): 12345.6
○ Use computer's number format: 12,345.6

Export text delimiter: Space   ▼

Export preview:

```
*
* Freq(Hz)  SPL(dB)  Phase(degrees)
20.000000 37.906399 39.8954
20.600000 39.817200 -84.4484
21.200001 43.625690 -167.3813
21.799999 46.285683 129.6498
22.400000 47.445747 82.7037
```

OK   Cancel

The entire measurement can be exported (**Use range of measurement**) or a range can be selected, if the selected range starts before or ends after the range of data within the measurement only the range that overlaps the measurement will be exported. If the data is exported at the measurement resolution that may be 96 PPO log spaced if the Allow 96 PPO log spacing Analysis preference is selected or linear spaced at the FFT resolution. Linear spaced exports are very large files (typically more than 1 MB). If a custom (log spaced) export resolution is selected REW picks a start frequency that is as close as possible to the desired start but that ensures the exported frequencies include the nominal octave centres (e.g. 20 Hz, 40 Hz etc). Smoothing should be applied when using log exports, the dialog will warn if insufficient smoothing has been selected for the chosen output resolution. Exports can use whatever smoothing is currently applied to the measurement, or any other smoothing selected from the dialog (the selection will not change the smoothing of the measurement in REW, only the smoothing for the exported data). By default REW's export number format uses dot as the decimal delimiter and no digit grouping, for ease of import in other packages subsequently, but **Use**

computer's number format will mean the current format of the computer on which REW is running will be used. The delimiter between exported values can also be selected here, this is the same selection as offered in the export menu. A preview of the export output is shown in the dialog including the first few lines of exported data. The format is compatible with the .FRD format. Comment lines start with *, data lines begin with the frequency, then the SPL in dB and finally the phase in degrees (0.0 if the measurement does not have phase information).

Export all measurements as text

Export the measured data for all loaded measurements as text files in a chosen directory. The file names use the measurement titles. The same settings dialog appears as for a single measurement, the settings are applied to all the files. Any note entered will appear in all the files, followed by any notes from each individual measurement. There is an option to export to individual files or to a single file for all measurements.

Export measurement as MLSSA .frq

Export the measurement in the MLSSA .frq format. The path to the file is remembered for the next time the dialogue appears.

Export distortion data as text

Export the distortion data for the current measurement as a text file. A dialog is displayed to control the content and format of the export

An example of the file format is shown below. Comment lines start with *. Data lines begin with the frequency or, for stepped sine measurements stepped in level, the generator level, then the fundamental SPL (if selected for export) in dB, then the selected distortion measurements. The fundamental will be 0 dB if the distortion data is normalised to the fundamental. Note that if the end frequency for the export is not configured to automatically adjust to suit the highest harmonic selected, the number of harmonics on each line will reduce as frequency increases, eventually leaving only the second harmonic and THD values.

```
    * Distortion data saved by REW V5.01
    * Dated: 16-Sep-2012 18:52:27
    * Measurement: Sep 16 18:52:27
    * Frequency Step: 3 ppo
    * Freq(Hz) Fundamental (dB) THD (%) H2 (%) H3 (%) H4 (%) H5 (%) H6 (%)
H7 (%) H8 (%) H9 (%) H10 (%)
    20.000 76.068 0.016 0.014 0.002 0.001 0.004 0.003 0.001 0.002 0.002
0.001
```

```
    25.198 76.133 0.011 0.008 0.000 0.001 0.001 0.006 0.001 0.001 0.001
0.003
    31.748 76.181 0.015 0.013 0.002 0.002 0.002 0.001 0.002 0.005 0.001
0.001
    40.000 76.215 0.014 0.013 0.001 0.003 0.001 0.002 0.002 0.001 0.003
0.001
```

Export distortion spectrum data as text

Export the distortion spectrum data for the current measurement as a text file, if it is a stepped sine measurement that includes spectrum data. Data is exported for the spectrum captured at each test frequency in a stepped sine measurement and for the noise floor. The spectrum data is log spaced, the frequency of each data point is included at the start of each row. The data header includes the frequency steps in the measurement. Note that exported spectrum data should not be used to try and calculate distortion levels as it does not have sufficient frequency resolution to accurately calculate the energy in each harmonic.

Export impulse response as text

Export the impulse response data for the current measurement as a text file. A dialog is shown to enter a note that is displayed when the data set is next loaded and preview the output. The file includes the speaker setting, target level and settings for the measurement. The path to the file is remembered for the next time the dialogue appears.
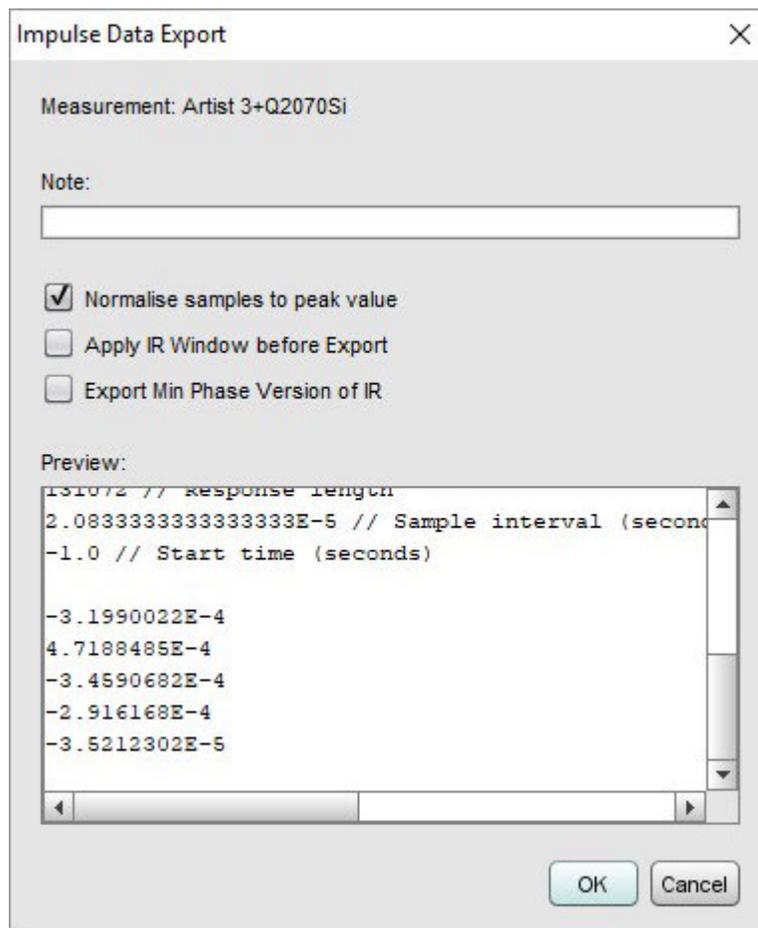
An example of the format is shown below.

```
Impulse Response data saved by REW V5.19
IR is normalised
IR window has not been applied
IR is not the min phase version
Source: Microphone (Umik-1  Gain:  0dB  ), MICROPHONE (Master Volume),
Right channel, volume: 0.550
Dated: 05-May-2014 14:12:40
Measurement: Artist 3+Q2070Si
Excitation:  256k Log Swept Sine, 1 sweep at -12.0 dB FS
Response measured over: 2.2 to 24,000.0 Hz
4.33353241533041E-4 // Peak value before normalisation
48000 // Peak index
131072 // Response length
2.0833333333333333E-5 // Sample interval (seconds)
-1.0 // Start time (seconds)

-3.1990022E-4
4.7188485E-4
-3.4590682E-4
-2.916168E-4
-3.5212302E-5
```

Export all impulse responses as text

Export the impulse responses for all loaded measurements as text files in a chosen directory. The file names use the measurement titles. The same settings dialog appears as for a single impulse response, the settings are applied to all the files.

Export RT60 data as text

Export the RT60 data for the current measurement as a text file. A note can be entered that is displayed when the data set is next loaded. The file includes both octave and one-third octave results if those have been generated. If RT60 Decay model data has been generated for octave or one-third octave a T60M figure will be included in the export, otherwise it will be zero. The path to the file is remembered for the next time the dialogue appears. A dialog is displayed to control the format of the export

RT60 Data Export

Measurement: arlo-04-0374RR

Export text delimiter: Space

○ Use REW export number format: 12345.6
○ Use computer's number format: 12,345.6

Note:

Preview:

```
                                                                        ▲
125 1/1 2.184 -0.979 1.941 -0.993 2.167 -0.995 2.179 -0.996 -5.000 -39.000 Forward -8.55 -4.
250 1/1 1.691 -0.970 1.499 -0.996 1.708 -0.996 1.725 -0.997 -5.000 -43.000 Forward -9.66 -5.
500 1/1 1.237 -0.994 1.156 -0.999 1.234 -0.999 1.169 -0.999 -5.000 -26.000 Forward -0.60 2.0
1000 1/1 1.119 -0.995 1.200 -0.999 1.177 -1.000 1.185 -1.000 -5.000 -38.000 Forward -0.69 1.
2000 1/1 1.105 -0.998 1.148 -1.000 1.170 -1.000 1.171 -1.000 -5.000 -37.000 Forward 0.15 3.4
4000 1/1 1.075 -0.997 1.114 -0.999 1.143 -0.999 1.110 -1.000 -5.000 -29.000 Forward -0.24 3.
8000 1/1 0.785 -0.996 0.982 -0.999 □ □ 0.982 -0.999 -5.000 -25.000 Forward 2.12 6.17 62.0 0.
```

OK   Cancel

Export RT60 data for all measurements as text

Export the RT60 data for all the measurements in individual text files, in the same way as the current measurement RT60 export described above. RT50 data for each measurement will only be exported if it has been generated, to quickly generate the data for all measurements view the Overlay RT60 graph.

Export → Set text delimiter

Set the character to use as a delimiter between values on a line when exporting a measurement as a text file.

Remove Current Measurement `Ctrl+Backspace`

Remove the current measurement. An unsaved measurement can be recovered using **Restore Last Removed**.

Remove All Measurements `Ctrl+Shift+Backspace`

Remove all measurements. Unsaved measurements cannot be recovered!
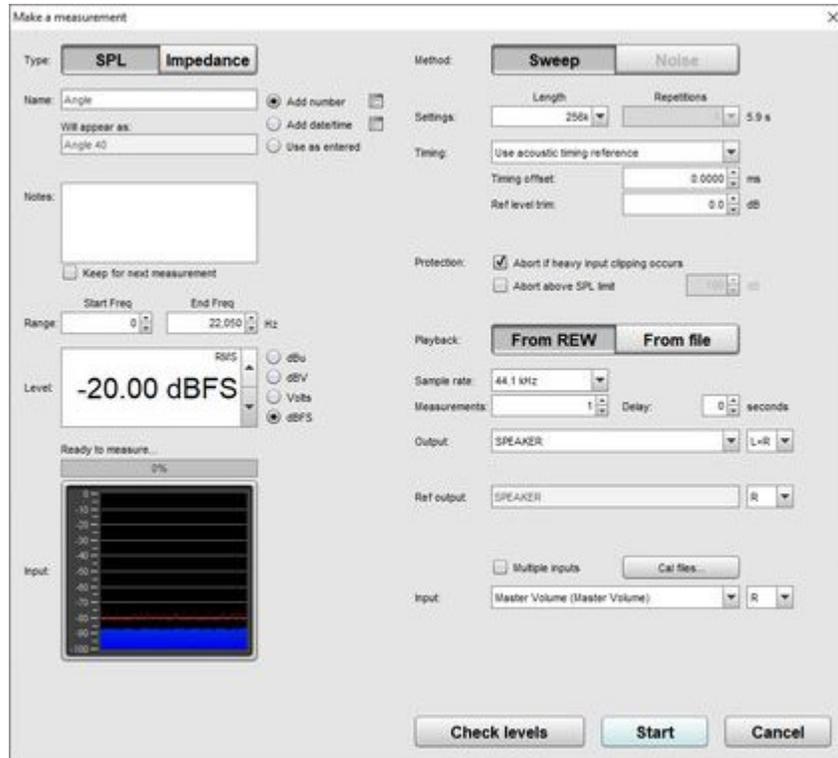
Restore Last Removed

Restore the last measurement that was removed, placing it at the end of the measurement list. Note that this is only available for individual measurements, if **Remove All Measurements** is used it is not possible to restore.

Help Index

Tools Menu

Measure `Ctrl+M`

[Make a new measurement](#) (brings up the measurement dialog)

IR Windows `Ctrl+Shift+W`

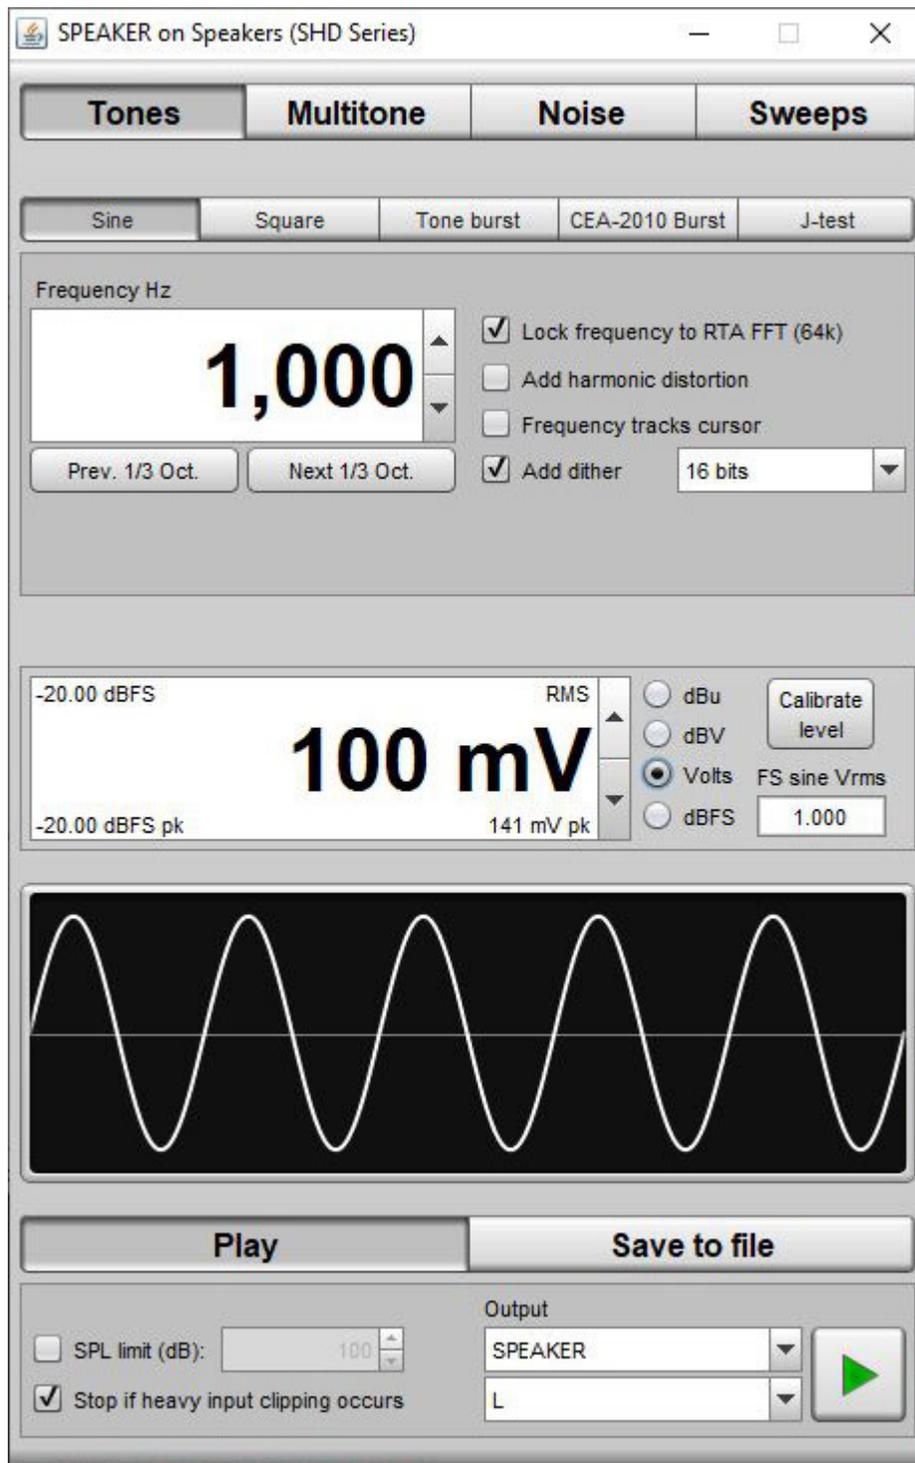Show the Impulse Response windows dialog

SPL `Ctrl+Shift+M`

Show the SPL Meter

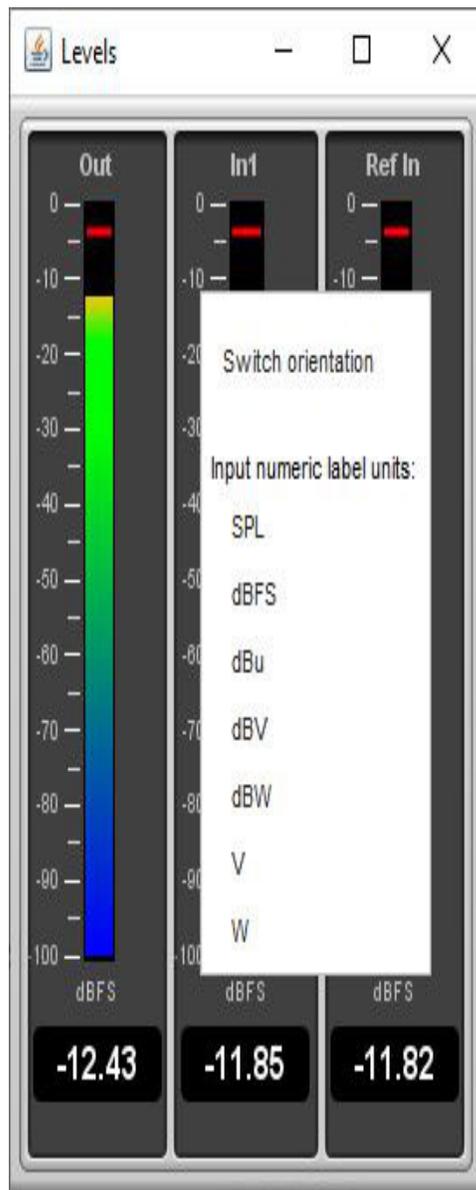### SPL Logger `Ctrl+Shift+O`

Show the SPL Logger window

Generator `Ctrl+Shift+R`

Show the Signal Generator

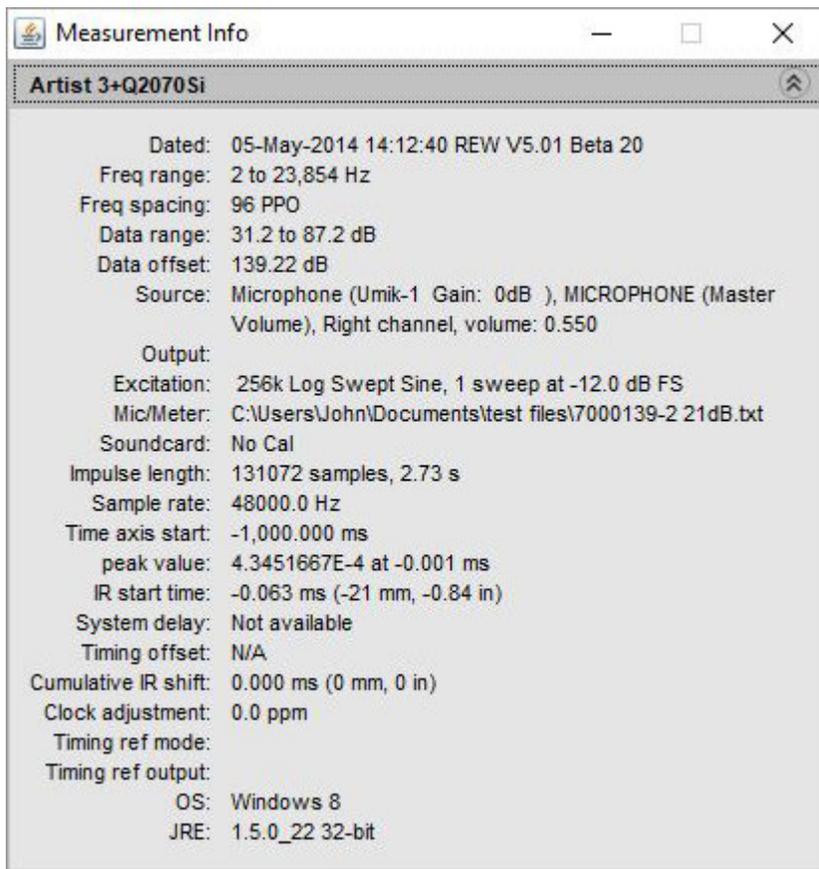***Levels*** `Ctrl+Shift+L`

Show the Level Meters

EQ Ctrl+Shift+Q

Show the EQ window

Overlays Ctrl+Shift+V

Show the Overlays window

Info Ctrl+N

Show the Measurement Info window

RTA `Ctrl+Shift+T`

Show the RTA window

Thiele Small Parameters `Ctrl+Shift+P`

Show the Thiele Small Parameters window

Room Sim `Ctrl+Shift+K`

Show the Room Simulator window

Relative gain and phase

Show the Relative gain and phase panel

Preferences Menu

Preferences

Brings up the Preferences panel. Refer to the help panels on the individual preferences tabs for more info
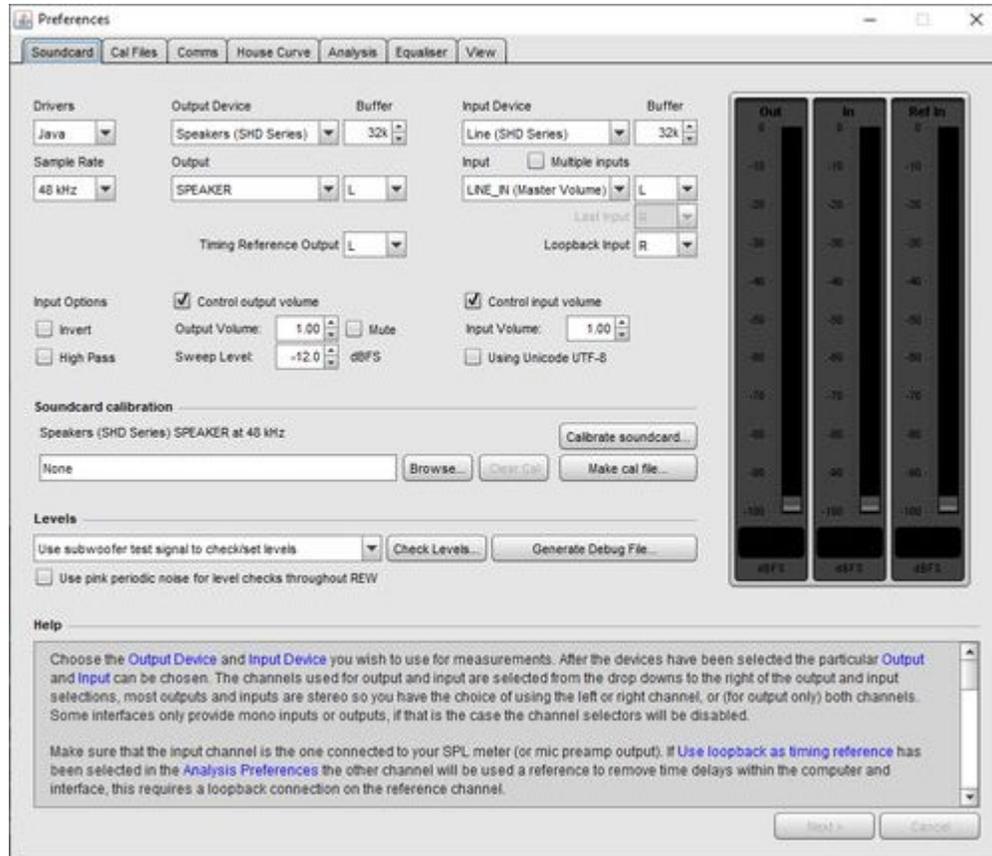
Delete Preferences and Shut Down

This command deletes all REW saved preferences and shuts down REW. When REW next starts up it will be as a clean installation. If you wish to remove REW and all its preferences from your system use this command before uninstalling REW. Preferences on Windows are stored in the registry at HKEY_CURRENT_USER\Software\JavaSoft\Prefs\room eq wizard, on macOS they are in the user's home directory at Library\Preferences in com.apple.java.util.prefs.plist in a node called roomeqwizard.

Set window size to 1024x720

This command resizes the main REW window to the specified size. Note that the OS may change the window size slightly from the requested size.

Help Index

Graph Menu

Show/Hide Grid `Ctrl+Shift+G`

Toggle the grid on/off by using this menu entry or the associated shortcut keys.

Frequency Axis Log/Linear `Ctrl+Shift+A`

Toggle the frequency axis between logarithmic and linear display. The button on the toolbar provides the same function.

Freq. Axis

Undo Zoom `Ctrl+Z`

Restore the graph axis settings that were in use when the middle or right mouse button was last pressed with the pointer in the graph area. The main application is to undo the last Variable Zoom or Zoom to Area action, but it can also be used to undo a pan or any other action that altered the axis settings since the right mouse button was last pressed in the graph area.

Fit graph to data `Ctrl+Alt+F`

Set the graph axes to show the full span of the data. The **Fit to data** button on the graph limits dialog has the same effect.

1/1 Octave Smoothing `Ctrl+Shift+1`
1/2 Octave Smoothing `Ctrl+Shift+2`
1/3 Octave Smoothing `Ctrl+Shift+3`
1/6 Octave Smoothing `Ctrl+Shift+6`
1/12 Octave Smoothing `Ctrl+Shift+7`
1/24 Octave Smoothing `Ctrl+Shift+8`
1/48 Octave Smoothing `Ctrl+Shift+9`
Variable Smoothing `Ctrl+Shift+X`
Psychoacoustic Smoothing `Ctrl+Shift+Y`
ERB Smoothing `Ctrl+Shift+Z`

Apply a smoothing filter to the current channel. Repeating the action removes the smoothing. Variable smoothing applies 1/48 octave below 100 Hz, 1/3 octave above 10 kHz and varies between 1/48 and 1/3 octave from 100 Hz to 10 kHz, reaching 1/6 octave at 1 kHz. Variable smoothing is recommended for responses that are to be equalised. Psychoacoustic smoothing uses 1/3 octave below 100Hz, 1/6 octave above 1 kHz and varies from 1/3 octave to 1/6 octave between 100 Hz and 1 kHz. It also applies more weighting to peaks by using a cubic mean (cube root of the average of the cubed values) to produce a plot that more closely corresponds to the perceived frequency response. ERB smoothing uses a variable smoothing bandwidth that corresponds to the ear's Equivalent Rectangular Bandwidth, which is (107.77f + 24.673) Hz, where f is in kHz. At low frequencies this gives heavy smoothing, about 1 octave at 50Hz, 1/2 octave at 100 Hz, 1/3 octave at 200 Hz then levelling out to approximately

1/6 octave above 1 kHz.

Remove Smoothing for Current Measurement `Ctrl+0`

Set smoothing to None. Pressing the smoothing shortcut again will also remove smoothing - for example, pressing `Ctrl+Shift+3` will smooth the data to 1/3 octave, pressing it again will remove the smoothing.
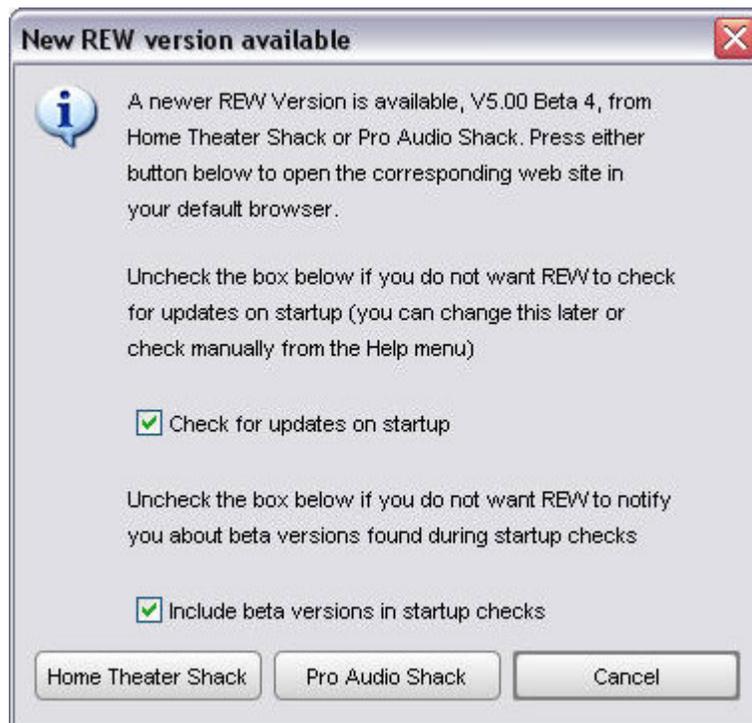
[Help Index](#)

Help Menu

Show Help `F1`

Brings up the help window

Show Help on Item `F2`

Brings up help for the next item clicked on, the mouse cursor has a "?" next to it while in this mode

Check for Updates on startup

If this is selected REW checks for a newer version on every startup

Check for Updates Now

Force a manual check for a newer version

About REW (in Application menu under macOS)

Show the About dialog
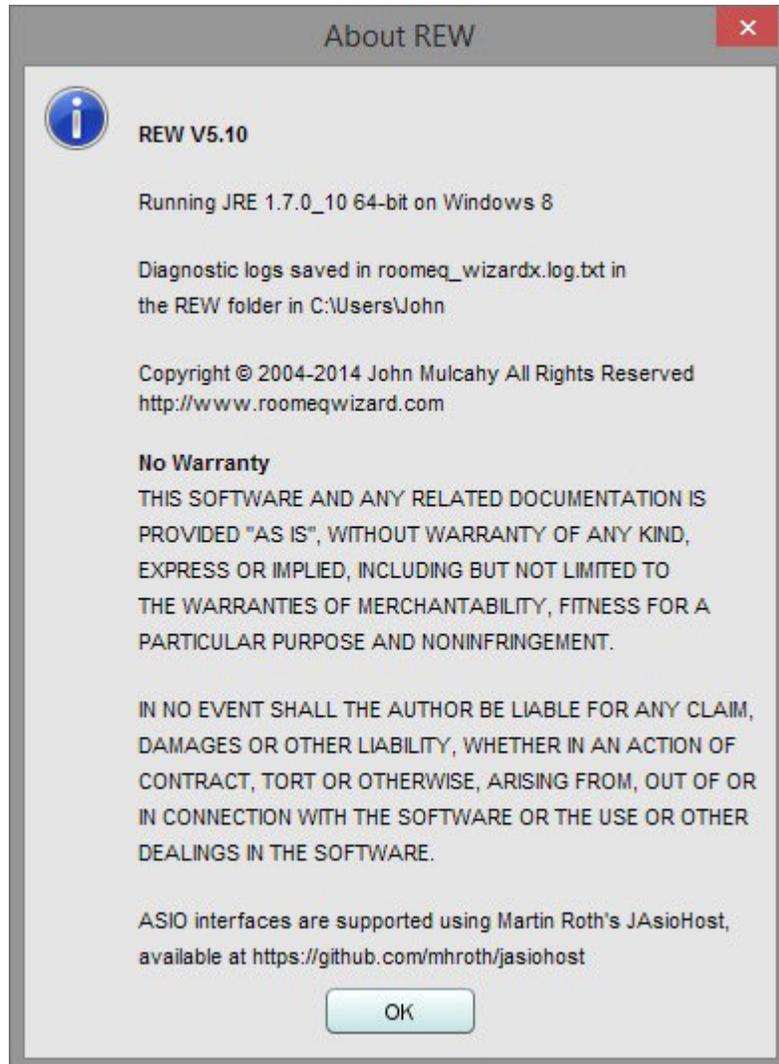
[Help Index](#)

Pro upgrades

The REW Pro Upgrade adds features for professional users. The first feature is support for capturing from multiple inputs simultaneously, further features will follow. The Pro Upgrade license unlocks multi-input capture and will also unlock all future pro features and is a one-time fee.

Multi-input capture

Multi-input Capture allows REW to capture multiple inputs simultaneously for measurement or for the RTA. Individual SPL metering is available for the first four inputs. Measurements made with multi-input capture will show an rms average of all the captured channel magnitudes with an underlay showing the span of the individual channels and can optionally include individual measurements for each channel. Each input can have its own calibration file. RTA measurements with multi-input capture show an rms average of the individual channel magnitudes.

Multi-input capture would typically be used with a mic array (such as the miniDSP UMIK-X) or a multi-channel interface with multiple analog microphones connected to it. REW captures a contiguous range of input channels. Whilst it is possible to use this feature with multiple USB microphones on Windows when using an ASIO wrapper, this is **NOT** recommended because:

- REW would only be able to access the input volume setting for the first USB mic and would use that to adjust SPL values for all the mics, so all the USB mics would need to have the same OS volume setting
- USB mics typically provide a stereo pair of inputs, each carrying the same data, so there would be two identical input captures for each mic

The first and last channels to be captured are selected on the Soundcard preferences or the Measure dialog. The example below shows selection using ASIO drivers, a similar selection is offered for Java drivers. Note that Java drivers on Windows are limited to two channels, on macOS or Linux Java drivers offer as many channels as the interface supports.

A license is valid for a single user on any computer they use.

The miniDSP UMIK-X multi-channel microphone system *includes* **an REW Pro Upgrade license.**

Before purchasing a license make sure that REW functions as required in its default single channel capture mode.

To enable your license, go to the REW Pro Upgrades menu and select Multi-input capture. Paste in your license code and click Verify license.

[Help Index](#)

Relative gain and phase

If the signal generator is playing a sine wave this panel shows the gain and phase of the measurement input relative to the reference input at the signal generator frequency, this may be useful for checking inter-channel gain and phase if both channels are fed the same signal.

[Help Index](#)