

## 1 Building Excitation Signals

Probably you've already heard that Monkey Forest offers a fantastical variety of measurement principles and excitation signals, but that it's very complicated to use and create them. This isn't true. Read this chapter carefully, and you will see that it's fairly easy to take advantage of a flexibility that hardly another measurement software offers you. The time you spend on reading these explanations is a good investment, as it will enable you to rent the congress center of a five star hotel and hold lucrative seminars for 40 Monkey Forest users or so, selling them your newly acquired knowledge about creating measurement stimuli. But beware of the snowball effect: Once you taught these MLS-eaters the fine art of custom tailoring sweep signals, they will try to snare your clients and hold their own seminars. So you better read the chapter right now to be one of the first to divulge the useful knowledge presented here.

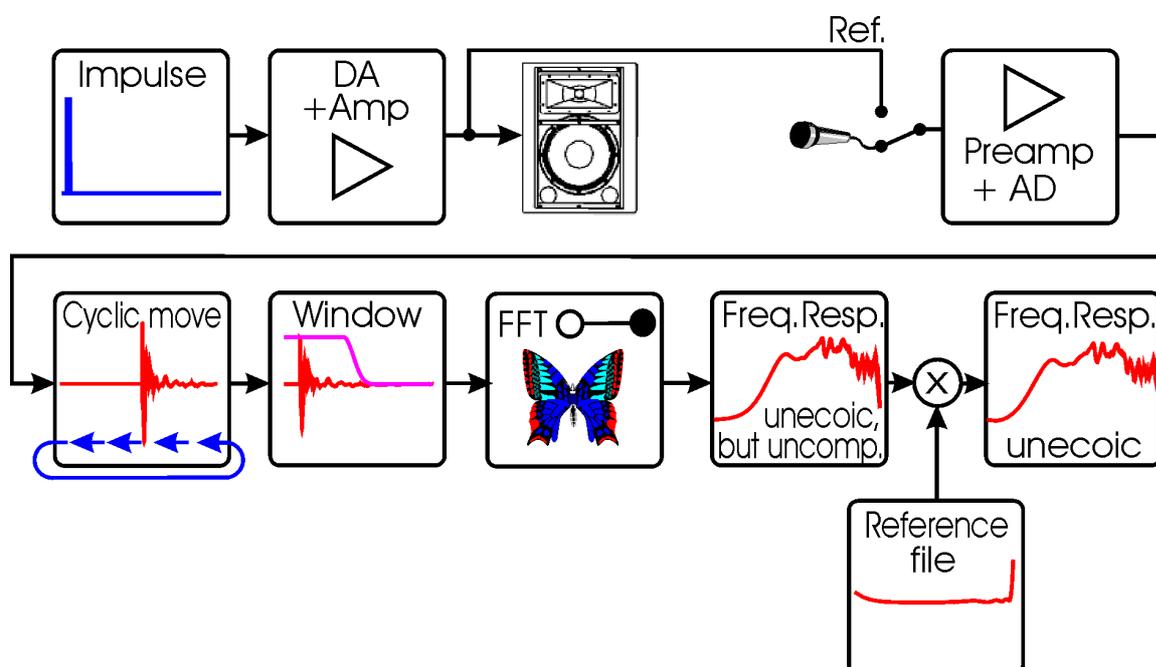
Before dropping into the "Build excitation Signal" menu, let's first review the pros and cons of the various measurement principles and the permanent need to pre-emphasize the excitation signals, in other words to give them non-white spectra.

### 1.1 Available excitation signal types

Monkey Forest offers the use of four different types of excitation signals: Pulses, noise (or generally any periodical signal of length  $2^n$ ), sweeps and maximum length sequences (MLS). Pulses and sweeps may be used in a periodic or non-periodic fashion, while noise and MLS will be used periodically and must have a fixed length ( $2^n$  for the noise,  $2^n-1$  for the MLS). Periodic use in practice means that the full length of the excitation signal will be ejected twice. In the first run, the device under test (DUT) is brought into a stationary state, and only in the second run, its response will be captured. The necessity for this proceeding lies in the fact that the desired amplitude response for noise and MLS is only valid for periodical repetition. Remember that FFTs in general assume that the  $2^n$  time signal samples to be analyzed are repeated continuously and yield the spectrum of this perpetually repeated signal. While in general, there is a drastic difference between the periodic and the non-periodic spectrum of a noise signal, the differences for impulses and sweeps are much smaller. Thus these can be used in a non-periodic manner, ejecting them only once and directly recording the response of the DUT. This will only impose a minor aberration from the desired spectrum of the excitation signal. Don't worry about errors in the measurements inflicted by this simplification: they are excluded by the use of the reference file (which, in general, is the inverse of the excitation signal's spectrum).

### 1.1.1 Impulses

Using impulses is the most straight-forward approach to perform FFT based transfer function measurements. The digitized impulse is simply sent out by the DA-converter and amplified. It feeds the DUT whose response is captured by the microphone, amplified and digitized by an AD converter. As the name implies, this captured response already is the wanted *impulse response* (IR), provided that a dirac style pulse with its linear frequency response has been used. The IR may be shifted to the left to cancel the delay introduced by the flying time between loudspeaker and microphone in an acoustical measurement. It may also be windowed to mute unwanted reflections that cause comb filter effects.



**Figure 1: Signal processing for transfer function measurement with impulses.**

The IR can then be transformed into the transfer function by FFT. To increase the precision of the measurement considerably, the result should now be multiplied with a reference spectrum. This reference spectrum is obtained by linking output and input of the measurement system and inverting the measured transfer function. Applying this technique offers the crucial freedom to pre-emphasize the used impulse to adapt it to the spectral contribution of background noise. This pre-emphasis will automatically be removed by performing the reference measurement and applying the obtained spectrum in all subsequent measurements.

Impulses are a simple and not too bad choice when the measurement is purely electrical (no acoustic path in the measurement chain), and the measurement should be as fast as possible. However, they require a low noise floor of the DUT to achieve reasonable measurement certainty. When measuring low noise audio equipment, this requirement is

often fulfilled. But despite of their far-from-optimal S/N performance, impulses can even be useful in acoustics. In an anechoic chamber, tweeters can be measured with reasonable S/N ratio, as due to their shortness, pulses can be fed with very high voltage without having to consider the danger of overheating the voice coil.

Impulse testing doesn't allow identifying distortion, but is immune against the detrimental effects of time variance that frequently haunt MLS or noise-based outdoor measurements.

### **1.1.2 Maximum length sequences (MLS)**

MLS are binary sequences that can be generated very easily by either soft- or hardware. Their length is  $2^n-1$ , with  $n$  being an integer number. Though not having an exact FFT block length, MLS have some unique properties that make them interesting for transfer function measurements. Their Auto-correlation comes close to a dirac pulse. Repeated periodically, all frequency components have exactly the same amplitude, meaning their spectrum is perfectly white. Compared to a pulse of same amplitude, much more energy can be fed to the DUT as it is now stretched out over the whole measurement period. This means increased S/N ratio.

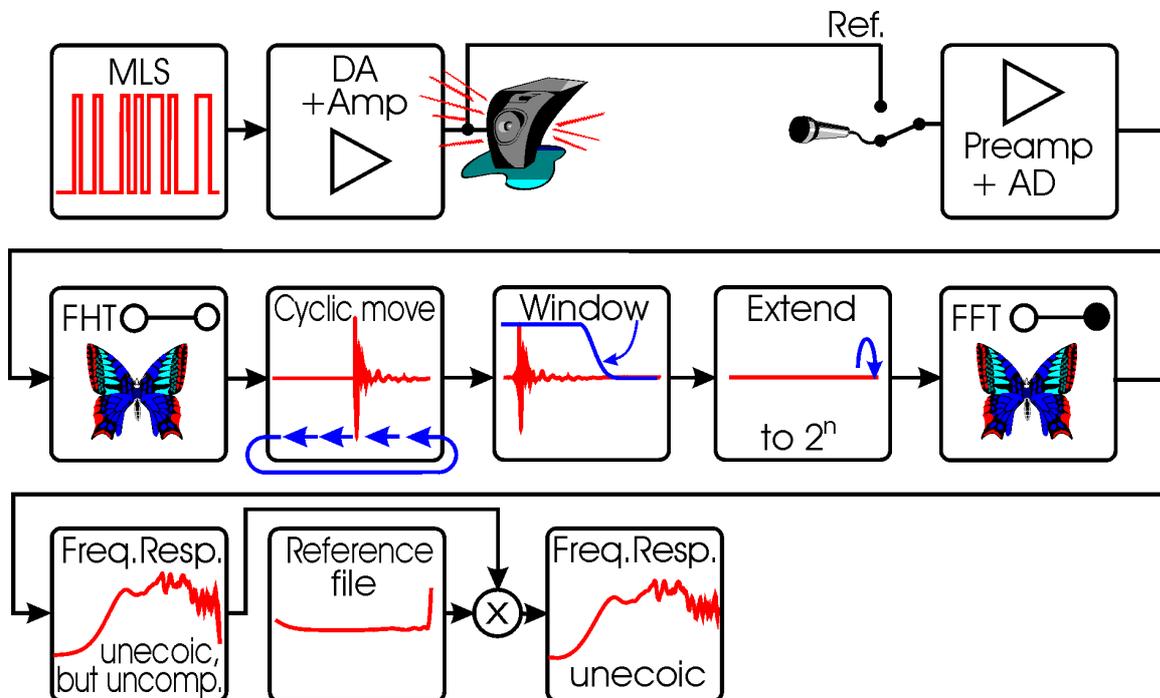
The most interesting aspect of MLS, apart from their easy generation, is the well-known fact that their use allows the application of an efficient cross correlation algorithm, the fast Hadamard transform (FHT), to calculate the impulse responses. The butterfly algorithm employed in the FHT only uses additions and subtractions and can operate on the integer data delivered by the AD converter. In former times, this meant calculation times that were much shorter than that of an FFT of similar length, but today this difference has shrunk a lot. Modern processors as those of the Pentium family are able to perform floating-point multiplications, additions and subtractions as fast as the respective integer operations.

In times when the time to calculate a long room impulse response was in the range of many seconds, this time-saving property of the Hadamard transform was very welcome. The advantage became especially prominent when just the impulse response, but not its associated transfer function was of interest. This, for example, holds for the evaluation of reverberation times by backward integration of (filtered) room impulse responses. In these cases, just one FHT is required to transform the MLS-response captured by the microphone into the desired IR. And this FHT is faster than an FFT. In contrast, using arbitrary noise signals or sweeps as stimulus requires at least one FFT and one IFFT to retrieve the IR, as will be seen in section 1.1.3.

In an MLS based measurement, the FHT is the first signal-processing step after digitization by the AD converter. The resulting impulse response can be shifted in a cyclic fashion and windowed, as with simple impulse testing.

If the transfer function is what is being looked for, it can be obtained by FFT. But as MLS have a length of  $2^n - 1$ , one sample has to be inserted to patch the impulse response to full  $2^n$  length. While computationally, this is a trivial action, care has to be exercised where to place this sample. It must be in a region where the impulse response has decayed to near zero to avoid gross errors. When using a window, it's easy to find such a region.

The acquired transfer function again can and should be corrected by multiplication with a reference spectrum obtained previously by a self-response measurement.

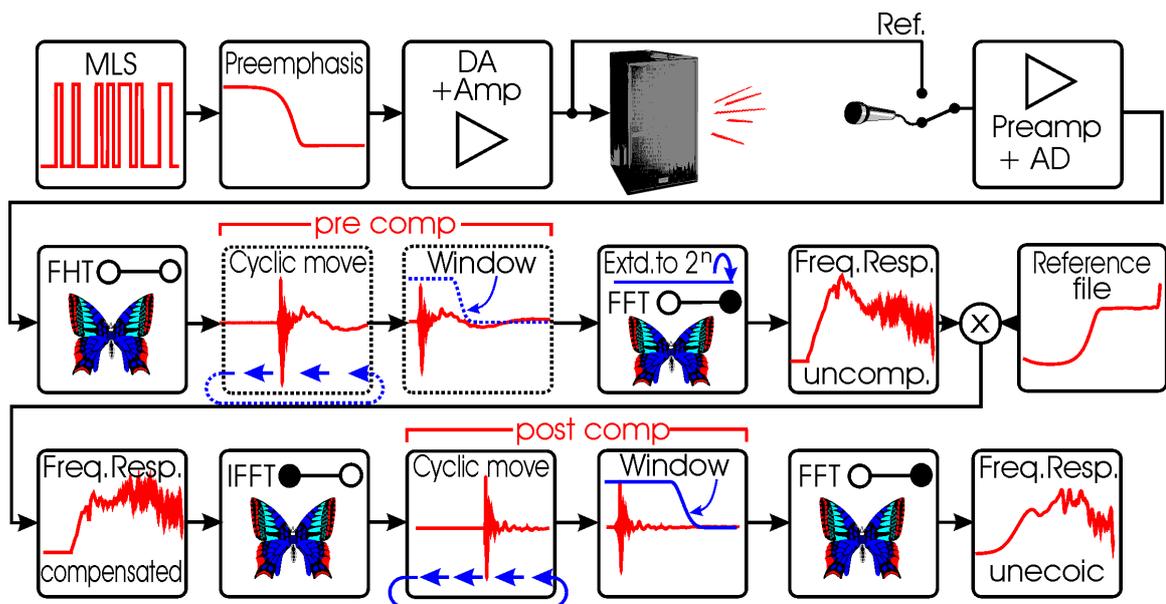


**Figure 2: Signal processing stages for transfer function measurements with MLS.**

MLS measurements have proved quite popular in acoustics, but have many drawbacks. Among high vulnerability against distortion and time variance, the most annoying property of MLS is their white spectrum. As will be revealed further in section 1.2, a non-white spectrum is desirable for almost all acoustical measurements. This requirement can be fulfilled by pre-filtering the MLS with an appropriate emphasis. Clearly, the MLS will lose its binary character by pre-filtering. Thus, this technique is only viable if the pre-filtered MLS can be output by a true DA converter, not just a one bit switching stage. Creating an emphasized MLS can be done most efficiently by means of the inverse fast Hadamard transform (IFHT). The IFHT simply consists in time inverting the impulse response of the desired emphasis filter (curtailed to  $2^n - 1$  samples), applying a normal FHT

on the impulse response and then time inverting the result again. This will yield an MLS periodically convolved with the emphasis filter.

When a pre-emphasized instead of a pure MLS is being used, obviously the impulse response obtained by FHT will yield the IR of the DUT convolved with the IR of the emphasis filter. In the common case of using MLS with a strong bass boost, this IR will be much broader than the one of the DUT alone. This broadening often constraints windowing, especially when reflections that shall be muted are in close proximity of the main peak. In these cases, applying a window to the non-equalized impulse response will noticeably attenuate the low frequency energy spread out in time. Thus, it is better to perform the windowing “post-comp”, meaning after transforming the uncorrected impulse response into the spectral domain, multiplying it with the inverse emphasis frequency response, and eventually back-transforming into the time domain. This will yield the true impulse response of the DUT that can now be windowed with lesser low frequency energy loss.



**Figure 3: Signal processing stages for transfer function measurements with pre-emphasized MLS.**

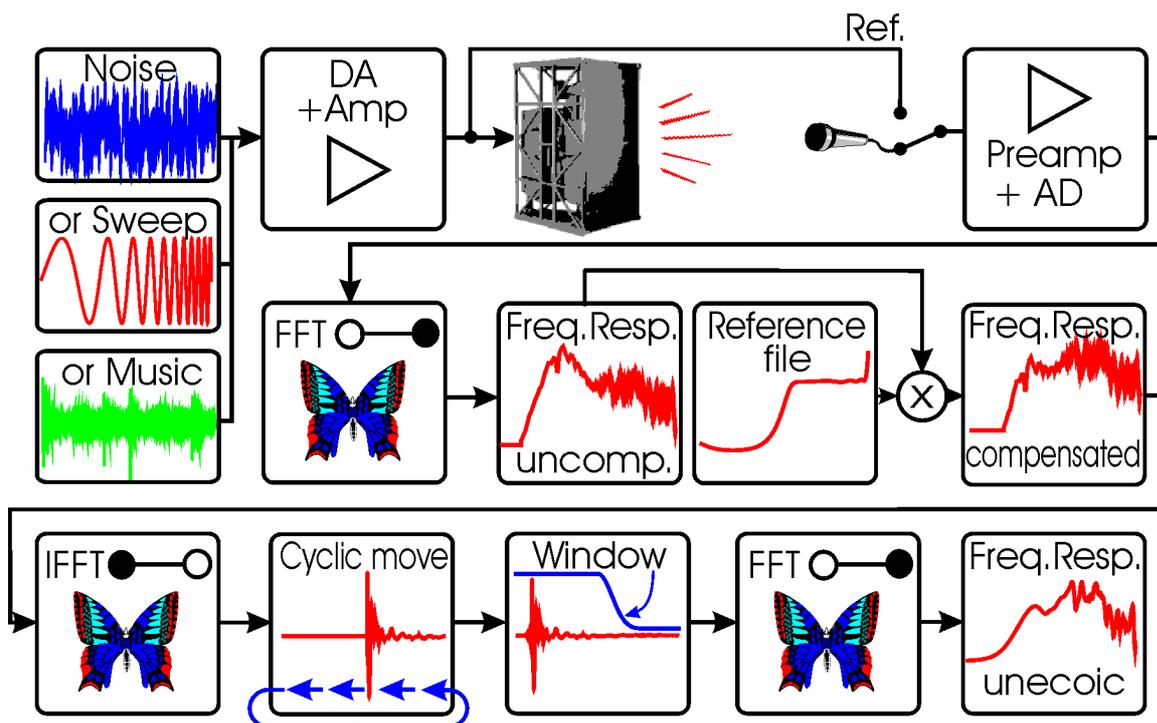
If the transfer function is the desired result of the measurement, another final FFT will have to be performed. So the number of transformations is totaling to 1 FHT and 3 FFTs when measuring transfer functions with pre-emphasized MLS and applying post-comp-windowing.

To summarize the properties of MLS, they have the advantage of increased power stretched out in time, but are very vulnerable to time variance and distortion. However,

they are still very popular in loudspeaker and room acoustical measurements, and in non-critical conditions, yield reasonable results.

### 1.1.3 Noise or just any kind of signal

Looking a bit closer to the MLS measurement setup of Figure 3 reveals that actually, the use of MLS and application of the FHT are pretty superfluous. More flexible and simpler is the use of a periodic signal that already has the required FFT block length of  $2^n$  samples. Omitting the FHT, the DUT's response to this signal can be transformed directly to the spectral domain. There, it is multiplied with the reference spectrum (the inverse of the product of the excitation signal's spectrum and the measurement system's frequency response). As this multiplication is a complex operation, not only the magnitude, but also the phases are corrected to yield the true complex transfer function of the DUT alone, regardless of the excitation signal's nature. Performing an IFFT on the compensated spectrum will produce the correspondent true impulse response of the DUT.



**Figure 4: Signal processing stages for transfer function measurements with any deterministic signal.**

When comparing this technique (FFT, compensation, IFFT) to the FHT, it becomes clear that it is far more powerful and flexible, allowing the use of arbitrary signals of length  $2^n$ . While the FHT is only able to equalize the phases of a special class of excitation signals, namely MLS, the FFT approach equalizes phases *and magnitude* of any excitation signal, be that noise, sweeps or even music.

Clearly, performing two FFTs consumes considerably more processing time than one single FFT. But with today's powerful microprocessors, this drawback is insignificant. Using the typical audio sampling rates, nowadays even very long FFTs can be performed many times faster than real time, allowing to emit the excitation signal continuously and to complete calculation and display updating within every period.

As we have seen, using pre-emphasized MLS and post-comp windowing even leads to longer calculation times than using arbitrary  $2n$  signals. Generally, it is far more beneficent to use emphasis and windowing "post-comp" than to save a few milliseconds processing time. And you are lucky enough to use software with extremely fast and efficient hand-optimized assembly language routines to do the number crunching.

As stated, just any signal of FFT block length can be employed as stimulus using this kind of signal processing. However, in order to guarantee a sufficient S/N ratio in the whole frequency region of interest, it is better to synthesize custom tailored signals by first defining their magnitude spectrum. According to the desired signal type, the corresponding phase spectrum will be constructed in quite different manners.

A noise signal can be generated easily by setting its phases to random values. The excitation signal is obtained by IFFT. Noise signals have similar properties as MLS, especially concerning their vulnerability against distortion and time variance. Some people refer to noise signals as "multi sine signals", but of course, just any non-pure tone is a multi sine signal.

#### **1.1.4 Sweeps**

Sweeps are the best choice for almost every measurement situation. They allow feeding the DUT with high power at little more than 3 dB Crest factor, and are pretty tolerant against time variance and distortion. The harmonic distortion products can be separated entirely from the impulse response. Moreover, they can be classified into single frequency-dependent harmonics, allowing a complete and ultra-fast THD(f) analysis together with the evaluation of the transfer function. When it comes to capturing room impulse responses for auralization purposes, there is no alternative to sweep measurements: The high dynamic range in turn of 90 dB required for this purpose is unattainable with MLS or noise measurements.

The construction of sweeps is somewhat more sophisticated and will be illuminated in detail in section 1.6.4.

### 1.1.5 Measurement duration

In all measurements, the AD capture period has to be at least as long as the impulse response itself (in practice: the time until the response has gone under in noise) to avoid errors. This is obvious for the measurement with a non-periodic pulse. All its energy is emitted at the very beginning, and the AD converter simply has to collect samples until the impulse response has decayed. In case of a non-periodic sweep being used as excitation signal, the capture period has to be a little bit longer, but in general not much. This is thanks to the sweep's nice property to start with the low frequencies. With normal DUTs such as loudspeakers, the largest signal delays will occur at the low frequencies. Thus, when sweeping through the high frequencies, there should be sufficient time for the delayed low components to stumble behind. For loudspeaker measurements, the decay for the highest frequencies is usually so short that the AD capturing can be stopped almost immediately after the excitation signal swept through (provided that the sweep is considerably longer than the delay at low frequencies, and of course observing the flying time between loudspeaker and microphone). In this case, the measurement time has to be only insignificantly longer than in the case of a single pulse.

In room acoustic measurements, the gap of silence following the emission of the sweep usually has to be as long as the reverberation at the highest frequencies.

In the case of periodic excitation, the period length and AD capture time has to be no longer than the impulse response length. Using a shorter length would lead to time aliasing, that is, folding back the tail of the IR that is crossing the period's end to its beginning where it is added to the start of the IR. Depending on the amount of energy folded back, this creates more or less tolerable errors. Compared, to non-periodic pulse or sweep-measurements, the need to emit the excitation signal twice means longer measurement durations than required. Moreover, only half of the total energy fed to the DUT is used for the analysis.

## 1.2 Pre-emphasis

In almost any acoustical measurement, it is not advisable to use an excitation signal with white spectral contents. When, for instance, measuring a loudspeaker in an anechoic chamber, two effects account for a heavy loss of S/N ratio at low frequencies: The loss of sensitivity (with 12 or 24 dB per octave) below the (lowest) resonance of the woofer, and the increase of ambient noise in this frequency region due to the wall's decreasing insulation against outside noise. So, in order to track the speaker's low frequency roll-off without the corrupting effects of low S/N, a strong emphasis of more than, say 20 dB is required to establish reasonable measurement certainty.

This also results in a better contribution of the power fed to a multi-way loudspeaker. The woofer typically endures much higher power than the tweeter. But when using a white excitation signal, the tweeter has to bear the brunt of the excitation signal's energy. A dome tweeter can be overheated and damaged with as little as a few watts, and this limit can easily be trespassed by any power amplifier. So also from this point of view, an emphasis of lower frequencies is highly desirable.

A third reason to mention is less of physical but of social nature: When measuring on site (concert hall, stadium etc...), an excitation signal with bass boost will sound warmer and more pleasant than white noise, and hence is more acceptable by people present. Moreover, a strong increase of power in the low band will not correspond to the impression of much increased loudness due to the decreased sensitivity of our hearing sense. The riggers climbing on some trusses near the PA system definitely appreciate unobtrusive excitation signals.

In room acoustical measurements (in which not the employed loudspeaker, but the room is the object of interest), using pre-emphasis is even more important. For this reason, we will talk about it in the separate next section (1.3).

Monkey Forest accommodates everything to handle the pre-emphasis issue with ease. As a matter of fact, only few cases are imaginable in which a deviation from the white spectral contents is *not* desirable. And after imbibing these pages, you will skillfully use the tools to create pre-emphasized stimuli adapted optimally to whatever specific measurement task.

### **1.3 Equalizing loudspeakers for room acoustical measurements**

Acquiring room impulse responses is one of the most important and ever repeating tasks in room and building acoustics. All typical parameters that describe the acoustical properties of a room (or, to be more precise, the acoustical transmission path between two points in the room) such as reverberation, clarity, definition, center time, STI and many others can be derived from it. A closer look on the impulse response can help to identify acoustical problems such as unwanted reflections or insufficient ratio between direct sound and reverberation. Visually examining the associated room transfer function (obtained by FFT) can reveal disturbing room modes or, of course, tonal misbalance of a sound reinforcement system.

Another very interesting application is the creation of “virtual reality” by convolution of dry audio material with binaural room impulse responses captured by a dummy head. We will later see that only sweeps are applicable to fulfill this task with sufficient dynamic range.

As a modern acoustician without shyness of using computers, you do not want to rely on techniques taken from the stone age, such as firing a pistol (delicate action especially in churches) or stinging balloons (more adequate for children birthday parties). You know that these methods have a very poor repeatability and produce unpredictable spectra. The only bet to avoid these severe drawbacks is the use of an electro-acoustical system, and hence bring a loudspeaker into play.

Obviously, when using a loudspeaker without any further precautions, the acquired room transfer functions will be colored by the loudspeaker's frequency response. This is particularly a problem when the RIA shall be used for auralization. To worsen things, the frequency response is direction-dependent. For room acoustical measurements, the ISO 3382 prescribes the use of a loudspeaker "as omni-directional" as possible (a condition that in practice cannot be satisfied over the full audio up to the highest frequencies).

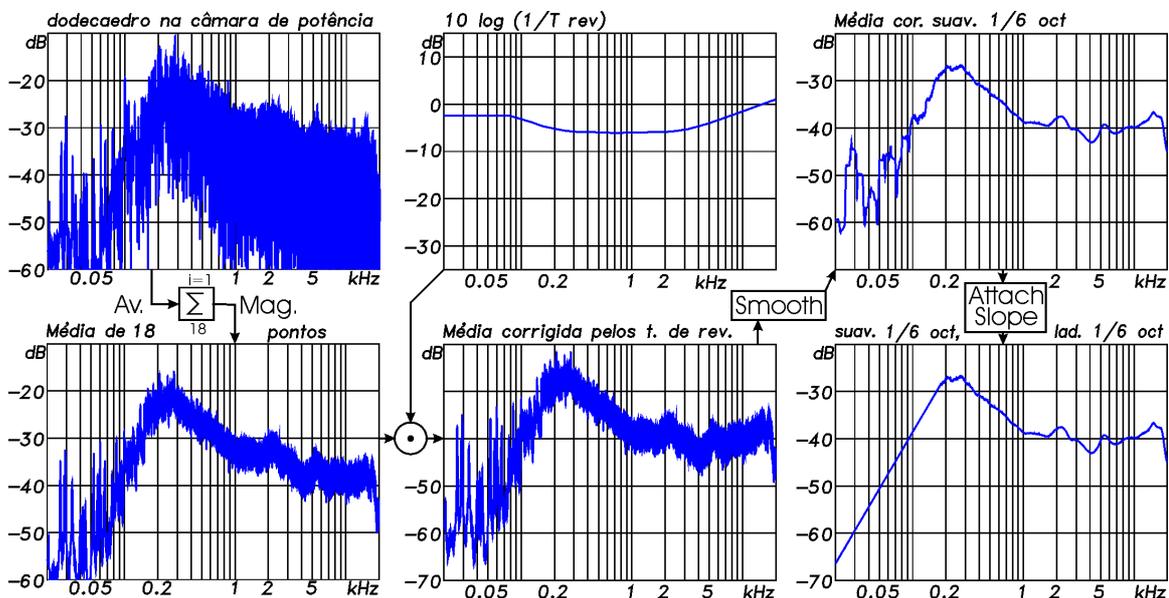
No loudspeaker will be able to produce a frequency-independent acoustical output. This is not a dramatic problem in reverberation time measurements in octaves or third of octaves, as long as the deviation within these filtered bands doesn't get too high. Alas the coloring of the RTFs by the loudspeaker's frequency response is highly undesirable for auralization purposes. In these cases, it's mandatory to use a pre-emphasis to remove the coloring. Of course, this equalization could also be done by post-processing the RTFs with the inverse of the speaker's response, but this wouldn't improve the poor S/N ratio at frequency regions where the acoustical output of the loudspeaker is low. So the better bet is to pre-filter the excitation signal to allow for a frequency-independent power output.

But this is only one component of the pre-emphasis that should be applied in room acoustical measurements. Additionally, they benefit a lot from adapting the emitted power spectrum to the ambient noise spectrum. And in most of the cases, this background noise tends to be much higher in the lower frequency regions. Thus, in order to achieve an S/N-ratio that is almost independent of frequency, there should be an extra pre-emphasis that tracks the background noise spectrum.

While a frequency independent S/N-ratio certainly is desirable for room acoustical analysis (especially for the extraction of reverberation times), it may be argued that in order to minimize audible noise, collected room impulse responses for auralization purposes better had an specifically high S/N ratio in the mid frequency region of highest ear sensitivity. On the other hand, giving a boost in this region will lead to particularly annoying excitation signals, while emphasizing the bass region will not do so. Anyway, the question which emphasis is best suited is a multifaceted problem and may be answered differently for every measurement scenario. Thus, it is desirable to be totally free in the choice of extra pre-emphasis added to the speaker equalization.

Only this second component of the pre-emphasis filter has to be equalized in the measured transfer functions (by multiplication with its inverse). Clearly, the speaker equalization component doesn't have to be removed. For this reason, the reference file used for room acoustical measurements cannot be created by an electrical reference measurement, as this would lead to canceling both emphasis curves.

A conundrum is how to handle measurements for acoustic power equalization of measurement loudspeakers. Usually, the acoustical power response of the speaker is obtained by magnitude averaging many transfer functions measured in the diffuse field of a reverberant chamber. A correction with the inverse of the reverberation times ( $10 \log \{1/T(f)\}$ ) converts the diffuse field sound pressure spectrum into a spectrum proportional to the acoustical power.



**Figure 5: Steps to construct the acoustical power spectrum of a loudspeaker**

After this, a smoothing with 1/6 or 1/3 octave is indispensable to obtain a non-corrugated curve. But this smoothing will not be sufficient for low frequencies, as can be seen in the top right curve in Figure 5. In this range, the chamber transfer function consists only in a few single high-Q modes. It's a good idea to replace these needles by a slope with the theoretical decrease of the speaker's sensitivity (12 dB/oct. for closed box or 24 dB/oct. for vented systems). Use Monkey Forest's "Edit/Here comes more/Slope" slope function to attach this synthetic slope.

The problem now is how to deal with the phase and the associated group delay. Evidently, the phase of a reverberant chamber's transfer function cannot be used. A feasible compromise is to combine the acoustic power magnitude response (as obtained in the reverberant chamber) with a free field phase response (as obtained by measuring

the speaker's sensitivity on-axis). In Monkey Forest, this can be accomplished without fuzz by first loading the diffuse field measurement and then performing an Edit/Read file/phase operation.

Obviously, merging the amplitude of one measurement with the phase of another leads to an artificial spectrum that will also correspond to a synthetic impulse response with some artifacts. However, this seems to be the closest way to acquire room transfer functions that are neither corrupted by amplitude nor by phase distortion.

It is interesting to note here that if a sweep shall be created, only the target's magnitude response will influence the excitation signal. The excitation signal's group delay can be synthesized by either using a fixed linear/logarithmic increment, or by making the increment proportional to the target's squared magnitude response to obtain a constant envelope sweep with arbitrary energy distribution. Having no influence on the creation of the sweep, the measurement loudspeaker's phase can thus only be equalized by post-processing (that is, applying the inverse of this phase to the reference file).

So, as we have seen, for general room acoustical measurements, the signal processing will always consist of a combination of pre- and post-processing.

#### 1.4 For quick starters

If you do not want to bother with the theory and all refinements of the art of creating excitation signals, here's an express manual to start measuring immediately with a custom tailored constant envelope sweep.

- Enter into the AD/DA measurement menu in which you would like to use a new excitation signal (e.g. frequency response, loudspeaker sensitivity etc..)
- Enter into the “**reference and excitation**” menu by typing “R”
- Enter into “**Build new sig...**” by typing “N”
- Don't let yourself confuse by the bewildering variety of switches and settings in the awful menu popping up. In fact you are now in the “**Build Excitation Signal**” menu that offers you all freedom and flexibility to create whatever excitation signal. But to get through it rapidly, only look at the top line and press ‘#’ to enter into the “**setup**” submenu.
- In the “**setup file menu**” evoked by this last keystroke, take a look at the list of presets on the middle of the left side. Chose one of the presets numbered 1..5 that matches best your measurement purpose.

- Back in the “**Build excitation Signal**” menu, simply press <Enter>. The desired excitation sweep will be generated along with a preliminary reference file.
- Back in the reference menu, press <Ctrl> <Enter> if you want to create a precision reference file (by measuring the self response of your AD/DA hardware). If you do not use the ITADDA16+Robo combination, prepare this reference measurement by connecting input and output first. The signal used as excitation will already be the one you just created.
- Back in the specific AD/DA measurement menu, adjust everything to your needs and simply start the measurement by typing <Enter>.

Was that really difficult? If you do not think so, let’s quickly plunge into a second session. That will give you the satisfactory feeling of not just blindly having accepted obscure standard settings, but of having interactively formed your own personal special unique measurement stimulus.

### 1.5 Second session for enthusiastic quick starters

- Perform the same 5 first steps as in chapter 1.4.
- Back in the “**Build excitation Signal**” menu, perform an audacious stroke on the key “D” to enter in the sinister “Create preemphasis curve” menu
- Risk an even more daring stroke on the “V” key to enter into the “**Interactive**” J-Filter mode.
- Now use the Cursor left/right key and the ‘,’ and ‘.’ key to alter the displayed pre-emphasis curve “online”. Once you like the result, press <Enter> to return to the “Create preemphasis curve” menu.
- Press ESC or <Enter> to return into the “**Build excitation Signal**” menu. Here, change the three last characters of the “**Signal file name**” prefix from “NEW” into something different. This way you won’t overwrite your precious excitation signal created in chapter 1.4.
- There is more to play around with, but to avoid overstressing your patience, type <Enter> now to create the second excitation signal.

## 1.6 The Menu “Build Excitation Signal” in detail

Now that you have made your first steps, let’s take a plunge into the deep end. Here’s a complete description of all settings you already saw by a glimpse when performing the sessions for quick starters. There are quite a few options in the BEX main menu, but don’t yourself discourage from this dazzling variety. Normally, you don’t have to bother about all settings.

### 1.6.1 Section “Excitation Signal properties”

Impulses, Noise, sweeps and MLS can be created, all of them with the same pre-emphasis tools. For each excitation signal, a corresponding reference spectrum has to be created. This reference file will later be applied multiplicatively to deconvolute the excitation signal’s and measurement system’s magnitude and phase responses, producing the pure transfer function (and correspondent impulse response) of the DUT alone. It may be created by measurement or by simulation, as described in 1.6.5

By default, either the created excitation signals and their corresponding reference files are stored in the same directory where MF.EXE is housed (typically C:\MF). This manner, the same excitation signals can be accessed by different users and from different directories. On the other hand, you have to be wary not to overwrite stimuli that other people still need. Monkey Forest will issue the usual warning when you attempt to write a file that bears a name already existing. If you do not like the signals to be saved there, nobody deters you from storing them in another directory. The excitation signal names proposed by MF adhere to the following rules: The first three characters characterize the kind of stimulus (IMP, NOI, SWP, MLS). The next two digits represent the degree of the signal. The remaining three characters are copied from the beginning of the target file name, if a target file is involved. If not, they are set to “NEW”. You can and should edit only these three last characters to create unique file names. Again, this nomenclature is merely a proposal and you can use whatever other file names.

The quantization with which the excitation signal will be stored on disk is by default 16 bits. Normally, there’s no need to mess with this value. But if you want to store a huge set of very long signals and are lacking disk space, try reducing this value to 8 bits. On the other hand, if you are experimenting with distortion measurements involving sweeps, you might like to raise the quantization to 24 bits to avoid introducing distortion to the signal when truncating it to integer format.

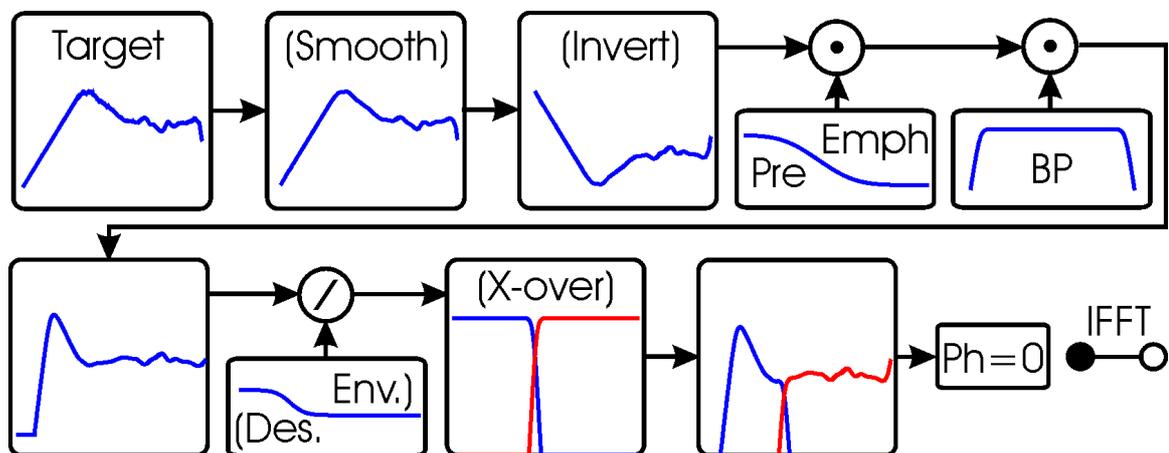
### 1.6.2 Section “Spectral processing”

The process of generating an excitation signal starts with the definition of its desired spectrum. This *target* spectrum may be taken from a normal SPK-file or from a three-staged parametric equalizer. A combination of both (that is, multiplying the target file with the pre-emphasis filter built by the PEQ) is also possible and very useful for loudspeaker equalization.

To realize a simple pre-emphasis at low frequencies, only one stage of the PEQ bank configured as low shelf filter is needed. The two remaining stages are spare. Should you like to create a very complex emphasis curve that cannot be accomplished with the 3-band filter, simply use J-Filter (Alt Edit/J-Filter) to create whatever response you desire, store it to disk and select it as the target file in the BEX menu.

As a first step, the target spectrum may optionally be smoothed. The smoothing parameters may be selected in the submenu evoked by the keystroke ‘T’. A logarithmic smoothing with 1/6 or 1/3 octave will prove useful in moist cases. A preferable alternative way of smoothing is to apply a window after inverse Fourier transforming the excitation spectrum. We’ll come to this in a few paragraphs.

After smoothing, the target spectrum can be inverted. This mode is used when the target file is a loudspeaker response that shall be equalized by the excitation signal.



**Figure 6: Spectral pre-processing of target response**

Following the inversion, the pre-emphasis adjusted by the 3-band parametric EQ accessible in the “Design” submenu is multiplied. If the “base” doesn’t specify a target file, the pre-emphasis curve alone will be injected here. The pre-emphasis may also be constructed conveniently with an interactive mode (“V” in the “design” submenu). For this mode, it is possible to merely visualize the pre-emphasis curve or the product target \*

preamp. In either mode, the following spectral processing stages (x-over, J-Filter) are already included to estimate the final result.

The excitation spectrum can now be confined by means of a band-pass whose parameters can be adjusted in the submenu accessible by a “J” stroke. The use of this band-pass is mandatory when the target is formed by an inversed loudspeaker spectrum to avoid excessive boost at lowest frequencies. To find a reasonable setting by visual control, again it is convenient to use the interactive mode (“V” in the J-Filter submenu).

In the “desired envelope” sweep mode, the processed target spectrum will now be divided by the spectrum defining the envelope. The expansion for this will follow in 1.6.5.

At this point, the signal may optionally be splitted for use with an actively controlled two-way speaker. If the target spectrum has 2 channels, channel 0 will be multiplied with a low-pass and channel 1 with the corresponding high pass. If the target was monaural, it will first be duplicated into channel 1 and then be treated with the LP/HP pair. Crossover frequency and slope may be edited in the “Select” submenu. An interactive mode to find the most suitable crossover point is also available. Again, the other signal processing stages (smoothing, inversion, pre-emphasis, J-Filter) will be included to see the overall result.

If you want to create a sweep and do use the x-over function with a 2 channel target file, bear in mind that not only the amplitude-, but also either phase- and group delay relationships between channel 0 and 1 are important. The 2-channel sweep to be created will have one channel delayed (by group delay manipulation in the frequency domain) to guarantee the same arrival time and phase for the sound originating from both speakers at the crossover frequency. This way, a drop of pressure is avoided when the sweeps transits the crossover region. For this reason, the two frequency responses comprising the target file should be measured at the same point (maintaining speaker and microphone at their locations).

### **1.6.3 Section “Impulse response processing”**

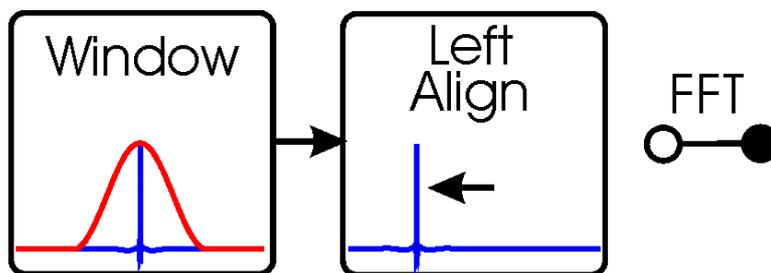
After completing the pre-processing steps in the spectral domain, the desired excitation spectrum will be transformed into an impulse response via IFFT. This is done primarily to make windowing possible. This windowing is equivalent to smoothing the spectrum. Avoiding sharp bents in the target spectrum helps a lot in creating sweep signals with minimum overshoot and crest factor.

As spectrum-smoothing is still handled quite primitively by Monkey Forest (by simply creating the average under a sliding rectangular window), applying a window in the time

domain should be preferred over smoothing the spectrum. But in contrast to the default logarithmic smoothing of the spectrum, windowing affects spectral resolution in a linear fashion. That means that unwanted removal of spectral details will at first occur at the low end of the sweep range, while it almost never will be a problem at high frequencies.

For Impulse- and MLS measurements, there's another problem: The DUT's impulse response retrieved before compensating the pre-emphasis will be broadened by the pre-emphasis impulse response. Hence, confining this impulse response to a period as small as possible makes it easier to separate unwanted reflections when using the time-saving "pre-comp" windowing mode. On the other hand, using a window too small in width will attenuate low frequencies too much. This effect becomes especially noticeable when using the "Zero DC" feature in the window menu. This tool is used to cancel the DC component in the excitation signal (which makes a lot of sense in acoustical measurements). Between a lower cut off frequency roughly equaling  $1/\text{width}$  and DC, a constant decrease of signal energy will be forced by the "Zero DC" mode, even if not indicated by the target.

For sweep- and noise excitation signals, the phase will *always* be deleted to establish a symmetrical impulse response that can be windowed most easily. The proceeding for MLS and impulses is slightly different. There's a selection for the phase to be minimal, linear, maximal or unchanged. In the first three cases, the phase will also be deleted before performing the first IFFT and windowing. A Hilbert transform will later be executed if either "minimal" or "maximal" have been selected (in case of "maximal", the resulting excitation impulse will finally be time-inverted in the borders imposed by the window length). When "unchanged" is selected, a peak search will be performed after IFFT-ing to position the middle of the window at the peak position.



**Figure 7: Windowing and left-aligning the excitation spectrum's impulse response**

After windowing, the impulse will be shifted so that the first non-zero value becomes the first sample of the time signal.

From this point, the signal processing splits, depending on the kind of excitation signal selected. Should this be an impulse – pronto, nothing more needs to be done except storing the excitation signal to disk.

An MLS will be created performing an inverse Hadamard transform (encompassing the steps: time inverting – FHT – time inverting again).

For noise signals and sweeps, the windowed impulse response will be transformed back into the spectral domain.

A  $2^n$ -noise signal (multi sine signal called by some people) is then simply created by randomizing the phases and back transforming to the time domain via IFFT.

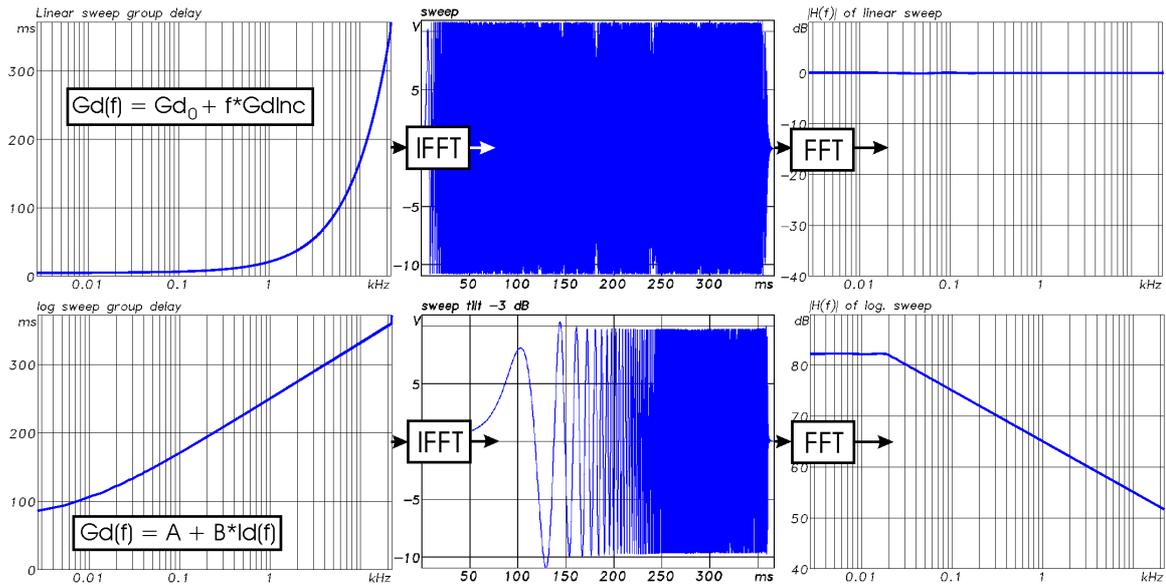
#### 1.6.4 Section “Sweep parameters”

The process for creating a sweep is a bit more sophisticated. The central point is the synthesis of the group delay (which is very well defined for a sweep, showing at which time which frequency occurs). The group delay for the first frequency bin of the excitation spectrum is set to the value given in the field “start margin”. The last frequency ( $f_s/2$ ) will have a group delay attached that equals to the FFT block length duration minus the value given in “stop margin”. It is important to neither let the sweep start at 0 sec nor let it stop at the last sample of the signal period, as the waveforms at the beginning and at the end will be smeared into each other (sort of aliasing, remember that the FFT calculation is only valid for periodic signals, meaning that start and end in reality are the same point).

Which course the group delay will then take within these two extreme frequencies depends on the “increment” selected.

With a “**linear**” increment, the group delay increases with a constant frequency/time rate. If you want to obtain a time signal waveform with constant envelope, the magnitude spectrum has to be chosen to be white (horizontal line in the spectrum, use the mode “premp”, choose “tilt” = 0 dB and set all PEQ gains to 0 dB in the “design” menu).

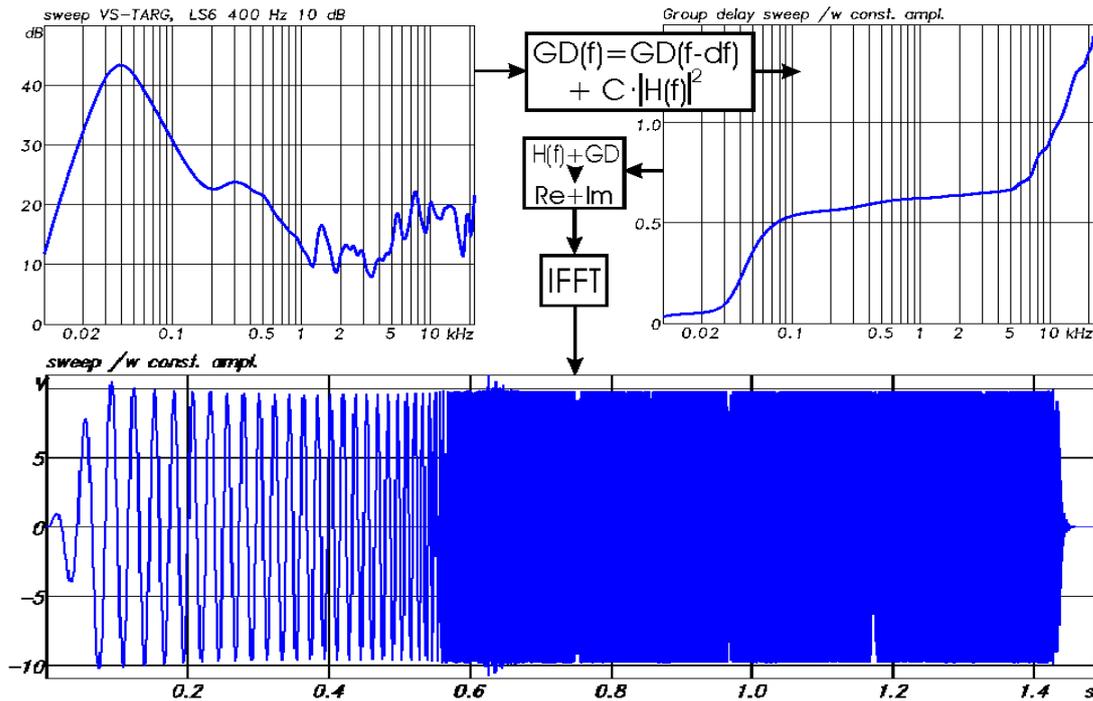
With a “**logarithmic**” increment, the group delay increases by a constant time/octave. That is, the relation  $f_2/f_1 / (GD_2 - GD_1)$  will be kept constant. If again, you’d like to achieve a constant envelope waveform, the magnitude spectrum has to be pink (meaning a magnitude course decreasing with  $-3$  dB/octave. Use the “tilt” feature with a slope of  $-3$  dB in the “design” sub menu).



**Figure 8: Creation of linear and logarithmic sweeps by formulation of group delay. The FFTs on the right side are performed to show the good agreement with the wanted magnitude response.**

Much more interesting and practical is the “**constant envelope**” setting. In this case, the group delay course will be dependent on the target spectrum’s magnitude distribution. The frequency-by-frequency increase of the group delay is made proportional to the target spectrum’s squared magnitude. This simple relationship engenders the coveted almost constant envelope of the corresponding sweep. Its frequency-dependent power now is not controlled by the amplitude of the momentary frequency, but by the speed with which the sine wave sweeps through. Where a high magnitude is prescribed by the target spectrum, the sweep slows down, allowing for more power to be fed to the DUT in this frequency region. The group delay increases fast at this point, meaning that a lot of time will pass by for merely a small increment of frequency.

To hit exactly the desired group delay (as imposed by the “stop margin”) at  $f_{\sqrt{2}}$ , a constant proportional to the target spectrum total energy (“C” in Figure 9) comes into play for the calculation of the group delay increment.



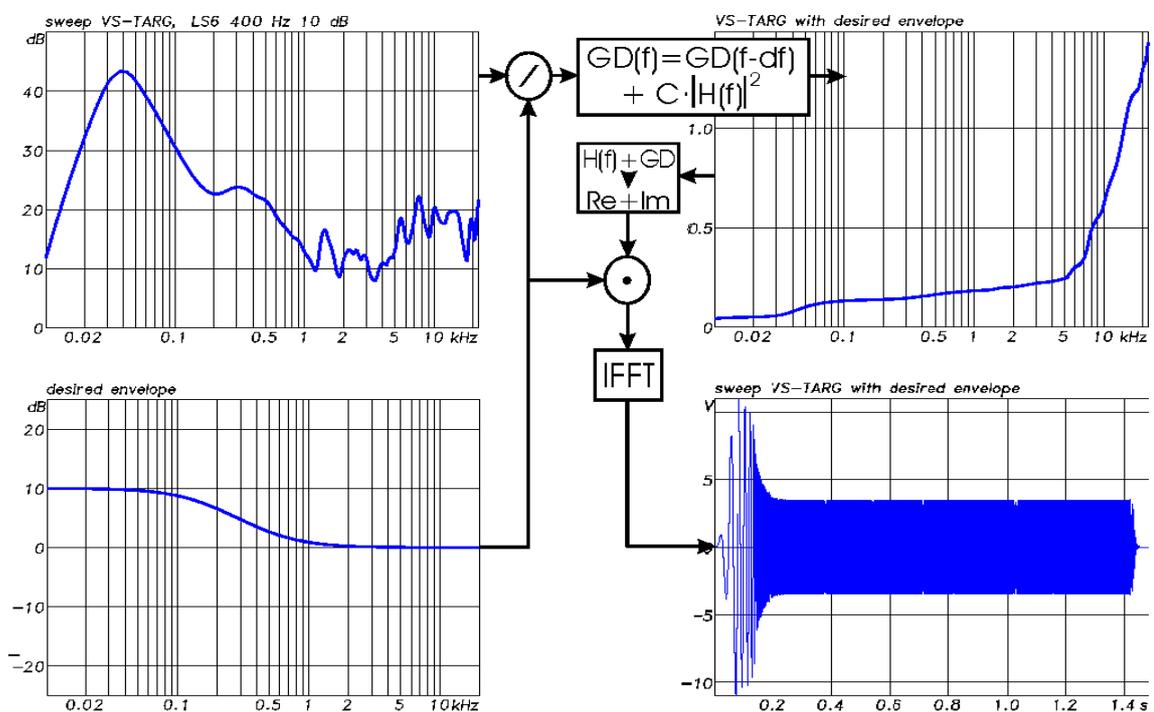
**Figure 9: Sweep creation with constant envelope and arbitrary magnitude response.**

After constructing the group delay curve, it will be transformed into a phase spectrum by integration. The phase/magnitude pair is then converted into a real/imaginary part spectrum that is finally inverse-Fourier-transformed to yield the fashionable sweep. This one will now be refined by fade in/out to reduce spurious effects. For this purpose, a left-half window between 0 sec and half of the “start margin” is positioned. To do the fade-out, a right-half window’s end will be located at exactly the *full FFT Block length* – “*stop margin*”. Its size is fixed to 2% of the sweep length. A small, but mostly negligible dip of magnitude at the high end of the sweep range can be observed due to this processing.

The use of the “constant envelope” mode is recommended for the bulk of measurement purposes. It produces excitation signals that can be ejected with high power due to their low crest factor of little more than 3 dB. Whenever a fixed electrical clipping level or amplifier power is the restricting maximum rating, this is the best excitation signal.

There are, however, cases in which the constant envelope of the sweep is still not optimal. A common example is a multi-way loudspeaker used for acoustical measurements. As already mentioned in 1.2, the tweeter generally will be the most vulnerable chassis. Not rarely, the woofer can handle 10 or more times the power of the tweeter. So, in order not to endanger the tweeter, a constant envelope sweep had to be fed with no more power than the latter can endure, leaving the woofer’s capabilities largely unexploited. This is annoying, as precisely in room acoustical measurements, most of the power is needed in the lower octaves for reasons already mentioned.

So, without doubt, it sometimes would be nice to control the sweep's amplitude in a frequency dependent fashion by means of a self-defined "desired envelope" spectrum. And this can be accomplished by only a minor modification of the sweep creation process. As depicted in Figure 10, the trick is to first *divide* the target spectrum by the "desired envelope" spectrum (whose phase is forced to  $0^\circ$  first). The resulting spectrum is the base for the group delay synthesis, using the same formula as in the "constant envelope" case. After creating the real/imaginary pair, the sweep spectrum will now be *multiplied* with the desired envelope spectrum, re-establishing the old magnitude response. The IFFT will now yield a sweep that faithfully features the frequency dependent amplitude course imposed by the "desired envelope" spectrum.



**Figure 10: Sweep creation with desired envelope and arbitrary magnitude response.**

Two degrees of freedom are offered here: An arbitrary spectral distribution along with an arbitrary definition of the frequency dependent envelope is transformed into a swept sine wave warped in amplitude and time.

It must be admitted that this approach doesn't remedy entirely the danger of overheating sensitive tweeters. Comparing a "desired envelope" to a "constant envelope" sweep with the same energy contents, it is obvious that the energy part going out to the tweeter is only stretched in time now, resulting in the same heat-up if the sweep length is smaller than the time constant of the voice coil's thermal capacity. Room acoustical measurements with long sweeps, however, benefit from this expanding. The "desired envelope" sweeps are also advantageous if the tweeter is not only threatened by

overheating, but mechanically caused damage (for example, excessive forces in compression drivers) as well. The reduced distortion due to the lesser level may also be motivating, although it is the big plus of sweep measurements that the distortion products can be separated so well from the actual impulse response.

For both the “constant envelope” and “desired envelope” case, you may freely adjust all associated parameters to create whatever type of absurd excitation sweep. However, too much freedom may lead to frustration. It is easy to misadjust the settings in a way that only rubbish occurs. Once this happens and you don’t know exactly which parameters you messed up, simply reselect the default or stored setting from which you began the process and start all over again. Frequently store setups (an easy task) as long you are satisfied with intermediate results. Following are a few hints to produce usable signals.

First of all, carefully watch out that the target spectrum is in reasonable shape, having no unwanted peaks inside the desired frequency range and no significant energy concentration outside. Use pre-emphasis and the band-pass with caution, taking advantage of the interactive mode. Some smoothing should always be applied to the target to avoid overshoot of the resulting sweep, preferably activating windowing under the “impulse processing”. You have to compromise between frequency resolution and non-overshoot of the sweep. If the sweep response at low frequencies is remote of what you wanted, either the window length, the “start margin” or the signal length defined by “degree” might be too small.

As already mentioned, start- and stop group delay should be away from the edges to avoid dirt effects at the beginning and the end. If the onset of the sweep shoots over too much, try increasing the “start margin” value. In contrast, if the sweep’s waveform evolves a bit late, reduce it.

### **1.6.5 Section “Reference file generation”**

The usual way to generate a reference file after having built a new excitation signal is by the well known “reference measurement”, available as submenu in Monkey Forest’s measurement menus. It entails connecting input and output, ejecting the excitation signal and creating the inverse of the FFT of the captured response. While this is also the most precise way, as it will exclude any frequency response imperfections of the measuring system itself, it is sometimes not feasible. This encompasses all cases in which the signal path to be equalized cannot be characterized by a single measurement. The equalization of a loudspeaker’s acoustical power output (described in section 1.3) with the requirement to average many measurements is a common example.

Sometimes, a reference measurement is possible, but the required precision is not very high. In these cases, it is simply convenient to use the automatic reference file generation presented here.

Monkey Forest uses three different kinds of reference files: one for the general measurement of frequency response (file suffix .SPF), one for loudspeaker sensitivity (file suffix .SPL) and one for impedance (file suffix .SPI). All three can be generated at a time, but of course it's disk space saving to create only the type you need.

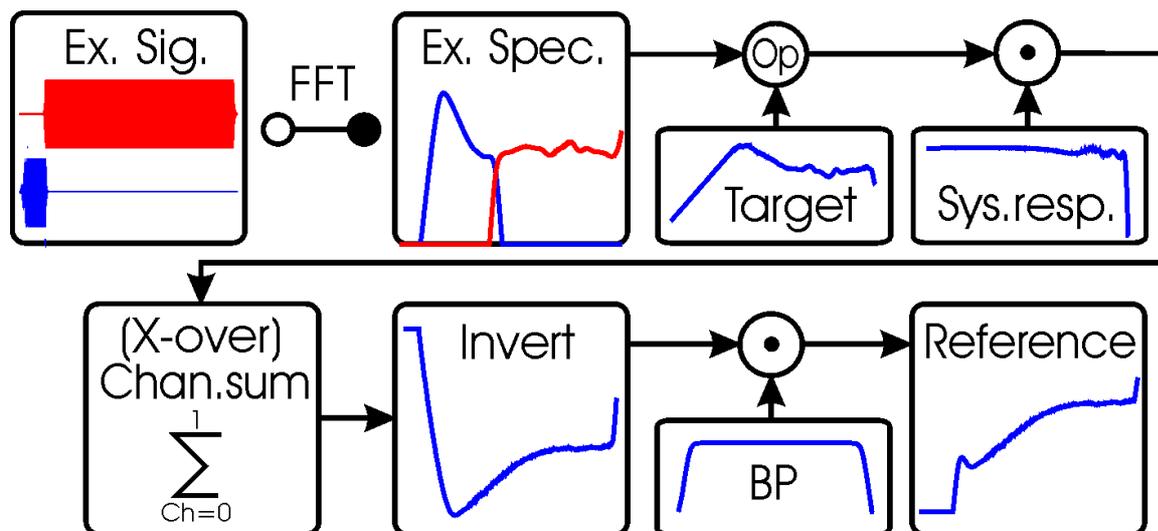
The reference calculation can be done in three different ways.

The most common is to simply invert the “**ExSig**”. When combining with the “**EQ** *<hardware>*” setting, this is a direct substitution for the reference file generation by self-measurement.

For room acoustical measurements in which the response of the source loudspeaker (given by the “target” file) shall be equalized, principally only the possible additional pre-emphasis (as defined in the “Design” menu) must be equalized in the captured room transfer function. In this case, “**preemp**” should be selected. The reference file's magnitude will then be the inverse of the pre-emphasis curve (optionally filtered by “K” Filter). This mode is suitable for general room acoustical measurements.

A third choice is “**Excitation signal · target spectrum**” if “invert target” is set to yes (or “**Excitation signal ÷ target spectrum**” if not). This will yield a reference file with a course similar to the “preemp” case, but with an additional correction for the softening caused by smoothing and windowing during the construction of the excitation signal.

Figure 11 shows how this works. The excitation signal is first transformed into the spectral domain. There, it is *multiplied* with the target spectrum if the “invert target” switch in the “spectral processing” section is on (this is the normal case for loudspeaker equalization). If the switch is not on, the excitation spectrum is *divided* by the target file.



**Figure 11: Construction of reference file for a 2-channel pre-emphasized sweep in “target signal ÷ target spectrum” mode**

In case of a loudspeaker equalization (“invert target” = yes), the result of this operation is the simulated spectrum of the excitation signal emitted by the loudspeaker and captured by the microphone. If a stereo excitation signal has been constructed, the spectrum will also have 2 channels here. The “target” would normally contain two channels as well, namely the sensitivity measurement of the two ways of the loudspeaker to be equalized. In Figure 11 however, the target used is monaural to avoid too much confusing traces.

The measurement system’s frequency response may be included into the reference file in two ways: By either using a file with a “**measured response**”, or by “**simulation**” of the measurement’s frequency response. But this latter mode is only available with the ITADDA16 AD/DA board (for which it’s worth the effort to equalize the relatively strong undulations caused by the primitive digital anti-aliasing filters).

If a stereo excitation signal has been created (as in the example of Figure 11), the two channels will now be added. The result will be the frequency response of the equalized loudspeaker as it would be captured by the microphone and passed through the AD converter. It should contain more or less only the additional part of pre-emphasis used to deviate from linear response (as defined in the “Design” menu). However, there will be some aberration due to the target spectral processing and windowing selected. Particularly, the roll-off imposed by the band pass will be evident.

This spectrum is inverted to form the reference spectrum. Applied now in subsequent measurements, it should yield an almost perfect straight magnitude response with zero phase distortion, provided there is enough dynamic range. This requirement, of course, is violated outside the excitation signal’s frequency range where the reference spectrum will

only boost noise. So it is indispensable to confine the reference spectrum to the same frequency range as the excitation spectrum.

So the final step of the generation process for the simulated reference spectrum is the application of a band pass to attenuate frequencies outside the excitation signal's range. In the measurements, this will inhibit unwanted noise especially in the subsonic region. When confining the excitation signal's frequency range in the "**J-Filter**" menu, the values chosen there are automatically copied to the "**K-Filter**" used to filter the reference spectrum (yes, you got it, "K" has only been chosen because it was one of the last free keys in this crammed menu). The filter order is then increased by 2 in order to achieve even higher rejection of unwanted out-of-band noise. You can override the settings in the "K-Filter" menu, but have to do so again every time you adjusted something in the "J-Filter" menu, due to the automatic copy process.

Select the "**Excitation signal ÷ target spectrum**" mode if you want to see a straight line in the spectrum when measuring the speaker at the same position where you previously took the "target" file (Example: Microphone calibration). Also select it when the phase and group delay distortion of the speaker shall be removed by the reference file.

### 1.6.6 Overview of “Build Excitation Signal” menu functions

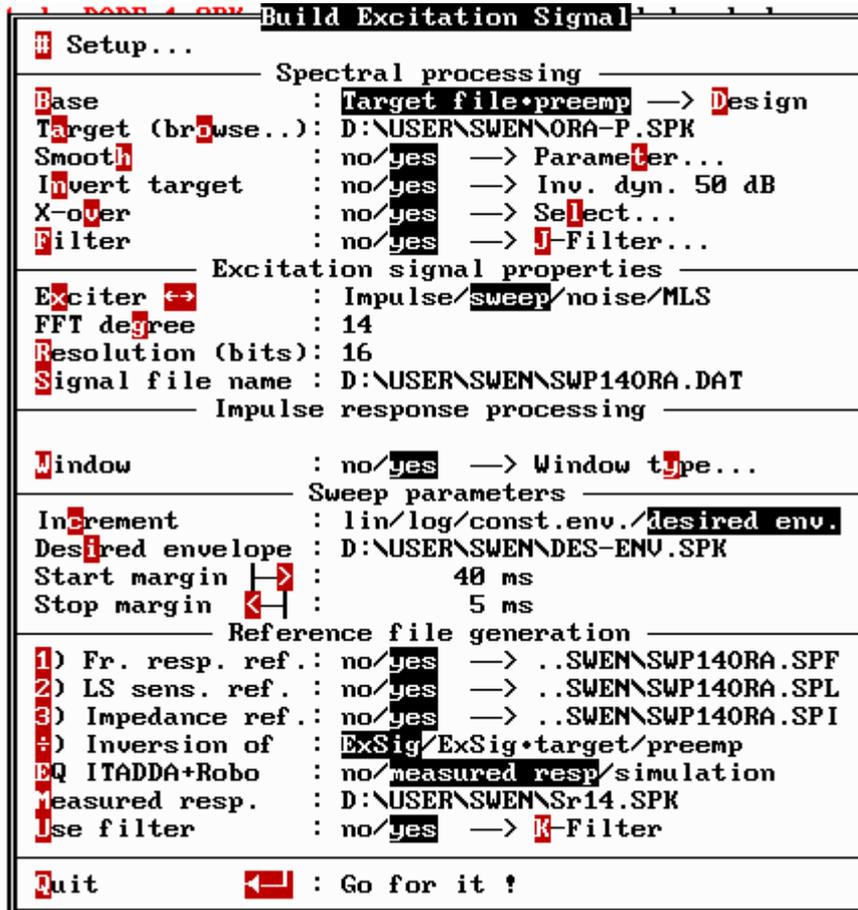


Figure 12: Screenshot from the Excitation Signal with most of the options set for maximum confusion

# Setup...		Enters into a submenu where you can load and store configurations or select a default setup
Spectral Processing		
Base $\gamma$ design		<p>Selects where the amplitude spectrum of the excitation originates from. There are 3 possibilities:</p> <ol style="list-style-type: none"> <li>1) Spectrum originates from a “<b>target file</b>”.</li> <li>2) A “<b>preemp</b>” spectrum will be created by means of a 3-band parametric EQ.</li> <li>3) A combination of both (“<b>Target file • preemp</b>”) is used. In the latter two cases, pressing D (“<b>Design</b>”) will guide you into a submenu where you can adjust the 3 EQ stages to adjust the pre-emphasis. Usually, just one low shelf stage set to the desired pre-emphasis gain will do the job.</li> </ol>
Target (browse...)		If a target file is in play, enter its name here. You may also

		search for the file using the browser function.
Smooth $\gamma$ Parameter...		If you would like to smooth the target spectrum, select the parameters in the submenu evoked by pressing “T”.
Invert target		This switch allows inverting the selected target file (an additional emphasis, if selected, will not be inverted).
X-Over $\gamma$ Select		Allows splitting the excitation signal into two ranges in order to feed a two-way loudspeaker. The crossover frequency and slope may be selected in the submenu. The target file, if involved, may be mono or stereo. In the stereo case, channel 0 will be treated with the low-pass and channel 1 with the high-pass.
Filter $\gamma$ J-Filter		Allows confining the spectral contents of the target spectrum by means of a band pass. Useful in acoustic measurements in order to alleviate the loudspeaker from frequency components far outside its transmission range
Excitation Signal Properties		
Exciter		Choose here which kind of excitation signal you are bound to create: Impulse, noise with arbitrary phase, sweep with controlled group delay, or MLS

<u>D</u> egree		<p>Corresponds to the length of the excitation signal. It will contain <math>2^{\text{Degree}}</math> samples when the excitation signal shall be an impulse or noise, <math>2^{\text{Degree}}-1</math> samples when an MLS is selected, and <math>2^{\text{Degree}}</math> samples minus a value corresponding to the “stop margin” in case of a sweep.</p> <p>As a rule of thumb, the length should be 2-4 times the expected reverberation time to facilitate the isolation of intermodulation products from the impulse response (especially rewarding in sweep measurements).</p>
<u>R</u> esolution (bits)		The quantization with which the excitation will be stored on disk.
<u>S</u> ignal file name		<p>Composed of three characters characterizing the kind of stimulus (IMP, NOI, SWP, MLS) plus two digits representing the degree of the signal.</p> <p>The remaining three characters are copied from the beginning of the target file name, if one is involved. If not, they are set to “NEW”. Edit only these three last characters to maintain the nomenclature.</p>
Impulse response processing		
<u>P</u> hase		<p>Appears only in Impulse or MLS mode. The phase of the target function may be erased (<b>lin</b>), set to minimum by Hilbert transform (<b>min</b>), set to maximum by Hilbert transform and time inverting inside the window range (<b>max</b>), or simply left <b>unchanged</b>.</p> <p>In Sweep and Noise mode, the phase will always be deleted to force the impulse response to be symmetrical, which relieves windowing.</p>
<u>W</u> indow (type...)		<p>The impulse response of the target function obtained by IFFT can and should be treated with a window here. This will provide a smoother excitation signal spectrum, which will exhibit minimum overshoot of the corresponding time signal. The length of the window should be selected considering the desired selectivity of the excitation spectrum. For the audio range, 50-100 ms are a good choice. If you want the excitation signal to be <b>DC-free</b>, select the appropriate switch in the window submenu.</p>

Sweep parameters		
Increment		<p>This switch is used to control the group delay of the sweep. There are four possibilities:</p> <p>”<b>lin</b>“ will create a sweep whose frequency increments at a constant rate [Hz/s].</p> <p>”<b>log</b>“ will generate a sweep whose frequency doubles at constant speed.</p> <p>”<b>constant envelope</b>“ will produce a sweep with constant amplitude in time (crest factor <math>\approx 3</math> dB), but with the desired spectral distribution as imposed by the target.</p> <p>”<b>desired envelope</b>“ allows to construct a sweep whose time signal amplitude is defined by the “desired envelope” spectrum for each momentary frequency.</p>
Desired envelope		Name of the spectrum file containing the desired amplitude for each frequency. The phase of this spectrum will be disregarded.
Start margin $\gamma$		The group delay for the lowest frequency contained in the excitation signal should not start at 0s, as this would cause overlapping of the low frequency components to the end of the sweep. A higher value should be selected in order to ease a soft fade-in by a left window half. The default values set by choosing a default setup are generally a good compromise. If you experience a corrugated magnitude structure at low frequencies of the excitation signal’s spectrum, try to increase this value. If the sweep evolves late, decrease it.
Stop margin $ \phi$		<p>The sweep should not be stretched out over the full length of the FFT block size, as this would deter capturing the full decay of the DUT’s response. The value to be entered here tells how much shorter than the full FFT block length the excitation signal will be.</p> <p>In case of a room acoustical measurement, this value should be set to at least:</p> $\text{expected reverberation time} \cdot \text{wanted dynamic range} / 60$ <p>For loudspeaker measurements, set the value to at least:</p> $LS\text{-Mic distance}/340m + \text{expected group delay maximum} + 5\text{ ms}$
Reference file generation		
1) Fr. resp. ref.		Set to yes if you want to generate a reference file for frequency response measurements

2) Ls. sens. Ref.		Set to yes if you want to generate a reference file for loudspeaker sensitivity measurements
3) Impedance ref.		Set to yes if you want to generate a reference file for impedance measurements
÷ Inversion of		<p>For normal measurements, select “<b>ExSig</b>”, meaning that the reference file is derived by simply inverting the excitation signal’s spectrum.</p> <p>Select “<b>preemp</b>” for the measurement of RIAs (room impulse responses). Then, the pre-emphasis magnitude and excitation signal phase will be removed from the measurements.</p> <p>The third choice is “<b>Excitation signal · target spectrum</b>” (if “invert target” is set to yes), or “<b>Excitation signal ÷ target spectrum</b>” (if not). It will remove pre-emphasis, imperfections caused in the construction process and phase distortion of the speaker (as manifested in the target file) from the measurements.</p>
<u>EQ hardware</u>		<p>This switch enables to take into account the self-response of the AD/DA hardware being used. This can be accomplished by connecting in- and output of the measuring system and measuring the transfer function of this link. The resulting self-response file can be used in all subsequent BEX sessions.</p> <p>If the ITADDA16 board is selected as the AD/DA hardware, the undulating response of the internal digital anti-aliasing filters may be emulated by setting this switch to “<b>simulation</b>”. This will enable the use of a sampled version of the Crystal Codec’s resulting impulse response. Note however that this is only a course approach to equalize the inherent self-response of your AD/DA hardware.</p>
Measured resp.		Enter the name of the measured self-response. It does not necessarily have to have the same size as the reference file, as an interpolation routine will adapt the spectrum to the required frequency resolution.
<u>Use filter γ</u> <u>K</u> -Filter		<p>It is highly recommended to confine the reference spectrum to the same range as the excitation spectrum.</p> <p>When changing settings in the “J” Filter menu, the values are copied into the “K”-Filter menu. The filter order is then incremented by 2. Normally, you can accept these settings, but you’re also free to change them.</p>

**Table 1: Overview of functions in the „Build Excitation Signal“ menu**