

Chapter 7: Impulse Response Measurement Basics

1: What is an Impulse Response?

In the most basic terms, an impulse response (IR) can be defined as the time domain (time vs amplitude) response of a system under test (SUT) to an impulsive stimulus. The word “system” in this case could mean something as small as a microphone or a single transducer, something as simple as a single filter on an equalizer. Or, it might mean something as big as a concert hall or sports arena, as complicated an entire sound system or a combination of the two. Smart users of course are most often concerned with sound systems and their acoustical environments.

In the context of acoustical analysis, you might also think of an impulse response as the acoustical “signature” of a system. The IR contains a wealth of information about an acoustical system including arrival times and frequency content of direct sound and discrete reflections, reverberant decay characteristics, signal-to-noise ratio and clues to its ability to reproduce intelligible human speech, even its overall frequency response. The impulse response of a system and its frequency-domain transfer function turn out to be each other’s forward and inverse Fourier transforms.

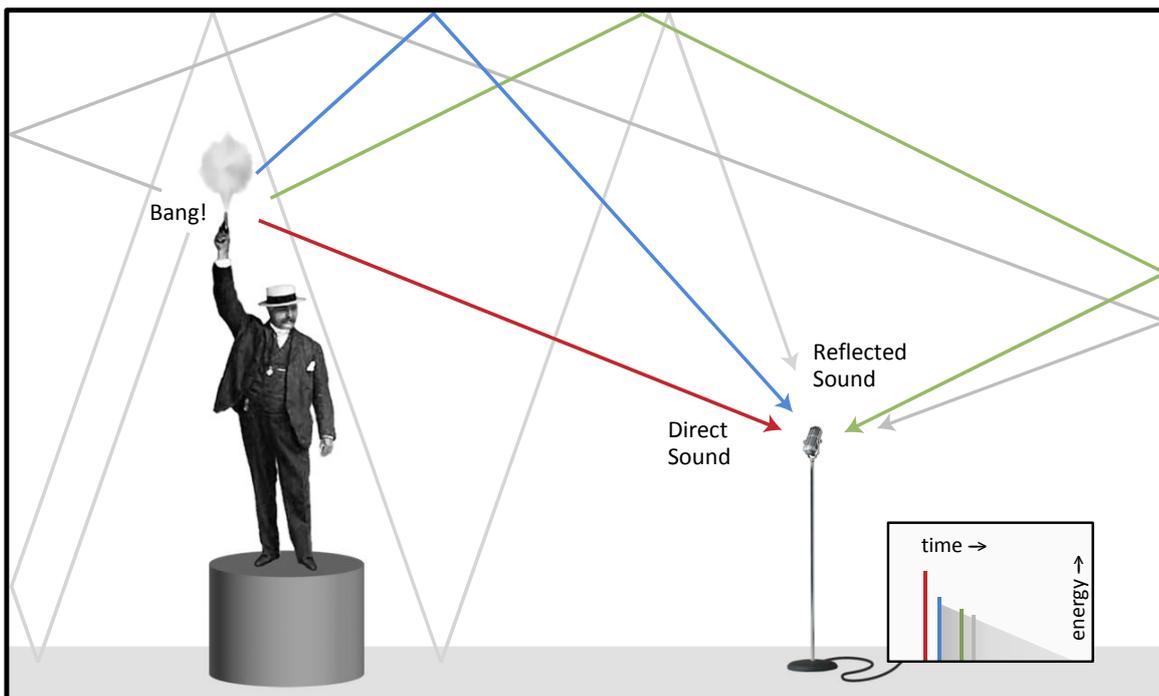


Figure 115: Conceptual illustration of an acoustical impulse. Sound from an excitation source arriving at a measurement position by multiple pathways, both direct and reflected. Here we see the path of direct sound from the source to the microphone in red, followed by a first order reflection in blue, a second order reflection in green, and higher order reflections in gray. Later arrivals tend to pile on top of each other forming a decay slope.

An acoustical impulse response is created by sound radiating outward from an excitation source and bouncing around the room. Sound traveling by the most direct path (a straight line from the source to a measurement position) arrives first and is expected to be the loudest. Reflected sound arrives later by a

multitude of paths, losing energy to air and surface absorption along the way, so that later arrivals tend to come in at lower and lower levels. In theory this process goes on forever. In practice, the part we care about happens within a few seconds – perhaps less than a second in smaller rooms and/or spaces that have been acoustically treated to reduce their reverberation times.

The arrival of direct sound and probably some of the earliest arriving reflections will be clearly distinguishable on a time-domain graph of the impulse response. As reflected copies of the original sound keep arriving later and later, at lower and lower amplitude levels, they start to run together and form an exponential decay slope that typically looks like something close to a straight line when displayed on a graph with a logarithmic amplitude scale.

Anatomy of an Acoustical Impulse Response

Although no two non-identical rooms ever have identical impulse responses, there are a few component features that we can identify in some combination in almost any acoustical impulse response. These include the arrival of direct sound, early reflections, reverberant build-up and decay, and the noise floor. Figure 116 shows an acoustical impulse with its component parts labeled. Descriptions for each follow.

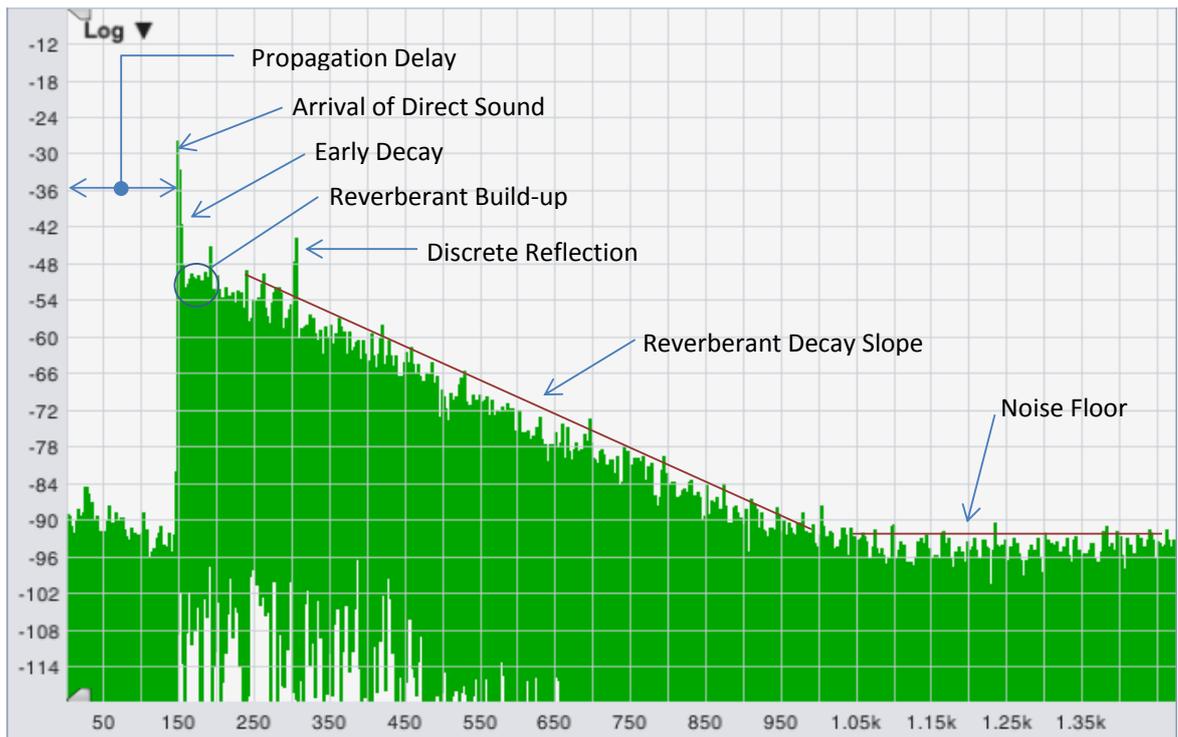


Figure 116: An acoustical impulse response with its common component parts labeled. This is a semi-log time domain chart with time in milliseconds on the x axis and magnitude in decibels on the y axis.

Propagation Delay

The time that it takes for direct sound from the sound source to reach the measurement position is the propagation delay time. This may include throughput delay for any DSP processors in the signal chain in addition to the time that it takes for sound to travel through the air.

Arrival of Direct Sound

Since the shortest distance between two points is always the straightest line, the first thing we expect to see when looking at an impulse response (IR) is the arrival of direct sound from whatever sound source we're using to stimulate the system under test. Depending on what we're trying to learn, the source could be an installed sound system, an omnidirectional loudspeaker brought in specifically for measurement purposes, a balloon pop or a shot from a blank pistol, or in a pinch, maybe hand claps or someone slamming a case lid shut.

In most cases, we would also expect the first arrival to be the loudest and correspond to the highest peak we can see in the IR, and in most cases we'd be right. There can be occasional circumstances where that might not turn out to be strictly true but in the vast majority of the cases, it should.

Discrete Reflections

After the arrival of direct sound, the next most prominent features we tend to see are sound arriving by the next most direct paths; the lowest order reflections. Sound that bounces off one surface to get from the excitation source to a measurement position is called a first-order reflection, two bounces gives you a second order reflection and so on. Reflected sound can be useful or detrimental, depending on factors such as its relative magnitude and timing in relation to the direct sound and the extent to which it is clearly distinguishable from the diffuse reverberant sound.

Early Decay, Reverberant Build-up, and Reverberant Decay

Following the arrival of direct sound and the lowest order reflections, sound in a reverberant space will continue bouncing around a room for a while, creating higher and higher order reflections. At any given listening position, some of this reflected energy will combine constructively over a relatively short period of time, resulting in a build-up of reverberant sound, before air loss and absorption by the materials that make up reflecting surfaces begins to take their toll. At that point, the reverberant decay phase begins.

In practice, you may or may not be able to see the reverberant build-up in an impulse response as distinct from the direct sound and early reflections. Sometimes it can be quite clearly visible, other times not so much. By convention, the first 10 dB of decay after the arrival of direct sound in the reverse-time integrated IR (we will get to that in *Chapter 9: Analyzing Impulse Response Data*) is considered to be early decay. Reverberant decay is conventionally measured over a range from 5 dB below the level of direct sound down to a point 30 dB below that on the reverse integrated IR, or 20 dB down in a pinch.

Noise Floor

In theory, the reverberant decay phase of the IR continues forever, as an ideally exponential curve that never quite reaches zero. In practice it reaches a point relatively quickly where we can no longer distinguish it from the noise floor of the measurement. Noise in an IR measurement can come from

several sources, including ambient acoustical noise, electrical noise in the SUT and the measurement system, quantization noise from digitizing the signal(s) for analysis, and artifacts from DSP processes used for analysis.

Uses for impulse response measurement data

Delay Time Measurement

The most common use for impulse response measurements in Smaart is in finding delay times for signal alignment in transfer function measurements and for aligning loudspeaker systems. Each time you click the delay locator in Smaart an IR measurement runs in the background. In this case all we really care about is the initial arrival of direct sound, which is typically so prominent that you can pick it out with high confidence even when signal-to-noise ratio of the IR is poor, so we don't even bother displaying the results. Smaart simply scans for the highest peak and assumes that to be the first arrival, and most of the time that works very well.

Occasions where automatic delay measurements might not work well include measurements of low-frequency devices or any case where you're trying to measure a directional full-range system well off axis, in a location where a prominent reflection can dominate the high frequencies. In the latter case, it's possible for reflected HF energy to form a higher peak later than the arrival of direct sound, requiring you to visually inspect the IR data to find the first arrival.

Reflection Analysis

Another common use for IR measurements is in evaluating the impact of problematic discrete reflections. Reflected sounds can be beneficial or detrimental to a listener's perception of sound quality and/or speech intelligibility, depending on a number of factors. These factors include the type of program material being presented (generally, speech or music), the arrival time and overall level of the reflected sound relative to the level of direct sound, and the frequency content and the direction from which they arrive. As a general rule, the later they arrive and the louder they are (relative to direct sound) the more problematic they tend to be.

Reverberation Time (T60, RT60...)

Reverberation time is kind of the grandfather of quantitative acoustical parameters. First proposed by Walter Sabine a century ago, T60 or RT60 reverberation time is the time that it takes for reverberant sound in a room to decay by 60 decibels from an excited state (after the excitation signals stops). It is one of the most widely used (and in some cases perhaps misused) quantities in room acoustics. Although it is quite possible for two rooms with identical reverberation times to sound very different, when evaluated band-by-band it can still give you some idea as to the overall character of the reverberant field in a given room. In concert halls it can give you an idea of perceived warmth and spaciousness for music. In auditoriums, it is often used as a rough predictor of speech intelligibility.

Early Decay Time (EDT)

Early decay time ends up being the decay time for direct sound and earliest, lowest-order reflections. Since the earliest reflections tend to be the most beneficial in terms of separating sounds we want to hear from reverberation and background noise, EDT can give you some clues about overall clarity and

Chapter 7: Impulse Response Measurement Basics

intelligibility in a room and/or system. EDT, like RT60, is conventionally normalized to the time it would take for the system to decay 60 dB at the measured rate of decay.

Early-to-late energy ratios

Early to late energy ratios are a direct measure of the sound energy arriving within some specified interval following the arrival of direct sound, vs the energy in the remaining part of the IR. These provide a more direct method of evaluating the relationship between beneficial direct sound and early reflections that a listener hears versus the amount of (potentially detrimental) reverberation and noise, than inferences made from the early and reverberant decay rates.

Speech Intelligibility Modeling

Early to late energy ratios such as C35 and C50 have long been used as objectively measurable predictors of subjective speech intelligibility. In the 1970s Victor Peutz came up with Articulation Loss of Consonants (ALCons), a predictive metric for intelligibility based on the volume of a room and its reverberation time, the directivity of loudspeakers and distance from source to the listener. Later on, Peutz revised the equation to use a direct-to-reverberant energy ratio in place of volume, distance and loudspeaker Q, making ALCons a directly measurable quantity. More recently, the speech transmission indexes (STI and STIPA) have emerged as metrics that are generally more robust. All of these can be calculated from the impulse response of a system.

Chapter 8: Impulse Response Mode User Interface

If you already know your way around IR mode in Smart 8 you can probably skip this section, but it might not hurt to at least skim over it. If you are new to IR mode then introductions are in order. To get to IR mode in Smart, select *IR Mode* from the *View* menu, press the “I” key on your keyboard or click the Impulse button that appears in the lower right corner of the main window in real-time mode (just below the signal level /sound level meter). You will find yourself confronted with a screen like the one below. (The colors may be darker but the layout is the same.) On the right side of the window is a vertical strip of controls. The rest of the window is devoted to the graph areas and cursor readout.

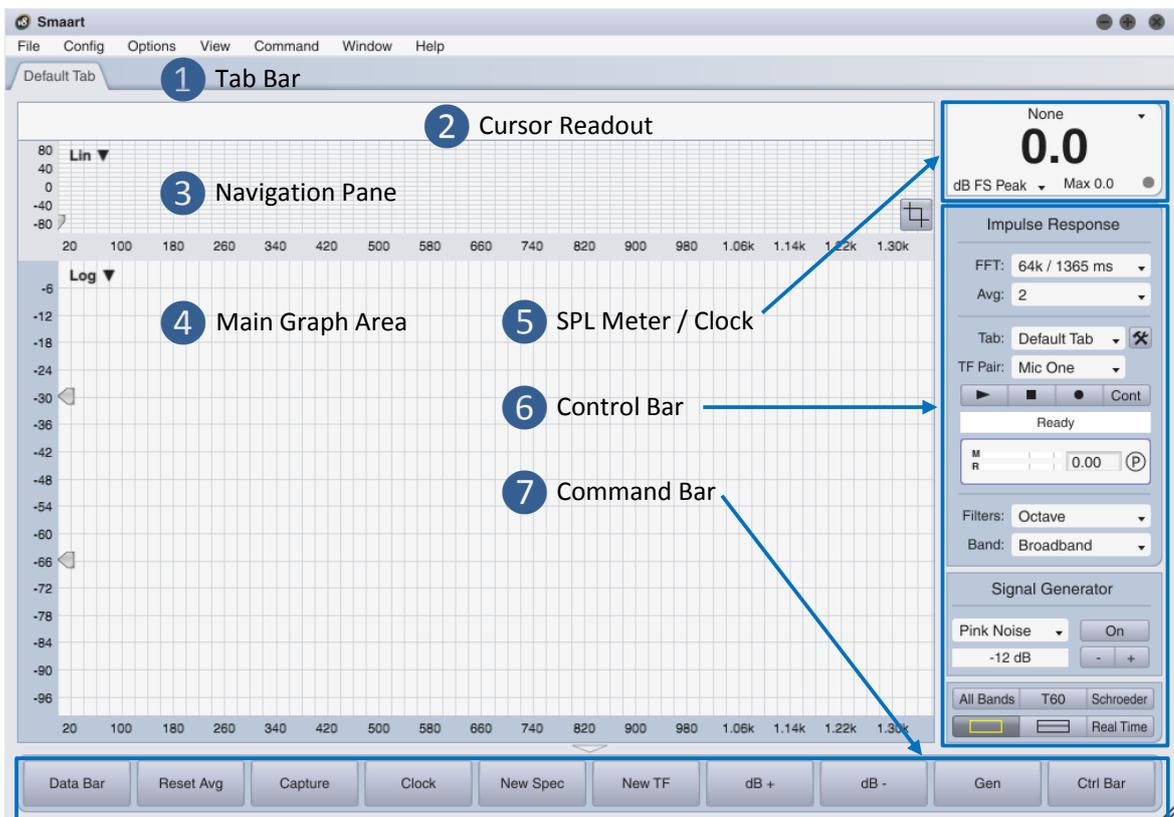
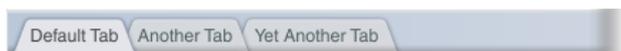


Figure 117: Anatomy of the main Smart window layout in Impulse Response mode

1 Tab Bar



Smart can run in multiple windows and each window can host multiple tabbed workspaces that we refer to simply as tabs. Each tab includes its own measurements, screen layout, and plot assignments and you can switch between them by clicking the tab-shaped buttons below the menu bar in the area we refer to as the Tab Bar. You can move a tab from one Smart window to another by clicking on its button in the Tab Bar with your mouse and dragging it to another window, then releasing the mouse button to drop it.

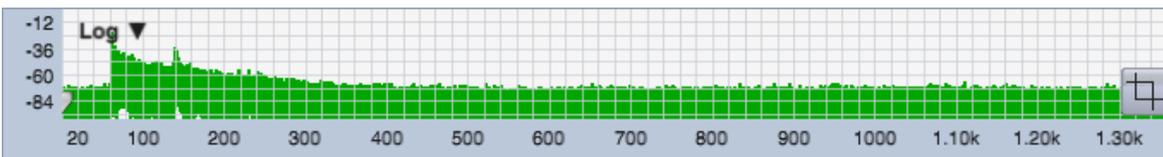
2 Cursor Readout

When measurement data is present, the cursor readout displays numeric coordinates for the cursor location(s) as you move your mouse over the graphs areas. Numeric coordinates are provided here for the cursor location in units of time, amplitude/magnitude and frequency, as applicable to graph type.

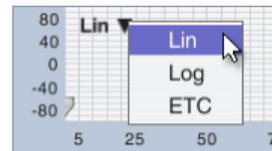


For time domain graphs (Lin, Log or ETC) in the main graph area(s), there are three sets of coordinates as show above. From left to right, they are the location of the locked cursor which typically marks the highest peak in the impulse response, the movable (mouse) cursor coordinates, and in brackets on the right, the difference between the locked and movable cursors. Note that time coordinates can optionally be displayed as both time (in milliseconds) and equivalent distance traveled, based on the currently specified speed of sound.

3 Navigation Pane



The small time-domain display in the upper part of the graph area is used for navigation and is always visible. Right-clicking and dragging (Ctrl + click and drag on Mac) across the graph in this pane selects a specific time range for display on the larger time-domain charts. The full IR time record remains visible in the navigation pane when you are zoomed in (unless you use the crop function). Clicking anywhere in the left margin of the plot clears the zoom range and returns any time-domain graphs in the main display pane(s) to the full IR time record.

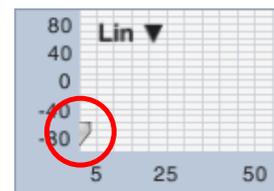


The selector control in the upper right corner of the navigation pane selects the graph type to be displayed in this area. The navigation pane is limited to time-domain graph types only (Lin, Log or ETC).

The *Crop* button in the lower right corner of the navigation pane can crop a file for display purposes to show only the selected time range – a very useful feature when working with IR measurements with long noise tails. Cropping is non-destructive and can be undone – clicking the Crop button again on a cropped measurement restores the full extent of the original time record – however if you save the IR to a file while cropped, the cropped version is written to file.



The little arrow shaped widget positioned in the lower left corner of the navigation pane (circled in red in the screen clip show the right) is the Time 0 marker. When you record a dual-channel impulse response measurement in Smaart, this marker is automatically set to match the reference signal delay time – if you are familiar with real-time transfer function measurement it's analogous to the center point of the Live IR. For single-channel measurements



or file-based data it is set to the beginning of the time record. Dragging it to the left or right moves the time-zero point for all time-domain graphs (Lin, Log or ETC) displayed in the main graph area.

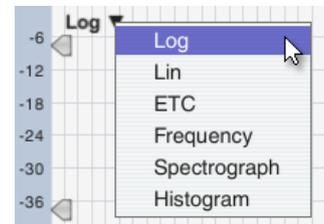
4 Main Graph Area

The larger (lower) portion of the graph area can be divided into one or two panes by clicking the buttons labeled with rectangles in area #5. Each pane can host any one of the six available main graph types:

- *Lin, Log, ETC* (time domain views)
- *Spectrograph* (frequency and level vs time),
- *Frequency* (spectrum)
- *Histogram* (bar chart of quantitative acoustics values by octave or 1/3-octave band.)

We will discuss each of these display types in the next chapter.

The graph type for each pane is selected by means of the drop-list control in the upper left corner of the pane. Some main graph types also have additional selector controls in their upper right corner that control display options specific to that graph. The two arrowhead-shaped widgets to the left of the list of graph types appear on all time-domain and spectrograph plots. They control the dynamic range for the spectrograph.



As you probably noticed, the main graph area and the navigation pane are blank when you first enter IR mode. You won't see any data in the graph areas until you record a measurement or load an impulse response from file (*File* menu > *Load Impulse Response*). Smart can open and analyze .wav and .aiff files containing any type of audio data, but IR mode is purpose-built for analyzing impulse responses. There is no multi-channel file support or optimization for working with files more than a few seconds in length – and of course a lot of the IR analysis capabilities are irrelevant for other types of audio data.

5 SPL Meter / Clock

The large numeric readout that appears (by default) at the top of the Control Bar in the upper right corner of each tab can be configured to function as a Sound Pressure Level (SPL) meter, an integrating Equivalent Sound Level (Leq) meter, a peak signal level meter calibrated to normalized digital full scale, or a clock. When the level meter is displayed, pressing the [K] key on your keyboard switches the display to a clock and vice versa. This display can be hidden if you don't need it by selecting *SPL Meter* from the *View* menu pressing [Alt/Option] + [K] on your keyboard. When hidden, repeating either of these actions will restore it.



The in-tab SPL Meter operates almost identically to the meter module in the *SPL Meters* window. Both are covered in detail in the section on *Sound Level Metering* on page 40. Note that in order to perform accurate SPL or Leq measurements, the input being monitored must be calibrated to SPL. Please see *Sound Level Calibration* on page 67 for more information.

6 Control Bar

The control bar in impulse response mode includes live measurement controls, bandpass filters signal generator and main display controls

Live Measurement Controls

The upper section of the Control bar is dedicated to live impulse measurement controls for recording impulse response measurements in the field. Note that you can also analyze file-based data in impulse response mode. To load and impulse response from a .wav or .aiff file, select *Load Impulse Response* from the *File* menu and then browse to the location of the file that you want to open in the *Load Impulse* dialog and open it.

The *FFT* size and averaging (*Avg*) controls together determine the measurement duration for dual-channel IR measurements. Notice that for each *FFT* size, the time constant is given along with the *FFT* size in samples. The *FFT* time constant, also called the time window, is the amount time it takes to record the required number of samples at the currently selected sampling rate.

Averages (*Avg*), sets the number of successive IR measurements to average together to improve the signal-to-noise ratio of dual-channel measurements. For deterministic IR measurements made using period-matched signals, the number of averages is normally set to a low number or even “None”. When measuring with random signals or in noisy environments, more averaging can greatly improve the signal-to-noise ratio of impulse response measurements.

The *Tab* selector can be used to switch between tabs if the *Tab Bar* is hidden. The button to its right labeled with the hammer and wrench icon opens the *Measurement Config* page in the *Configurator* dialog.

TF Pair selects the signals to be used for dual-channel IR measurements. Dual-channel IR measurements in *Smaart* are essentially transfer function measurements with the addition of an inverse Fourier transform at the end, and any of the reference and measurement signal pairs that you have set up in the current tab for real-time transfer function measurements may be used for IR measurements. Just click on this control and select the name of a measurement that uses the input channel that you want to record. For single-channel recordings, only the measurement signal channel is recorded.

The live measurement controls in IR mode are analogous to a transfer function “measurement engine” in real-time mode, with a couple of extra twists. Starting from the top left, the buttons marked with a triangle (▶) and a square (■) start and stop a measurement.

The button labeled with a circle (●) works like the record button on a tape deck or digital recorder, but in this case, it is a measurement *mode* control. Clicking the start (▶) button *without* the record button punched in kicks off a dual-channel IR measurement. With the record button (●) activated, *Smaart*

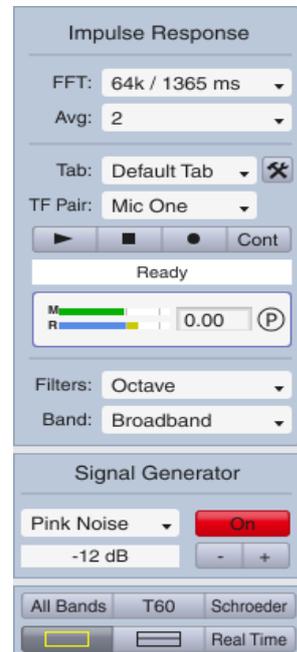


Figure 118: Impulse response mode Control Bar



becomes a single-channel digital recorder and records just the measurement signal channel of the selected signal *TF Pair*. The idea is that you would start the recording and pop a balloon or fire your starter pistol (or whatever), and then click the stop (■) button to end the recording and display your results.

The Continuous (*Cont*) button causes the dual-channel measurement routine to run continuously, starting over again automatically each time it finishes a measurement until you tell it to stop (by clicking the stop button). The results of the last measurement are displayed while it's recording and processing the next measurement. Click the stop button (■) to end the recording and display the recorded data.

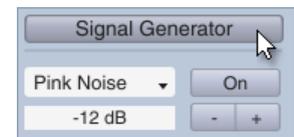
Bandpass Filters

The broadband impulse response is useful for finding delay times and discrete reflections, but for most acoustical analysis purposes, the IR needs to be filtered into octave bands or sometimes 1/3-octave bands. Smaart includes complete sets of octave and 1/3-octave bandpass filters for impulse response analysis. Bandpass filtering is non-destructive and is done on the fly whenever you need it. All that you have to do is select the filter set that you want to use (Octave or 1/3-Octave) using the *Filters* selector on the Control Bar and then choose the center frequency for the band that you want to analyze from the *Band* list.



Signal Generator Controls

The next group of controls on the Control Bar is for the signal generator. The label at the top of this section is actually a hover button (it turns into a button when your mouse cursor passes over it) that opens the *Signal Generator* dialog, which contains a lot more options for the signal generator than we could fit on the Control Bar. Below are a signal type selector (*Pink Noise* is selected in the example to the right) and an output level field that shows the current output level in normalized dB full scale. The *On* button turns the generator on or off – it glows an angry red when the generator is running. The minus and plus (-|+) buttons to the right of the output level field bump the signal level down or up by 1 dB.



Main Display Controls

The last group of controls in the control strip on the right side of the real-time mode window is devoted to data display functions. Starting from the top left of the screen clip shown here on the right:



- The *All Bands* button opens the All Bands table, where you will find nearly all of the quantitative acoustical metrics that Smaart can calculate automatically for an impulse response. See *Histogram and All Bands Table* for more in this feature.
- Clicking the *T60* button displays level marker widgets used for calculating reverberation time and early decay time on Log IR or ETC plots.
- The *Schroeder* button displays a reverse time integration curve on Log IR or ETC plots.
- The two buttons labeled with rectangles divide the main plot area into one or two graph panes: One rectangle, one pane; two rectangles, two panes.

Chapter 8: Impulse Response Mode User Interface

- The *Real Time* button exits IR mode and takes you back to real-time frequency domain measurement mode. (In real-time mode it changes to an *Impulse* button that will bring you back to IR mode.)

7 Command Bar



The Command Bar is a user-configurable button bar that runs across the bottom of a Smart window. You can hide and restore it by clicking the triangular button in the border area just above it. This show/hide button remains visible in the window border when the Command Bar is hidden and clicking it again will restore it to visibility. You can also hide/restore the Command Bar by selecting *Command Bar* in the *View* menu or by pressing the [U] key on your keyboard. To customize the command bar, select *Command Bar Config* from the *Config* menu (see *Configuring the Command Bar* on page 38 for details).

Additional Options for Impulse Response Measurement and Analysis and Display

Some additional options pertaining to IR measurement that don't appear on the main screen can be found in the *General* and *Impulse Response* options pages, accessible by selecting *General* from the *Options* menu or pressing [Alt/Option] + [O] on your keyboard.

General Options Pertaining to Impulse Response Displays

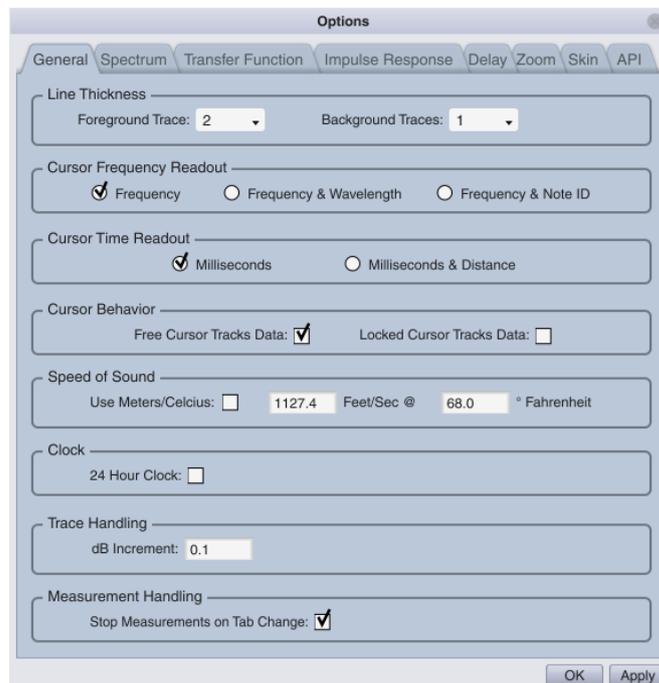
Cursor Time Readout

The cursor time readout setting applies to the both IR Mode and the Live IR graph on Transfer Function displays in Real-Time mode.

- *Milliseconds* – displays time coordinates and relative time differences in milliseconds only.
- *Milliseconds & Distance* – displays time coordinates as milliseconds and equivalent distance, based on the Speed of Sound settings.

Speed of Sound

The settings in this section determine the speed of sound that Smart uses for calculating equivalent distances for time coordinates and also whether distances are displayed in feet or meters. It can also serve as a handy speed of sound calculator any time you need to know the speed of sound for a given air temperature.



- *Use Meters/Celsius* – when this option is selected, Smart displays distances in meters and the temperature used for calculating speed of sound in degrees Celsius. Otherwise Smart displays distances in feet and uses degrees Fahrenheit for temperature.
- *Speed of sound* ([unit]/sec) and temperature – At elevations where humans are comfortable breathing, the speed of sound is mainly a function of temperature, and so the two inputs are linked. Changing the temperature setting automatically recalculates the corresponding speed of sound and vice versa.

Impulse Response Options

Most of the controls in the Impulse Response options dialog tab echo the settings of the controls in the main window. We have already discussed those in some detail so we will concentrate here on the ones that do not.

Time Domain Display Settings

- *FFT Size* and *Averages* echo the settings of the on-screen controls for dual-channel IR measurements that we talked about on page 140.
- *Overlap %* (for Averaging, not to be confused with overlap for the *Spectrograph* display) – When overlap is set to any value other than zero, each successive measurement going into an averaged dual-channel IR measurement shares the specified percentage of data with the previous frame(s).
- *Show IR Peak* sets the locked cursor in IR mode to the highest peak in the impulse response each time you run a new measurement.
- *ALCons Split Time* sets the split time for the early-to-late energy ratio used in calculating ALCons (a type of speech intelligibility estimation that can be calculated from an impulse response). There is no real standard for this parameter but common settings are 10 or 20 milliseconds.
- *Mag Threshold (dBFS)* is similar to magnitude thresholding in transfer function measurements. It is turned off when set to zero. When set to any other value (in dB FS), Smart will zero out the transfer function at any frequency where the reference signal does not cross threshold before calculating the (dual-channel) impulse response.

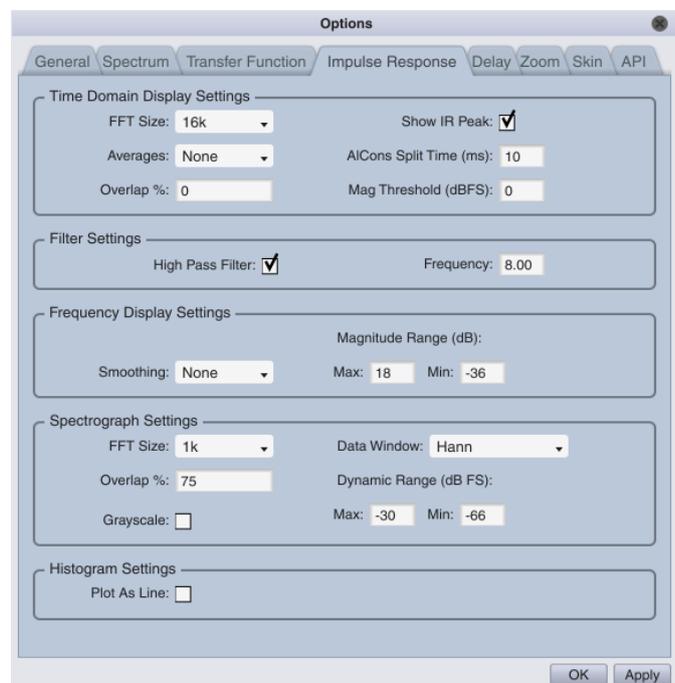


Figure 119: Impulse Response options

Filter Settings

Smaart includes a sweepable highpass filter for IR measurements that can be handy when analyzing IR measurements that include a lot of very low-frequency noise or in cases where the reference signal being used in a dual-channel measurement is band-limited somehow. The filter is applied post process to IR data for display purposes – meaning it can be used for file-based or newly measured data – and only affects what you see on the screen. It does not change the underlying measurement. The cutoff frequency for the filter can be set to any value between 0 Hz and one half of the Nyquist frequency (equal to one half of the sampling frequency) for the currently selected audio sample rate.

Frequency Display Settings

- *Smoothing* echoes the setting of the Smoothing control in the upper right corner of the *Frequency* graph in IR mode.
- The *Magnitude Range (dB)* controls are useful for setting a specific decibel range for the Frequency graph in IR mode, however you can also resize the range using the +/- keys or by right-click-and-drag mouse zooming, as you can with other graph types in Smaart.

Spectrograph Settings

- The *FFT Size* and *Overlap* controls echo the settings of the controls found in the upper right corner of the Spectrograph display in IR mode. Together they determine the time resolution of the spectrograph.
- *Grayscale* plots the spectrograph using varying shades of gray instead of color to represent magnitude.
- *Data Window* – sets the data window function used in calculating the individual FFTs used to create the spectrograph display. You can leave this set to Hann unless you have some good reason to change it.
- *Dynamic Range* echoes the settings of the slider control widgets found on the left edge of time domain and spectrograph displays on IR mode. The spectrograph scales its color (or grayscale) spectrum to the range between the Min and Max values and plots decibel values above the Max thresholds in white and below the Min in black.

Histogram Settings

The *Histogram* chart in IR mode plots the values found for any column in the *All Bands* table band-by-band for all octave or 1/3-octave bands. Selector controls are found in the upper right corner of the chart in the main window. By default, Smaart plots the histogram chart as a bar graph. Selecting *Plot as Line* in *Histogram Settings* causes this chart to be plotted as a line graph instead.

Chapter 9: Analyzing Impulse Response Data

Smaart provides a powerful set of tools for analyzing impulse response data in both the time and frequency domains. Time-domain analysis tools include logarithmic and linear time-domain views, Energy Time Curves, octave and 1/3-octave bandpass filters, reverse time integration and automatic calculation of common acoustical parameters such as EDT, RT60 and clarity factors. Frequency domain analysis tools include spectrum analysis of arbitrary time ranges and the Spectrograph.

If we were discussing Smaart's real-time measurement and analysis mode, we would almost have to pause at this point to set-up and start actively measuring some kind of sound source in order to have something to analyze. But in IR mode, measuring and analyzing are generally two separate things that we can talk about separately. Data analysis in IR mode is an off-line, post-process affair that works the same whether we're onsite actively measuring a system or working with an impulse response recorded in a .wav or .aiff file. Since we just talked about the IR mode user interface in the previous chapter, let's dive right into actually using it.

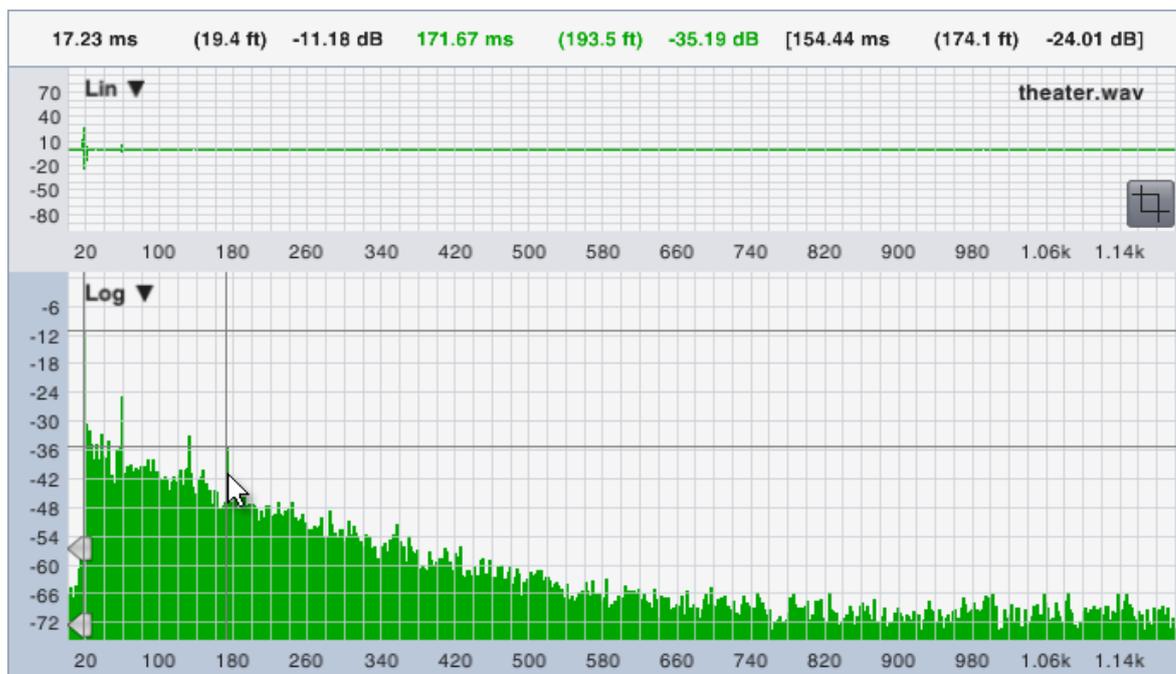


Figure 120: The logarithmic (Log) time domain IR graph plots time on the x axis and magnitude in decibels on the y axis. The combination of locked and movable cursors enables you to find time and level differences between any two points on the plot. Time coordinates can optionally be plotted with equivalent distances as shown above. The pair of coordinates on the far left in the cursor readout at the top of the frame is the locked cursor position, which is set to the highest peak in the IR. The middle pair of coordinates in green is the absolute location for the movable cursor and the rightmost pair in brackets is the difference between the first two.

Most of the examples in this chapter were created using a handful of .wav files that you can download from the Rational Acoustics web site. Wherever applicable, we will tell you which file was used and how to duplicate our settings, so that you can gain a little hands-on experience as we go.

Chapter 9: Analyzing Impulse Response Data

To download the files, visit www.rationalacoustics.com/support/168699-Smaart-v8-Documentation and select *Sample IR Wave Files* from the *Smaart v8 Documentation* section.

Our first example uses *theater.wav*, an IR measurement of a 400-seat historical vaudeville theater. The measurement was taken from the main floor seating area, about 20 feet (6 m) from the stage, using a small horn-loaded PA speaker positioned on the stage lip as the excitation source. If you would like to load the file yourself, open up Smaart, switch to IR mode, then select *Load Impulse Response* from the *File* menu and navigate to wherever the file resides on your hard disk to open it.

Time Domain Analysis

Logarithmic Time Domain Display

The time domain IR display with logarithmic (Log) amplitude scaling is probably the most familiar to anyone much accustomed to looking at acoustical impulse responses. In this view you can find the arrival times of direct sound and early reflections and overlay the reverse time integration of the IR, along with interactive widgets to calculate EDT and reverberation time (on Log and ETC displays only). Smaart provides octave- and 1/3-octave bandpass filters that you can use to filter the IR on the fly, to see how reverberant decay and other characteristics change with frequency.

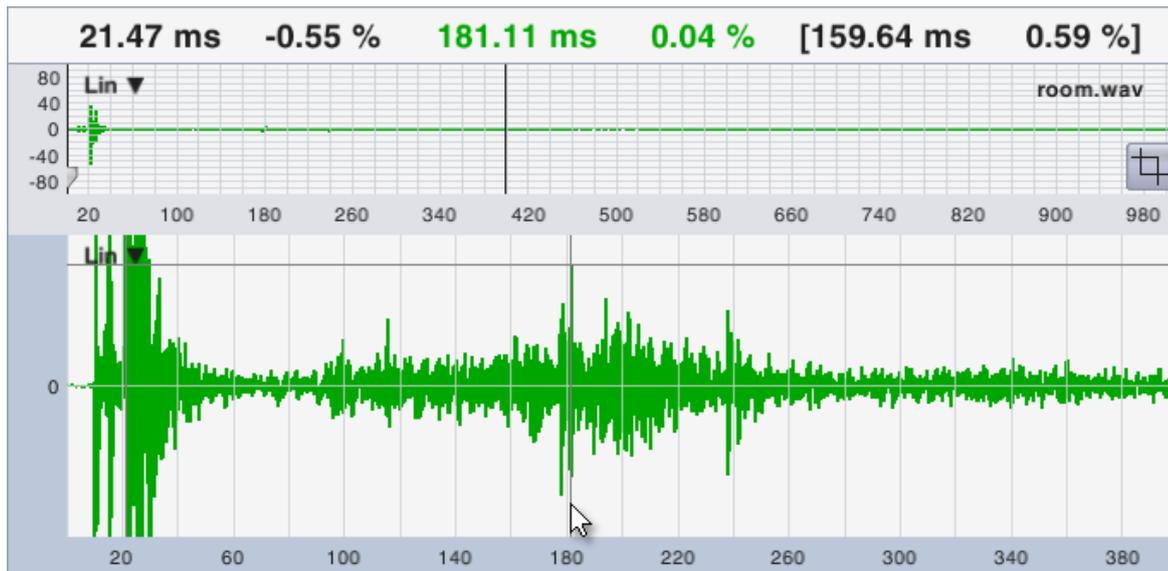


Figure 121: Zooming in on a Linear (Lin) time domain view of *room.wav* and using the cursor readout to find the relative arrival time of a prominent discrete reflection

The combination of locked and free cursors on time domain displays enables you to find the relative arrival time and amplitude differences between any two points on the plot. The difference between the two is shown in the cursor readout. If the Milliseconds and Distance option is selected in the *Cursor Time Readout* section of the *General* options page (*Options* menu > *General*), Smaart will also give you equivalent distances for time coordinates, based on the current Speed of Sound settings. To move the locked cursor to an arbitrary point on the plot, hold down the Ctrl key (Cmd key on Mac) on your

keyboard while clicking with your mouse on a point that you want to mark. Pressing Ctrl/Cmd + “P” resets the locked cursor to the highest peak in the IR.

Linear Time Domain Display

A linear (Lin) time domain chart plots the same data as the Log IR but on a normalized linear amplitude scale, where amplitude values are given as a percentage of digital full scale. This view tends to be of limited usefulness for acoustical analysis in general, however it can be a very good tool for finding discrete reflections, particularly when measuring in an empty hall before an audience arrives. In this case, using the linear IR view can help you to identify hard reflections that might be masked by the diffuse reverberant field on a logarithmic display, only to become much more obvious and audible (often on stage, to the consternation of opera singers) once there is an audience in place and the reverberant levels decrease.

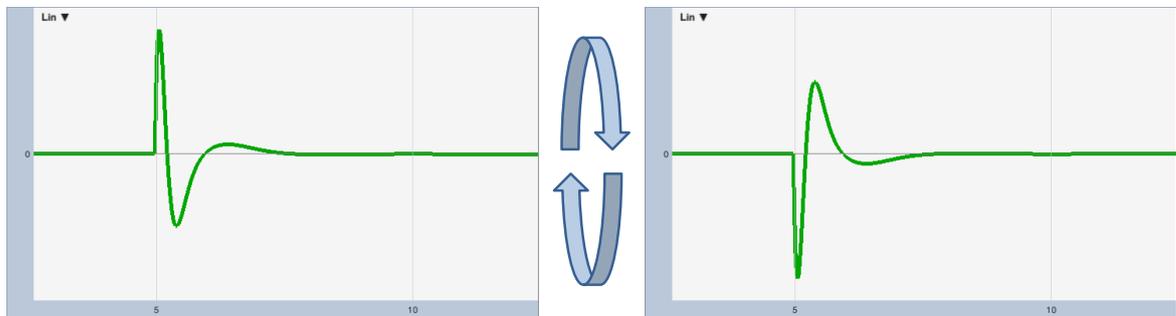


Figure 122: Zoomed in views of the linear impulse response of a bandpass filter, with normal and inverse polarity

Another thing the Linear IR can tell you that the Log and ETC graphs can't is relative polarity. For example you could measure two midrange drivers or other like devices and determine if they are wired with the same or different polarity by noting which direction the prominent peaks in the impulse are pointed. Figure 122 shows a zoomed in view of the linear (Lin) scaled impulse response of a 2nd order Butterworth bandpass filter with normal and inverse polarity. Cutoff frequencies for the filter are 400 and 1600 Hz. It's easy to see that the peaks in the two IRs are pointed in different directions relative to each other. Unfortunately, this doesn't necessarily tell you which one is correct. But if you measured three like devices and one was different, you might reasonably say that the majority rules. Or if you measured two like devices and found opposite polarity and one of them sounded better, it's possible you might have found the problem. Linear view can also come in handy for looking at other types of signals in the time domain other than impulse responses.

Energy Time Curve (ETC)

The impulse response represents a 2-D graph of a 3-D event: the magnitude and phase of the energy arrival over time. With magnitude on the vertical (y) axis and time on the horizontal (x) axis, phase ends up being represented on the z-axis, which is effectively lost in this view. Consequently, in the linear and log views of the IR, energy arrival that is 90° or 270° shifted shows as a zero crossings, thereby making a single arrival that is spread out over time and phase appear to be multiple arrivals.

Chapter 9: Analyzing Impulse Response Data

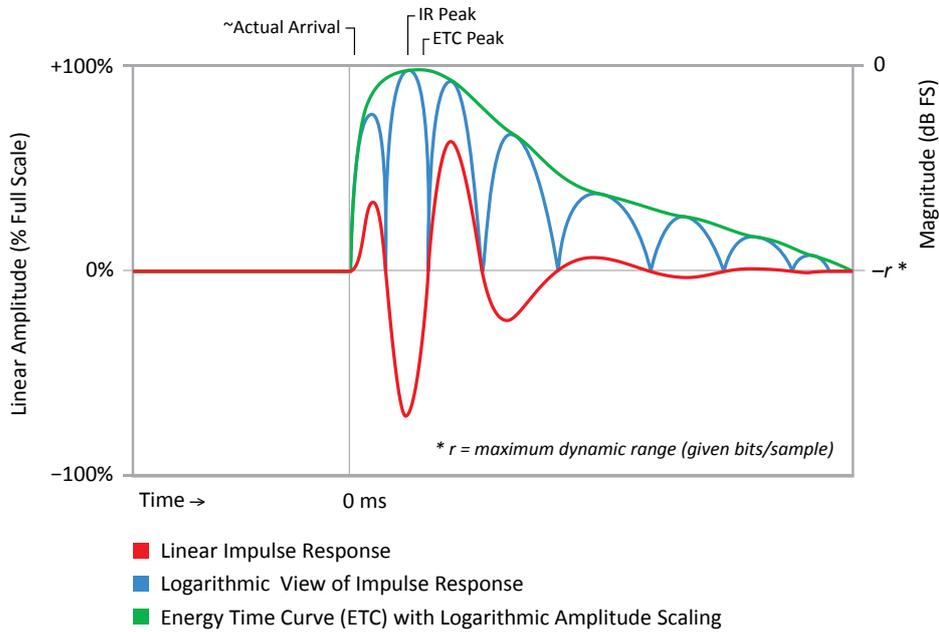


Figure 123: A comparison of the ETC and the impulse response with linear and logarithmic amplitude scaling

The Energy Time Curve, also called envelope of the impulse response, represents the magnitude of the energy arrival over time by effectively ignoring phase. The textbook description is the real impulse response combined with its Hilbert transform – a copy of itself that has been rotated 90° in phase. In practical terms, the summation of the two tends to fill in zero crossings seen in the Log IR, producing a signal that can be a lot easier to look at than the Log IR by virtue of being less squiggly. At higher frequencies the Log IR and ETC may look very similar – both are plotted on a logarithmic magnitude scale – but the ETC is particularly useful for sizing up the arrival of direct sound at low frequencies.

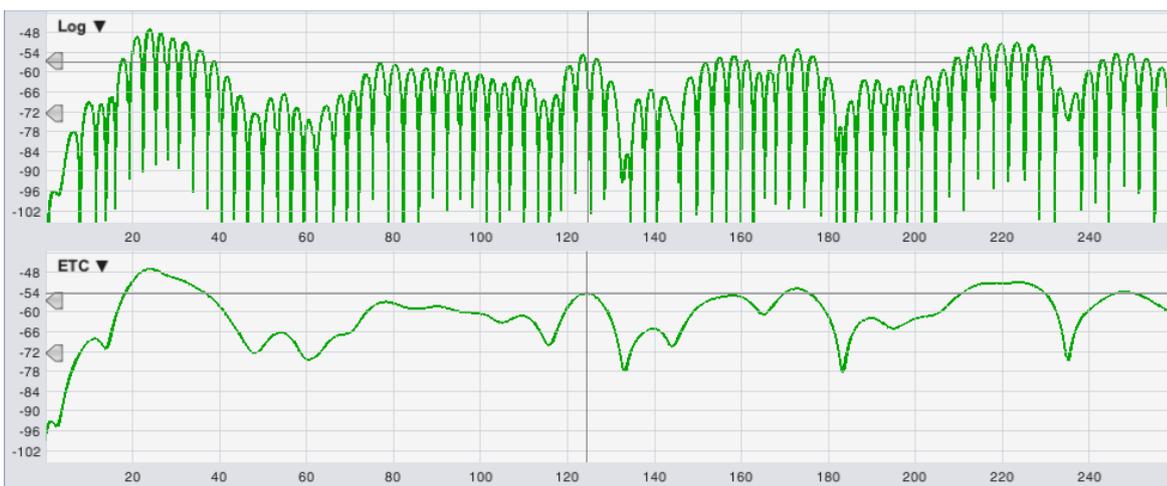


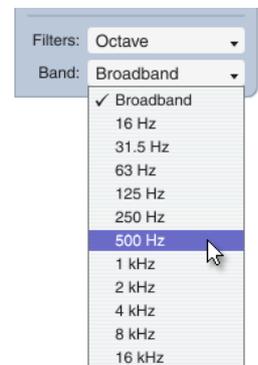
Figure 124: A comparison of the Log IR and ETC graphs in Smaart for the 125 Hz octave band in room.wav

If you zoom in on the first 250 ms of room.wav and switch to the 125 Hz octave band, the difference between the Log IR and ETC is pretty striking (see Figure 124). To scale your display to look like Figure 124, press the plus [+] key on your keyboard a few times to zoom in on the magnitude range then use the up/down arrow keys to move the range up and down.

Note that when using peak locations to find delays, the ETC can sometimes give you a slightly different answer than the Log IR, because of the way it effectively interpolates between peaks in the IR. If you look at the smaller peak in the ETC, at about 124 ms in Figure 124, you can see that it falls in between two lobes in the Log IR. We have found that the ETC can be more effective than the Log IR tool for finding subwoofer delay times. But that is better done in real-time mode, using the ETC on the Live IR in conjunction with the frequency domain transfer function displays, where you can see phase as well as magnitude and watch changes happening in real time as you adjust processor settings

Bandpass Filtering

Up to now we have mainly been looking at the broadband IR, but quite a lot of acoustical analysis is conventionally done using octave, or sometimes 1/3-octave bands, especially as we get into reverberation times and early-to-late energy ratios. Smaart includes complete sets of octave and 1/3-octave bandpass filters for the octaves between 16 Hz and 16 kHz (assuming 48k or higher sampling rate, at lower sample rates you lose some of the upper bands). Bandpass filtering in Smaart is done non-destructively, on demand. To see a filtered version of the IR, select which set of filters to use (Octave or 1/3 Octave) on the Filters selector, then select the band that you want to look at from the Band list.



Smaart's bandpass filters have linear phase response and their magnitude response satisfies the most stringent (Type 0) tolerances for octave and fractional-octave bandpass filters specified in IEC 61260 and ANSI S1.11. If you would like to see the magnitude response of the bandpass filters you can load the wave file *1samplePulse.wav* and bring up the Frequency graph, then step through the *Bands* list to see each filter. Bandpass filtering applies to all main display types except the Histogram chart (which is already filtered into bands). It does not affect the small graph in the navigation pane. Note that filtering the impulse response will clear the Spectrograph display if present and require a recalculation (by clicking the *Calc* button again).

Discrete Reflections

Reflections are a complicated subject because humans are very good at processing them. They may be useful or detrimental, depending on such factors as their arrival time and loudness relative to direct sound (the two biggies), their frequency content and even the angle they arrive from. Discrete reflections can cause audible problems ranging from coloration (timbre change) to image shift to audible echoes, but trying to figure out which reflections are friend or enemy by looking at squiggly lines on a computer screen can be a bit of a dicey prospect.

Short reflections arriving within the first 30 milliseconds or so after the direct sound at relatively high levels are notorious for producing comb filters that muck up our real-time frequency domain analysis;

Chapter 9: Analyzing Impulse Response Data

but humans actually find them beneficial, enhancing the intelligibility of speech and the clarity of music. Outside that early integration window, reflections can still contribute to subjective impressions of presence, warmth, spaciousness, etc. However, the rules are a little different for speech vs music.

Individual broadband reflections arriving at 95 ms or more can destroy speech intelligibility and make life difficult for presenters and performers if they reach the stage. This is the threshold of where strong reflected sounds begin to be heard as separate events (echoes) and can be disorienting for anyone trying to speak or sing. This happens to have been the problem being investigated in the IR measurement shown in in Figure 134 on page 160, where a high-level reflection was arriving at about 160 ms, which is close to the average syllabic rate for normal, conversational speech.

Low-order, early reflections may be visible on time domain plots as individual peaks following the arrival of direct sound. Later arrivals can show up as spikes protruding from the reverberant decay slope. On the Spectrograph plot, higher-level broadband reflections can often be identified as distinct vertical streaks when you run the dynamic range controls up and down, particularly the Max setting. They tend to be most problematic when arriving at longer delay times and relatively high levels, compared to the level of the diffuse reverberant field.

A pretty good rule of thumb is that the later the arrival, the lower in level it needs to be in order to be perceived as beneficial or neutral. Another is that our tolerances for reflected sounds and reverberation tend to be wider for music than speech. Smaart is very useful for identifying problematic reflections; however your ears are probably still the best tool for evaluating their relative significance or severity.

Reverberation Time

Reverberation time (commonly referred to as T60 or RT60, or somewhat less commonly as T30, T20 or simply T) is the time required for reverberant sound energy in a space to decay by 60 dB from an excited level. It is regarded as an important metric in the acoustics of musical performance spaces and also classrooms, auditoriums and cinemas, where it is used as a rough predictor of speech intelligibility.

Reverse Time Integration

Reverberation time is calculated from the reverse integration of an impulse response that has been filtered into octave bands. Conventionally, the 125 Hz to 4 kHz bands are evaluated. Reverse time integration is also called Schroeder integration, after Dr. Manfred Schroeder whose brain child it was. It is a simple thing in concept, but it can be a little tricky to do well.

In theory, you just start at the end of the time record and work your way back to the beginning, tallying up the squares of each sample in the IR as you go. A common problem however, is that the integration will flatten out when the reverberant decay slope runs into the noise floor of the IR. This can lead to overestimation of the reverberation time, particularly if the IR has limited dynamic range, and/or a lengthy noise tail.

The most straightforward solution for this problem is to find the point in the IR where the decay slope meets the noise floor, sometimes referred to as the “saddle point,” and begin the integration there, rather than some arbitrary point such as the end of the recording. The location of the saddle point in an IR is notoriously difficult to estimate automatically though. Smaart 8 uses a proprietary algorithm for IR

saddle point estimation that has proven quite robust, but it is not completely foolproof. Therefore, it is still a good idea to check each band to make sure that you agree with the choices the software makes – particularly if there are any large anomalies in the tail of the IR, such as a prominent spike or distortion products piled up at the end of the record by a sweep signal.

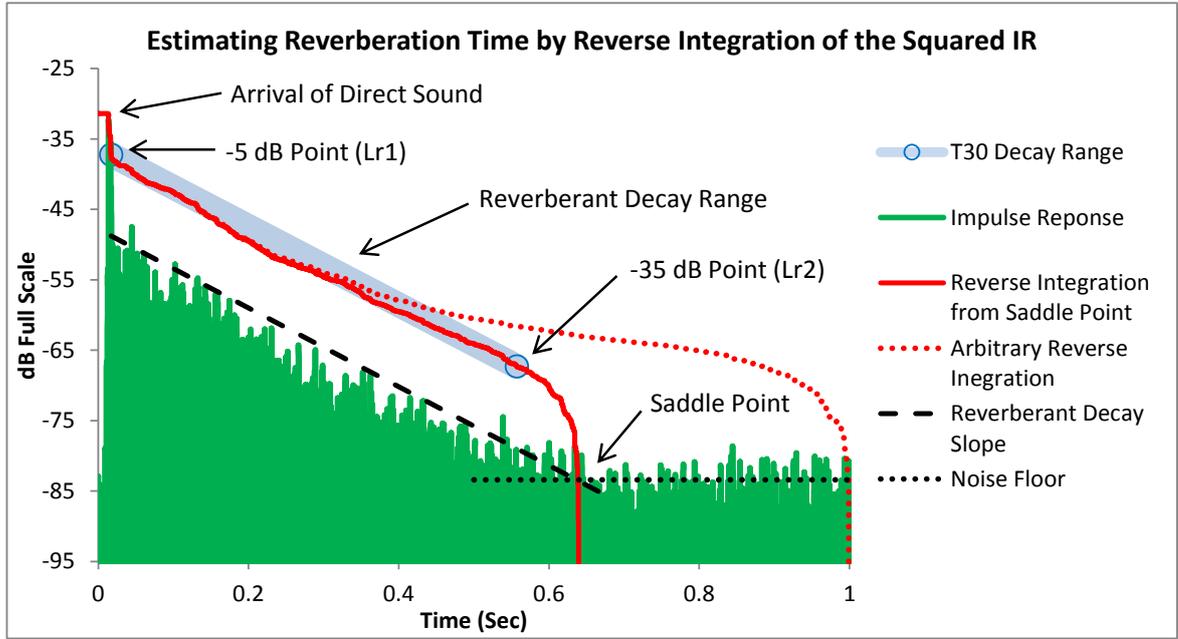


Figure 125: Estimating Reverberation time by reverse integration of the impulse response. Reverse integration of the IR from the “saddle point” – the approximate point where the reverberant decay slope meets the noise floor of the measurement – provides a very good estimation of reverberant decay time. Starting the reverse integration from an arbitrary point such as the end of the file, may result in overestimation of decay time.

Evaluation Ranges (EDT, T20, T30)

Because it is rarely possible to actually measure a full 60 dB of reverberant decay in acoustical systems, reverberation is typically evaluated over a smaller range. The starting point is always 5 dB down on the reverse integration curve from the point corresponding to the arrival of direct sound. The end point of the range is 30 dB down the curve from the starting point, provided that it is at least 10 dB above the noise floor – if not, a 20 dB range may be used. In either case, the measured decay time is extrapolated to the equivalent 60 dB decay time. In ISO 3382 parlance, these are referred to as T20 or T30. Early decay time (EDT) is conventionally measured from the arrival of direct sound down to 10 dB below it on the integration curve. Like reverberation time, EDT is also normalized to 60 dB decay time.

Notice the five level marker widgets shown on the plot in Figure 126. If you were wondering about the cryptic labels, your secret decoder and the default positions for each of the markers is as follows.

- *Ld* = Level Direct. This marker is positioned on the reverse integration curve at the point corresponding to the arrival time of direct sound.
- *Le* = Level Early (Decay). This marker is automatically positioned 10 dB down from the *Ld* marker on the reverse integration curve. The slope between *Ld* and *Le* is used to calculate EDT.

Chapter 9: Analyzing Impulse Response Data

- $Lr1$ = Level Reverberant 1. This marker designates the top of the reverberant decay range, 5 dB down the integration curve from the Ld marker. All of the level markers are user adjustable but positioning these three is pretty cut and dried. You *should* rarely find any need to touch them.
- $Lr2$ = Level Reverberant 2. This marker designates the end point for the reverberant decay slope. If there is sufficient dynamic range it should be placed 30 dB down the reverse integration curve from $Lr1$. If not, 20 dB will do. $Lr2$ is one of the two markers that you may sometimes want to adjust by hand; the other is Ln (below).
- Ln = Level of Noise. This is typically the most subjective of the five markers in terms of placement. The time location determines the start point for the reverse time integration curve, which is the basis for positioning all of the other markers. Ideally this will roughly correspond to the saddle point in the impulse response. The magnitude coordinate is used to estimate the level of the noise floor of the measurement and the $Lr2$ marker needs to be at least 10 dB above that. Smart does a pretty good job of placing the Ln marker most of the time. However, it may still benefit from a human touch in some cases – particularly if the dynamic range of the measurement is marginal or there are significant distortion artifacts from a swept sine measurement or any other prominent anomalies in the noise tail of the IR being analyzed.

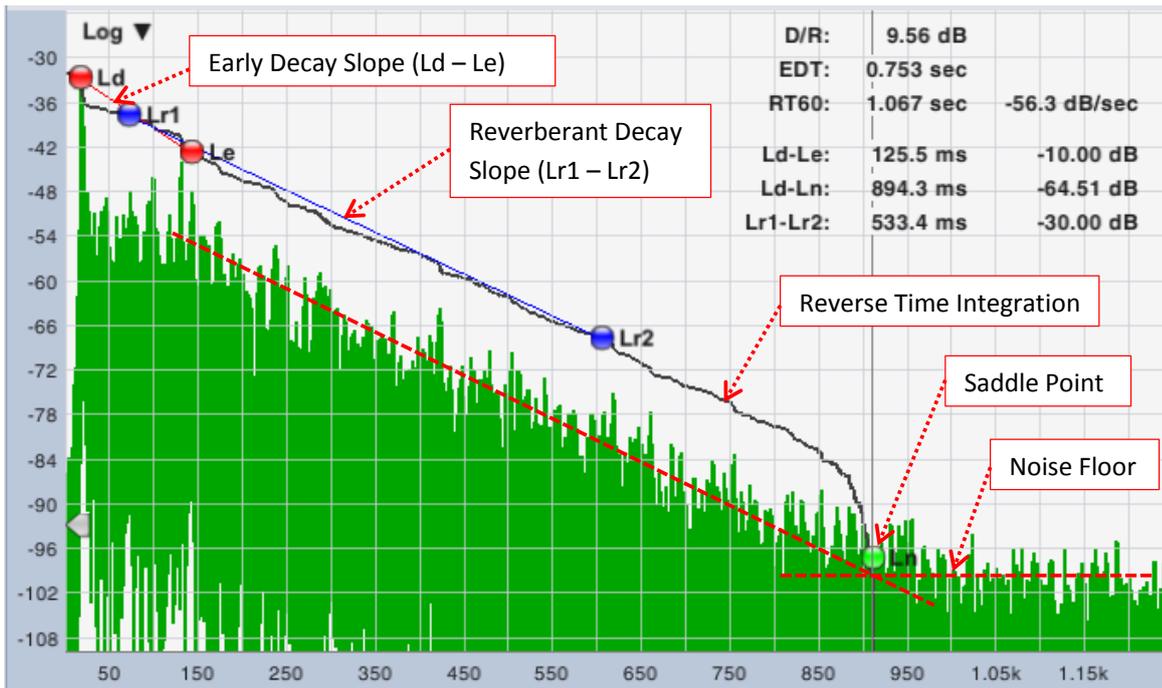


Figure 126: A log IR display with all the bells and whistles. The impulse response shown here is the 500 Hz octave band of theater.wav. Clicking the Schroeder and RT60 buttons displays the reverse time integration curve and the start and end points for the EDT and RT60 evaluation ranges on Log or ETC charts. The positions of all the level markers (Ld , Le , $Lr1$, $Lr2$ and Ln) are user adjustable, however they should not require very much adjusting under most circumstances.

When the level marker widgets for reverberation time are visible, a block of vital statistics also appears in the upper right corner of the plot. These include the 60 dB reverberation and early decay times (RT60 and EDT) and the time and level differences between three pairs of markers. The *Ld-Le* level difference should always be 10 dB. *Lr1-Lr2* should be either 20 or 30 dB – this number is convenient for checking your work if you end up adjusting *Lr2* is by hand. The *Ld-Ln* delta is interesting also, as it gives you the dynamic range of the measurement. D/R stands for direct/reverberant ratio. It is an early-to-late energy ratio that gets its split time from the time coordinate of the *Le* marker.

Saving your work

If you do end up adjusting any of the Level marker positions by hand, you can save their positions to a comma-separated values (.csv) text file, by selecting *Save Decay Markers* from the *File* menu. To reload them again later, first load the .wav or .aiff file containing the impulse response, then select *Load Decay Markers* from the *File* menu to open the marker position file.

Reporting Results for Reverberation Time

Reverberation time ideally should be measured from several locations throughout the room and the results from each measurement position averaged together, octave band by octave band to get an average decay time for each octave. Smart doesn't do that part for you, but the All Bands table does make it easy to get the data from each measurement into a spreadsheet.

Frequency Ranges

The standard evaluation range for reverberation time is the six one-octave bands from 125 Hz 4 kHz. Average times for each octave band can be presented in a table or on a graph. When presenting reverberation times on a graph, the frequency axis of the graph should be labeled with the IEC standard nominal octave band center frequencies. The y-axis of the graph should have an origin of 0 and be labeled in seconds. It should be noted both in the table and on the graph whether T20 or T30 was used. ISO 3382-1 specifies that if a graph is presented it should be a line graph with a standardized aspect ratio of 2.5 cm per second and 1.5 cm per octave. ISO 3382-2 isn't so picky. It just says "a graph."

Reverberation times for the 125 and 250 Hz bands may be averaged together to get a T_{Low} figure. The average of the 500 Hz and 1 kHz band is called T_{Mid} . When a single number figure is given for reverberation time, it is assumed to be T_{Mid} unless otherwise stated. Smart calculates these values for you automatically and displays them in the All Bands Table.

Dividing T_{Low} by T_{Mid} gives you the Bass Ratio. Bass ratio quantifies the "warmth" of sound in a venue and is a particularly important parameter for concert halls. The word "Bass" in this case refers to vocal or instrument bass registers and should not be confused with PA-type sub-bass frequencies. Acceptable values are dependent on expectations. A Bass Ratio of 1.1-1.25 would be regarded as good for fairly reverberant concert halls (RT 60 greater than 1.8 seconds) but the upper figure could be increased to 1.45 for less reverberant spaces.

As for what to look for in reverberation time results, preferred reverberation times vary according to room size and purpose and the type of program material being presented. In general, you would like to see shorter reverberation times for auditoriums, classrooms, theaters and cinemas – ideally from about 0.4-0.5 seconds for smaller rooms, up to 0.8 to 1.2 seconds for larger rooms. Opera houses and mixed-

Chapter 9: Analyzing Impulse Response Data

use performance spaces where both speech intelligibility and musical appreciation are equally important typically aim for the lower end of the 1.2 to 1.8 second range. Spaces intended for symphonic performances and organ music can range from about 1.8 seconds up to three seconds or more in very large halls.

Reverberation times that are roughly equal across all frequencies are generally preferable for most purposes. The exceptions are things like choral, organ and romantic classical musical music, where a reverberation time curve weighted more toward the lower frequencies may be preferred. It is pretty normal for higher frequencies to decay faster than lows but you don't want to see times that are wildly different in neighboring octaves. In general though, acoustical treatments and/or physical changes to the sound system are typically required to effectively address problems any problems you may find.

Early-to-Late Energy Ratios

Early-to-late energy ratios are another way of objectively characterizing the reverberant characteristics of a room. They are arguably a better measure than reverberation time for any venue where a sound system is an organic part of the acoustical equation. They are simple to calculate automatically and are not subject to the kinds of complications that can make measurement of reverberation times somewhat subjective, but they are a more recent innovation and may be less widely understood than RT60.

Clarity Ratios (C35, C50, C80...)

Clarity indexes are early-to-late energy ratios that compare the integral of the energy arriving within the first n milliseconds of the arrival of direct sound (inclusive) to the energy in the remainder of the reverberant decay period. The two most commonly used are C50 and C80, which use at the 50 or 80 milliseconds respectively as their split times. The result of the comparison is expressed as a decibel ratio.

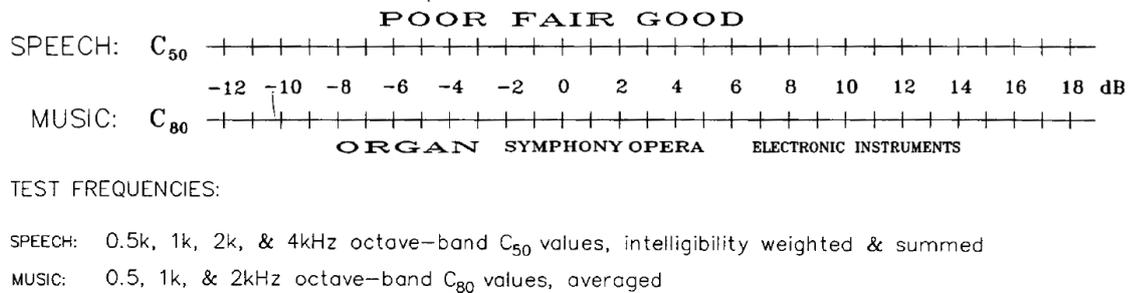


Figure 127: A scale for interpreting C_{50} and C_{80} measurement results for speech and music

Shorter split times such as 35 or 50 ms are regarded as better predictors of speech intelligibility. C_{80} is more useful for music. In terms of what kinds of numbers to look for, Gerald Marshall provided the table shown in Figure 127 in the in a 1996 Journal of the AES article titled, *An Analysis Procedure for Room Acoustics and Sound Amplification Systems Based on the Early-to-Late Sound Energy Ratio*.

For the speech intelligibility side of the graph, Marshall used a weighted average of the 500 Hz to 4 kHz octave bands, with the following weights assigned to each band: 15% for 500 Hz, 25% for 1 kHz, 35% for 2 kHz and 25% for 4 kHz. Others have used the weighting tables for Articulation Index, STI and other

scales of their own devising with similar results. For music, he used a simple average of the 500 Hz, 1 kHz and 2 kHz octave bands. We know of no applicable standards for this metric and it has been suggested extending the frequency ranges that Marshall used an octave higher for speech and two octaves higher for music might be useful, but hopefully this example provides a useful starting point for evaluation.

The Histogram Display

Selecting *Histogram* as your display type for an IR mode graph plots a chart of all reverberation times or early-to-late energy ratios by octave or 1/3 octave bands. The type of data the Histogram displays and the resolution are selected by means of the list control in the upper right of the graph. You can change the Histogram to a line chart by opening the *Impulse Response* options page (*Options* menu > *Impulse Response*) and selecting *Plot as Line* under *Histogram Settings*.

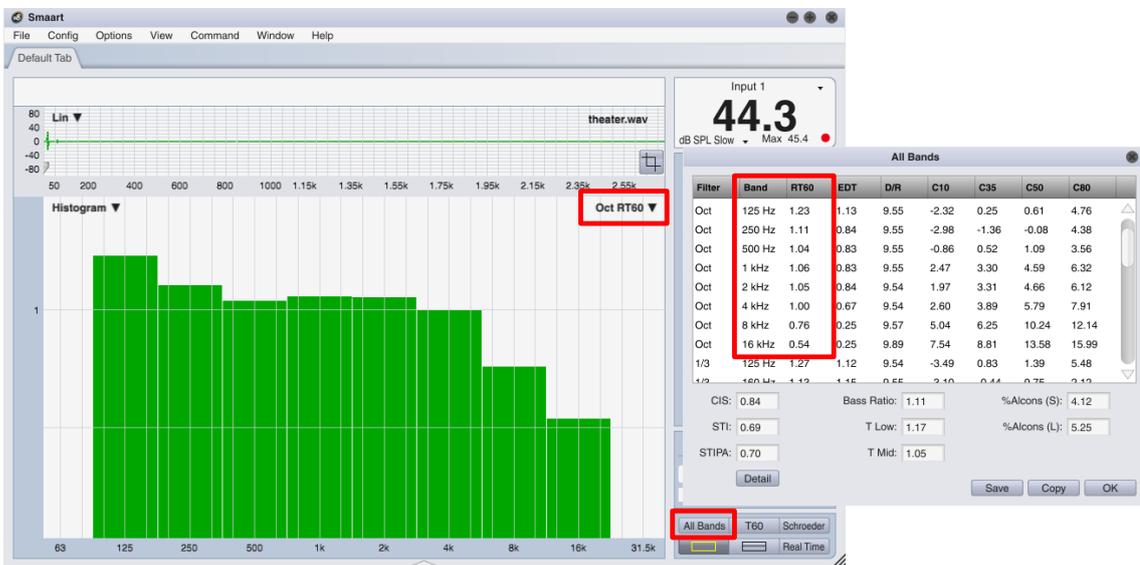


Figure 128: The Histogram graph and All Bands Table. The All Bands Table in Smart collects reverberation times and early-to-late energy ratios for all octave and 1/3-octave bands in a single table. Speech intelligibility metrics (STI and ALCons) are displayed here as well. The Histogram chart can plot any column of the All Bands Table as a bar graph or line chart in octave or 1/3-octave resolution.

The All Bands Table

Clicking the *All Bands* button in Impulse mode brings up a report window containing just about every acoustical quantity that Smart can calculate automatically from an IR, for each octave band and 1/3-octave band where applicable (see Figure 128). Speech intelligibility metrics (STI and ALCons) are displayed here along with *Bass Ratio*, *T Mid* and *T Low*, which are calculated from reverberation times for the 125 Hz to 1 kHz octave bands.

Clicking the *Copy* button in this window copies the entire table to the operating system’s clipboard in tab-delimited ASCII format suitable for pasting into a spreadsheet or any other program that accepts ASCII text. You can also save it to a text file by clicking the *Save* button.

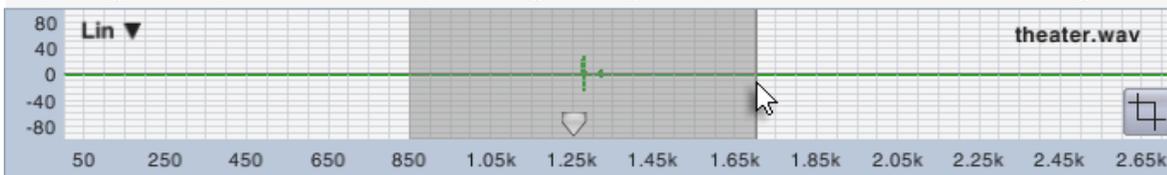
Frequency Domain Analysis

Selecting Frequency as your graph type in the main graph area automatically transforms the IR into the frequency domain to show you its spectrum. The *Frequency* graph has frequency in Hertz on the (horizontal) x axis and magnitude in decibels on (vertical) y axis. The Smoothing control in the upper right corner of the *Frequency* graph works exactly the same way as smoothing on the real-time transfer function display.

Step 1: Click and drag Time 0 slider to right.



Step 2: Right-click (Ctrl + click on Mac) and drag in navigation pane to select a desired time range.



Result: The spectrum of the selected time range is displayed in the Frequency graph.



Figure 129: Moving the time 0 point and selecting a time range for display. The time range selected in the navigation pane applies to the Frequency graph as well as time domain graphs (Lin, Log or ETC). Note that Smart uses a tapered data window when transforming any subset of the full IR time range. We've drawn the outline of a Hann window in red on the navigation pane of the "result" portion of the illustration above to help visualize this.

Smart can calculate arbitrary-length DFTs in IR mode to give you the spectrum of virtually any subset of the IR time record that you care to zoom in on. Time and frequency domain displays are linked so that zooming in on a time-domain graph (Lin, Log or ETC) automatically changes the Frequency display to

match. When the entire time record is selected, there's an assumption that you are analyzing a dual-FFT IR measured in Smaart and so no data window is used in calculating the spectrum in that case.

Smaart automatically uses a tapered data window when transforming any subset of the time record, so if you are analyzing an IR file from some other source or a file that has been cropped to less than its original length, you may see better results if you zoom in slightly in the time domain. Tapered data windows significantly attenuate data at the edges of a selected time range – many go all the way to 0 – so you generally want to position any peaks that you want to examine near the center of a selected range. The Time 0 slider in the navigation pane can be used to move peak structures nearer to the center of the time window if they are too close to the edge to center up in the range. Clicking in the right margin of the navigation clears a time zoom.

If you zoom way in on the IR and select a very narrow time range centered on the arrival of direct sound it's possible to see the magnitude response of loudspeaker without comb filters caused by early reflections, at least at high frequencies. In practice, the usefulness of this strategy may be limited to how far away both the loudspeaker and microphone are from the nearest reflecting surfaces. The frequency response of a DFT is limited by its time constant, so you may find that by the time you squeeze the time window in enough to get rid of first order reflections, you can't really see much detail in the frequency domain. But it's something that people used to do quite often in the days before lab-measured anechoic response data for most professional loudspeakers became commonly available.

The Spectrograph

The Spectrograph display in impulse response mode is essentially the same display as the real-time spectrograph. If you understand one, then you understand the other – and if you don't understand either you may want to review *Spectrograph Basics* on page 95. The principle difference between the two is that the IR mode spectrograph is rotated 90° relative to the real-time version, to put time on the x axis instead of frequency. In real-time mode in Smaart we want to relate the spectrograph to other frequency-domain graphs, but in IR mode, we most often want to look at it in the context of other time-domain graphs. The other main difference is that the number of "slices" in the IR mode spectrograph is determined by FFT size and Overlap parameters that you select.

To bring up the spectrograph in IR mode, click on the graph selector in the upper left corner of a main graph area pane and select Spectrograph. Initially, you are presented with a blank chart area until you click the *Calc* (calculate) button in the upper right corner of the graph pane. The Spectrograph can eat up a lot of graphics resources when it repaints and so we try to paint it only when necessary. Changing the time range selection in the navigation pane does not affect the spectrograph as it does the other time domain (Lin, Log, and ETC) and Frequency graphs in IR mode but moving Time 0, cropping the time record or filtering the IR will clear the spectrograph and require clicking the *Calc* button again. You can resize the spectrograph and move its range using the [+], [-], and arrow keys or right-click-and-drag on the plot to zoom in on a selected range as you can with any other graph in Smaart. Also, as with other graph types in Smaart, clicking in the left margin of the plot after zooming in will clear the zoom and return the plot to its previous x/y range.

Chapter 9: Analyzing Impulse Response Data

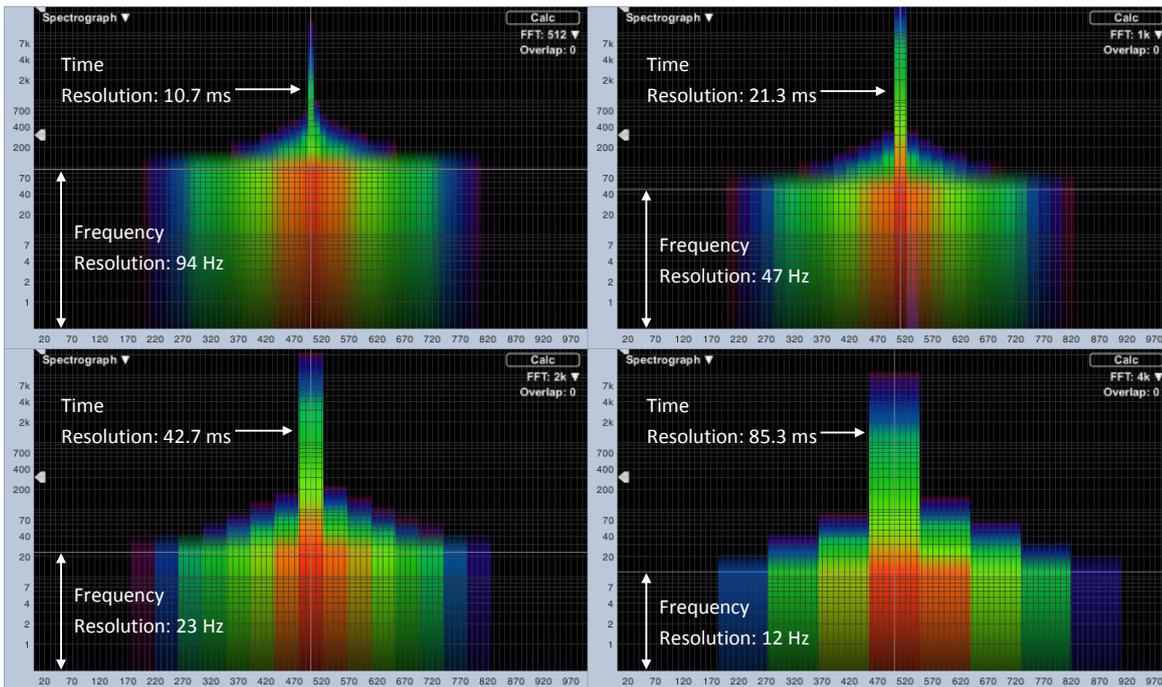


Figure 130: Spectrograph time and frequency resolution as a function of FFT size. FFT sizes ranging from 512 samples to 4K samples are compared, using 0% Overlap. As the FFT size is increased, frequency resolution improved but the peak of the IR is smeared out over a wider time range. The x axis of the graph is time, with frequency on the y axis.

Spectrograph Time and Frequency Resolution

In *Spectrograph Basics* in Chapter 5 (see page 96), we used an octave band analyzer as an example for the sake of simplicity. The IR mode spectrograph is narrowband but the concept is the same. Each vertical stripe in the IR mode spectrograph represents one FFT, meaning the FFT size determines both the time and frequency resolution of the display. Larger FFTs provide greater detail on the frequency axis but may mask transient events on the time axis, so it's a trade-off in that respect. The *FFT* and *Overlap* controls that appear below the *Calc* button determine the time and frequency resolution of the spectrograph display.

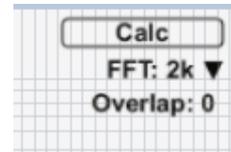


Figure 130 illustrates how this relationship works. It was created using the file *6dbOctImpulse.wav*; the impulse response of a linear phase lowpass filter with 6 dB per octave roll-off. The sharpest part of the peak in the impulse response, where most of the HF energy lives, is probably no more than a few milliseconds wide but notice how its energy is spread across the full FFT time constant in every case. At 512 points, time resolution (the FFT time constant) is a respectable 10.7 milliseconds

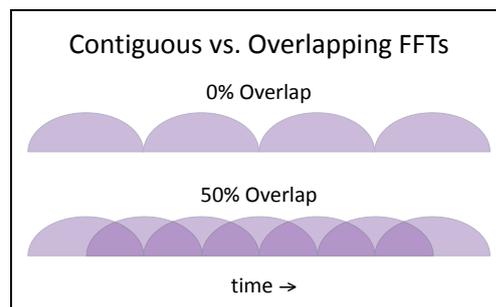


Figure 131: FFT overlap. When set to 0%, each FFT is calculated from unique data. At 50% overlap, the darker shaded areas are shared by successive FFTs. Our "FFT"s are drawn as flattened half circles to suggest a tapered FFT data window.

but the FFT frequency bins are spaced almost 100 Hz apart. Increasing the FFT size to 4K points gets you 12 Hz frequency resolution but smears the peak in the IR over an 85 ms time range.

The other factor affecting the time resolution of the spectrograph is the Overlap percentage that you specify in the upper right corner of the graph. When Overlap is set to zero as in Figure 130, each successive FFT “slice” of the Spectrograph is calculated from unique time domain data, each frame beginning where the last one ended. When Overlap is set to any non-zero value, each successive FFT frame shares some percentage of its data in common with the previous frame(s). The FFT time constant is still the FFT time constant but more overlap can sometimes allude to, if not exactly restore some missing detail on the time axis as FFT size is increased, in addition to producing smoother blending between “slices” (see Figure 131).

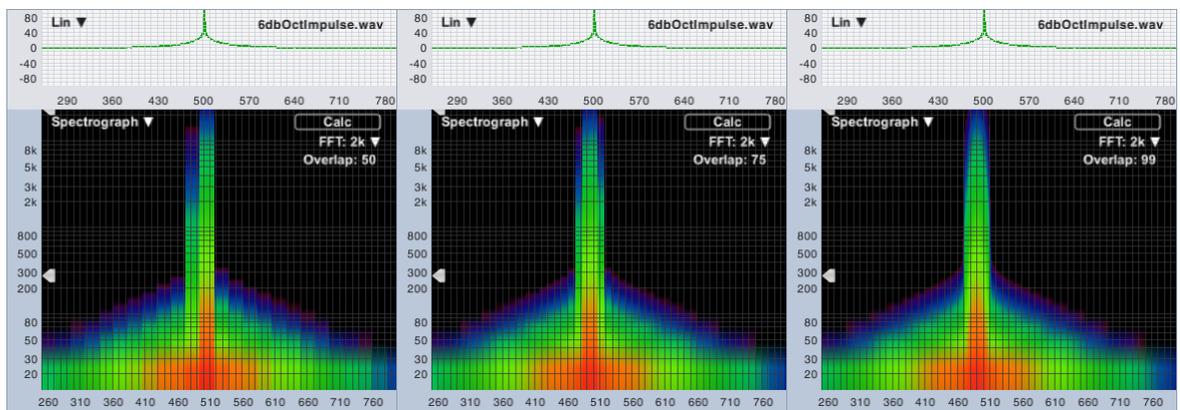


Figure 132: The 2K FFT example from Figure 130, shown with 50%, 75% and 99% overlap(left to right)

Spectrograph Dynamic Range

The dynamic range of the Spectrograph is controlled by two arrowhead-shaped widgets that appear on the left edge of the Spectrograph chart. These controls are echoed on the Log IR and ETC displays where they relate directly to the decibel levels on the graph. The upper of the two widgets sets the maximum (Max) threshold; the lower one sets the minimum (Min). The Spectrograph scales its color spectrum between these two extremes. Any FFT bin whose magnitude exceeds the specified maximum is mapped to the color white. Values falling below the minimum are mapped to black.

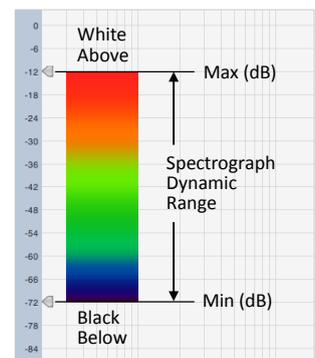


Figure 133: Spectrograph dynamic range and color mapping

Spectrograph Analysis of an Acoustical Impulse Response

Long-time Smaart users may recognize the impulse response measurement in Figure 134 as the *room.wav* file that was distributed with early versions of Smaart. The measurement was recorded on the stage of a 6000 seat performance space using an overhead PA cluster as the excitation source. It features a very prominent reverberant build-up phase and problematic late reflected energy arriving about 160 ms after the direct sound.

Chapter 9: Analyzing Impulse Response Data

This is a good file to experiment with to see how changing the FFT size, Overlap and dynamic range can reveal different aspects of the IR. You can see that we've set the navigation pane graph type to ETC and moved Time 0 to about 100 ms. The FFT size is 2K, overlap is 95% and dynamic range is -20 to -60 dB.

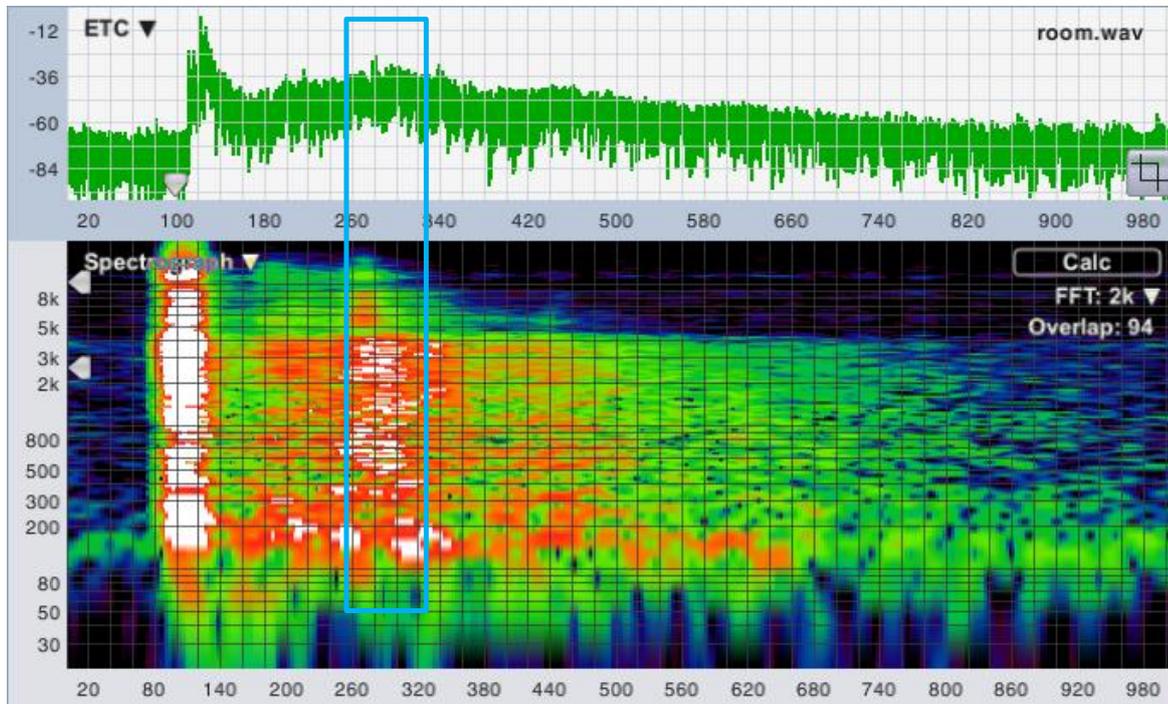


Figure 134: Broadband ETC and Spectrograph of a room impulse response showing a problematic back wall reflection. The spectrograph can be extremely useful for examining both the level and frequency content of features in the IR such as reverberant build-up and discrete reflections.

Speech Intelligibility Metrics

We have discussed the applicability of early-to-late energy ratios as good predictors of subjective speech intelligibility and RT60 as a somewhat rougher gauge. There also exist some purpose-built metrics created specifically for predicting speech intelligibility. Two of these, ALCons and STI, can be calculated from the impulse response of an acoustical system.

STI and STIPA

The Speech Transmission Index (STI) has emerged over the last decade or so as the go-to metric for objective estimation of speech intelligibility in acoustics. STI is a relative of the Articulation Index (AI), which is based on speech-to-noise ratios across a wide range of frequencies. But unlike AI (and its successor, SII), STI works well for reverberant environments in addition to communication systems.

Rather than estimating intelligibility based on direct-to-reverberant or signal-to-noise ratios, STI starts with the concept of speech as a carrier wave (sound from our vocal cords) that is modulated by very low frequency fluctuations as the speaker's mouth and tongue move and change shape to form words (or more precisely, the phonemes from which spoken words are constructed). Looking at Figure 135, a segment of actual human speech, it's not hard to see how someone might arrive at that conclusion.

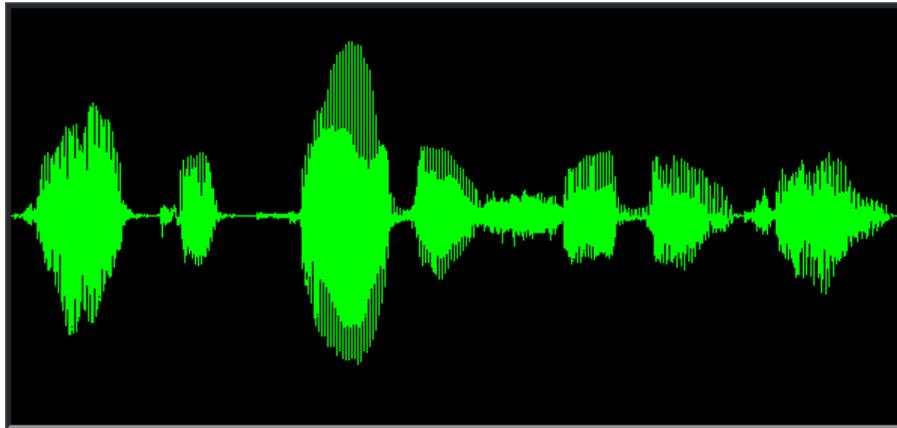


Figure 135: Recording of a male voice saying, "Joe took father's shoe bench out"

The basic idea is that most of the information in speech is carried in these low-frequency modulations, and anything that reduces the depth of the modulations must negatively affect speech intelligibility. The real advantage of this approach is that STI ends up being sensitive to just about any factor that works to degrade speech intelligibility in a sound system and/or a room, including noise, excessive reverberation, distortion and audible echoes.

The basis for STI is the modulation transfer function (MTF) which quantifies the depth of modulation in the received signal relative to the transmitted signal at specified frequencies. The modulation transfer function can be measured directly using specialized, "speech-like" test signals or calculated indirectly, from the impulse response or ETC of a system under test. In either case, it is measured over a range of seven octaves, from 125 Hz to 8 kHz, at fourteen modulation frequencies per band. The modulation frequencies range from 0.63 Hz to 12.5 Hz in 1/3-octave intervals.

The current IEC standard on STI (60268-16, Edition 4.0 2011-06) says that when measuring STI indirectly from an impulse response, *"The duration of the impulse response shall not be less than half the reverberation time and at least 1,6 s to ensure a reliable calculation of the modulation indices for the lowest modulation frequency of 0,63 Hz."* This, however, seems to overlook the fact that frequency bins in a DFT are linearly spaced, meaning that the two lowest frequency bins are always an octave apart, whereas the modulation frequencies for STI are on 1/3-octave intervals. If the DFT time window were 1.6 seconds then the two lowest bins would be at 0.63 and 1.25 Hz, whereas the first three STI modulation frequencies are at 0.63, 0.8 and 1.0 Hz.

This is to say that hitting all of the lowest STI modulation frequencies would require a measurement window significantly longer than 1.6 seconds. We experimented and found that a measurement time constant of 5 seconds produces data points very close to all of the STI modulation frequencies. That is the reason for the 5000 ms DFT size for IR measurements. (The exact DFT size in samples depends on the selected sampling rate, for example at 48k sample rate, 5000 ms works out to 240,000 samples.)

IEC 60268-16 qualifies IR-based STI measurement techniques for "noise-free" measurements, where a minimum signal-to-noise ratio of 20 dB in all seven octave bands is obtainable. The standard specifically qualifies MLS and swept sine test signals for used with indirect measurement techniques but also states

that, “Theoretically, other mathematically deterministic pseudo-noise (random phase) signals could be employed.” Period matched pseudorandom noise in Smart fits that description and works well for STI.

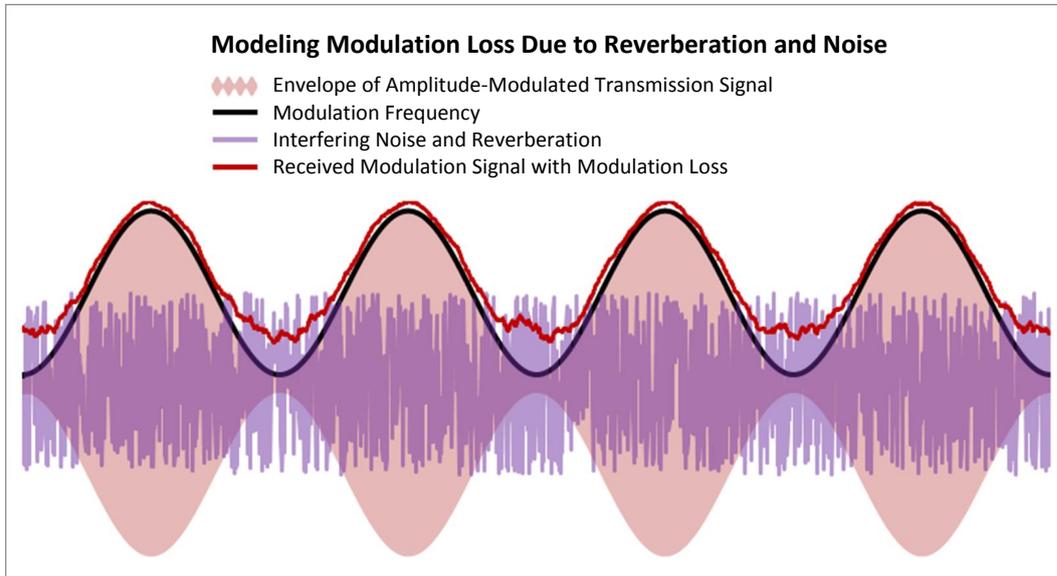


Figure 136: STI estimates speech intelligibility through a transmission channel as a function of modulation loss. The black line in the figure above is the modulation signal for a transmitted signal. The red line is the modulation signal of the received transmission. The difference between the two (the modulation transfer function or MTF), quantifies loss of modulation due to factors such as reverberation and noise.

There are a couple of potential advantages to using period-matched noise over sweeps. One is that it's important to conduct STI measurements at sound levels that reflect actual use of the system under test, and the absolute level of sweep signals that start and stop can be ambiguous. Another is the issue of distortion products potentially piling up at the end of the time record FFT-based IR measurements made with logarithmic sweep signals. Those need to be dealt with in some fashion, either cropped off or windowed out, lest they corrupt the measurement. That isn't an issue with period matched noise although you still need to take care not to overdrive the SUT.

General caveats to using STI are that it is sensitive to strongly fluctuating background noise levels, which can lead to overestimation of low intelligibility systems or underestimation of scores on the high end. When measuring in the presence of fluctuating background noise, at least three measurements should be taken and their results averaged to reduce measurement uncertainty. Also, if the speech source and some prominent source of interfering background noise are widely separated, STI may underestimate intelligibility – human hearing can be smarter than machines about that kind of thing. STI is also sensitive to clipping or amplitude compression in the transmission channel, but in our case, those would also violate the linear time-invariant system rule for transfer function measurements. So don't do that.

Qualitative Thresholds for STI and STIPA

In terms of specific levels to look for, anything better than 0.80 STI is considered excellent. STI scores between 0.61 and 0.80 are “Good,” 0.45-0.60 is “Fair,” 0.35 to 0.44 is “Poor” and anything less than 0.35

is atrocious. Note that there are male and female versions of STI in the more recent versions of the standard. Any time an STI number is stated without specifying whether it's for a male or female speaker, the assumption is that it's the male version.

STIPA

One of the problems with direct measurement of STI is that it takes a lot of time to make a measurement. The modulation frequencies are so closely spaced that each one had to be measured separately and there are 98 modulation frequencies in all (14 x 7). The full direct measurement takes about 15 minutes to perform as a result. STI for Public Address systems (STIPA) was developed to get around this problem.

STIPA is essentially the same measurement as STI, but uses a subset of its modulation frequencies; two per octave, for a total of 14. STIPA is typically measured directly, using a special test signal that excites all 14 frequencies at the same time, so that the measurement can be completed in a single pass. STIPA measurements can be completed in a few seconds and have been found to correlate very well with the more rigorous, full STI. STIPA is currently validated only for male speakers.

In Smart, of course, we measure STI indirectly from the impulse response and the full STI measurement takes no longer to perform than a typical direct measurement of STIPA. Smart does provide figures for both STI and STIPA, however STIPA in our case is more properly termed STIPA(IR), since it is based on IR data rather than measured directly. It is provided for informational purposes, e.g., to facilitate comparison with readings from hand-held STIPA meters and is literally just a subset of the full STI measurement, calculated from exactly the same measurement data.

ALCons

ALCons, sometimes called %ALCons because it is stated as a percentage, stands for Articulation Loss of Consonants. Consonant sounds are critical to speech intelligibility because they are short in duration and tend to get lost more easily than vowel sounds that are voiced over a longer period of time and have more total energy. ALCons was originally conceived as an estimate based on distance of the listener from a sound source, room volume and reverberation time of the room. This is commonly called the "architectural" form of ALCons, due to its reliance on room dimensions and distance. Later forms of the calculation using direct-to-reverberant ratio in place of volume and distance made ALCons suitable for direct measurement from an impulse response and found their way into a number of acoustical measurement platforms including Smart.

The later forms use an early-to-late energy ratio with a relatively short split time – typically in the range of 10-15 milliseconds – to estimate direct-to-reverberant ratio. The earlier of the two most commonly used directly measurable forms did not take background noise into account, making it suitable only for cases where noise was not significant factor affecting speech intelligibility. A later version that does account for noise is informally called "long form" ALCons, to differentiate it from the earlier, "short form" calculation. The directly measurable forms of ALCons can be calculated for any frequency range but ALCons is regarded being most meaningful in the octave band centered on 2 kHz, as this is where most of the energy in consonant sounds is found.

Chapter 9: Analyzing Impulse Response Data

Advantages of ALCons include the fact that in its original “architectural” form, it made estimation of speech intelligibility possible for sound system designs not yet installed, in rooms not yet been built, without the aid of acoustic modeling programs not yet available in the 1970s and ‘80s. In its later form, it could be directly measured in existing installations based on ETC or IR data produced by TDS and MLS/FHT-based measurement systems that were prevalent before the computing power needed to calculate FFTs large enough for room acoustics work became widely accessible. The main disadvantages of ALCons are its reliance on assumptions that the sound field is statistically well behaved and without audible echoes and also, in the case of electroacoustical systems, that the system is free of audible distortion that could affect intelligibility.

Qualitative Thresholds for ALCons

ALCons is upside-down relative to other intelligibility metrics. Smaller numbers mean better scores. Anything less than 5% is considered excellent. Between 5-10% is “Good”, 10-15% is rated as “Fair” and anything more than 15% is problematic.

Chapter 10: Measuring an Acoustical Impulse Response

You can make an impulse response measurement in Smaart by clicking a single button. Whether or not that measurement returns any useful information depends on decisions made beforehand. The process for measuring an acoustical impulse response can be summarized as follows:

- Selection of measurement technique and stimulus type.
- Selection of excitation source(s) and position(s).
- Selection of measurement position(s).
- Estimation of the reverberation time and background noise.
- Selection of measurement parameters (measurement length/duration, excitation level)
- Exciting the system and recording results

What are we measuring, and why?

Before you set out to make any acoustical measurement, it is always helpful to define your objectives clearly. The trip back to the site that you save may be your own. Obviously we want to measure the acoustical impulse response of a system under test (SUT) for some reason, but what exactly is the system? Is it a room? Is it a sound system? Is it a combination of a sound system its acoustical environment? What do you want to know about the system? What equipment and measurements will be needed to make sure you get the information you need?

If you want to measure the reverberation time of a room with an installed sound system, are you more interested in the room or the system? Consider that using a directional loudspeaker to excite the space may affect reverberation times in locations that are on axis with the speaker. Consider also that when using different speakers to measure from different points in the room, any significant differences between those speakers will show up in your measurement results. If the room is your target, the course of least resistance might be to bring in an omnidirectional loudspeaker specifically designed for acoustical measurement. On the other hand, if your objective is to measure the performance of a loudspeaker system installed in a room, you might be more concerned with early-to-late energy ratios and speech intelligibility metrics than reverberation time of the room, exclusive of the sound system.

Direct vs Indirect IR measurement

There are two basic ways to measure an impulse response in Smaart; direct or indirect. Or you could say there are three possible methods, because the indirect IR measurement method that Smaart employs can be used as a deterministic or non-deterministic measurement technique. Let's start with the simplest and easiest to understand – old fashioned direct IR measurement – and work our way up from there. This should not be construed as being in any sort of order of preference. In fact, we're starting with what is usually the least preferred method, but they all have their selling points.

Direct IR Measurement Using an Impulsive Stimulus

The most intuitive way to measure the impulse response of a system would be to use an impulsive stimulus of some kind and simply record what happens. And in fact, people have been doing just that for decades. The advantages are that you do not need a sound system or even a measurement system to perform the measurement. All you really need is some way to make a loud bang and some way to record it. The main problem with this approach – other than the fact that it doesn't tell you anything about an installed sound system (if applicable) – is the scarcity of really good impulsive stimulus sources.

An ideal impulsive stimulus would be a perfectly instantaneous, perfectly omnidirectional burst of energy, containing equal proportions of energy at all audible frequencies. In the time domain, it would appear as a single vertical spike no more than one sample in width. In the frequency domain, it would produce a perfectly flat magnitude and phase trace. When evaluating the response of your system under test (SUT) to this ideal stimulus, you could then confidently assume that anything you saw in the IR that wasn't an instantaneous spike, or anything in the frequency domain that wasn't a flat line, must be the response of your system.

If you loaded the wave file `1samplePulse.wav` to look at the frequency response of Smaart's bandpass filters in the previous chapter, you were actually performing a direct IR measurement on the filters using an ideal impulse. Unfortunately, stimulus signals like that do not exist in the physical world. When we need to measure the impulse response of an acoustical system directly, we end up using stimulus sources that are less than ideal. Blank pistols and balloon pops are common sources. Signal cannon, spark gaps, fireworks and even spot welders have been used. The problem with all of these is that their spectral content is not uniform, their envelopes are not instantaneous, they may not really be as omnidirectional as one might guess, and all of these factors will vary to some extent from one measurement to the next. This introduces uncertainty from the start as to which part of the completed measurement is stimulus and which is response. It also limits the repeatability of test results. For this reason, systems such as Smaart that indirectly infer the response of a system to an ideal impulse have become more the tools of choice these days.

Indirect (Dual Channel) IR Measurement

Indirect impulse response measurements are made using dual-channel measurement techniques that mathematically estimate the response of an SUT using continuous or periodic test signals. Three of the four indirect IR measurement methods we can name require specialized test signals such as sweeps (Time Delay Spectrometry and Direct Convolution) or specialized noise (Hadamard Transform/MLS).

The dual-channel transfer function method that Smaart uses for indirect IR measurement also works best using period matched test signals, but unlike the other three it can also produce very acceptable results using random test signals, provided that both the reference and measurement signals are captured. Transfer function-based IR measurement systems work by calculating the frequency-domain transfer function of a system under test (SUT) from the Fourier Transforms of two signals – the signal going into a system and the output of the system in response to this input – and then transforming the result back into the time domain using an inverse Fourier transform (IFT).

Remember that perfect impulse that we were lamenting didn't exist in the real world for direct IR measurements? Well, that happens to be what you get if you take the IFT of the transfer function of two identical signals. So it follows that when we take the transfer function of a stimulus signal and the SUT's response to it, we theoretically should get something very much like its response to an ideal impulse. And in fact that's pretty much what happens in practice when you use a period-matched excitation signal. When you use this same technique with effectively random signals you also get a lot of extra noise, but repeating the measurement several times and averaging the results generally takes care of that, and Smart makes this easy to do.

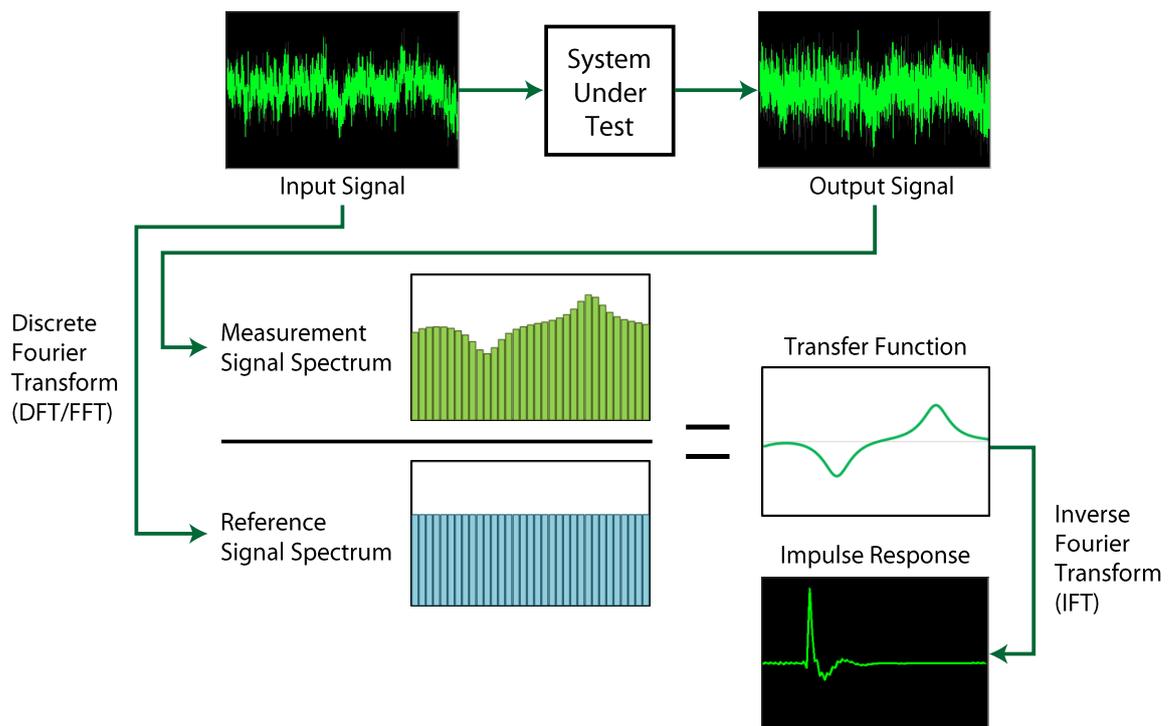


Figure 137: Block diagram of a Dual-FFT transfer function IR measurement

Dual Channel IR Measurement Using Period-Matched Signals

In a way, the discrete Fourier transform (DFT or FFT – all FFTs are DFTs, but not all DFTs are fast) is kind of a dirty mathematical trick. Fourier transforms of all types theoretically work only with signals of infinite length but the DFT gets around this by pretending that a finite chunk of signal being analyzed is really just one instance of an infinitely repeating series of chunks that look exactly like it.

The best way to get around this inherent assumption of this cyclicity in DFT/FFT analysis is to feed the DFT what it really wants to eat: a test signal that either fits completely within the measurement time window or cycles with periodicity equal to the length of the DFT time constant. Signals that meet these criteria can produce deterministic, highly repeatable measurements in a fraction of the time it takes to get comparable results using random signals.

Chapter 10: Measuring an Acoustical Impulse Response

When using matched periodic signals:

- No data window is required.
- Delay compensation is not a critical requirement.
- Considerably less averaging is typically required.
- Measurement time constants can be kept to reasonable lengths.
- Small time variances become less of a concern.
- Subjectivity in selecting measurement parameters is reduced.
- The measurement system doesn't necessarily need to be connected to the system under test. (When using a known test signal, the measurement system and the SUT can get their stimulus/reference signals from two different sources and the measurement will still work. You won't get an accurate propagation delay time without an audio feed from the signal source being used to excite the SUT, but if you don't really need delay times this can be a very handy option.)

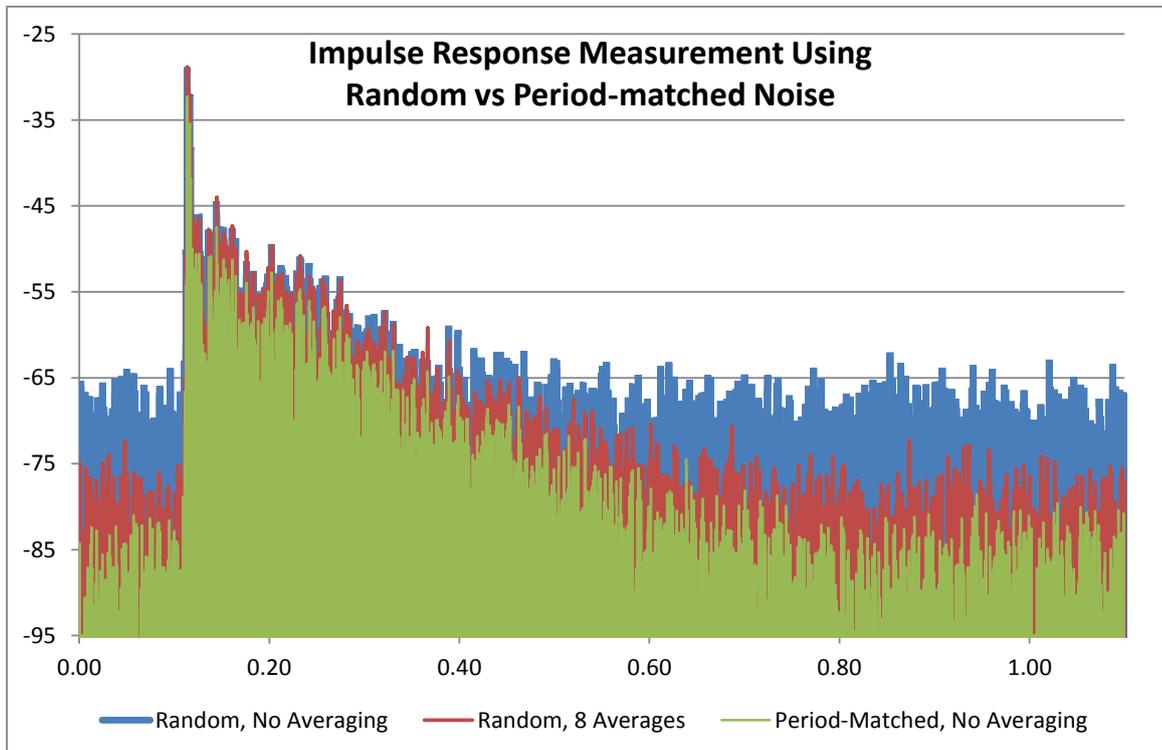


Figure 138: Three indirect IR measurements of the same room, taken from the same microphone position using effectively random noise vs period-matched pseudorandom noise. The period-matched noise measurement (in Green) takes the same amount of time as the unaveraged random measurement (in Blue) but has much better dynamic range. By repeating the random noise measurement eight times and averaging the results (the measurement in Red) we can greatly improve its signal-to-noise ratio, however the measurement takes eight times longer to perform.

Logarithmic Sweeps

Logarithmic sweeps are called Pink Sweeps in Smaart. When you select this signal type in the signal generator, Smaart drops the IR data window without being told. A data window in conjunction with a sweep signal would act as a filter on its frequency content, since each frequency appears at only a single point in time during the measurement.

Sweeps can be used as a circular or aperiodic signal source. If the Triggered by impulse response option is enabled in Smaart's signal generator, the sweep signal is triggered by starting an IR measurement. When you kick off the measurement, Smaart will insert a short period of silence before the sweep in case there's any lag in starting the recording device, then run the sweep and insert another period of silence afterward to let the SUT ring out. If the Triggered by impulse response option is unchecked, the sweep runs continuously when the generator is turned on. In this case you would start the generator before starting the measurement as you would with other test signals.

A peculiarity of dual-FFT-based IR measurements made with logarithmic sweeps is that distortion products in the excitation loudspeaker/SUT are washed out of the IR and show up as "pre-arrivals". Because the DFT is a circular function, these typically end up wrapped around past the beginning of the measurement and pile up near the end of the time record. The practical implication is that you may need to make the measurement time window a little larger than you would for a matched noise measurement, to ensure that these artifacts do not intrude on the reverberant decay slope.

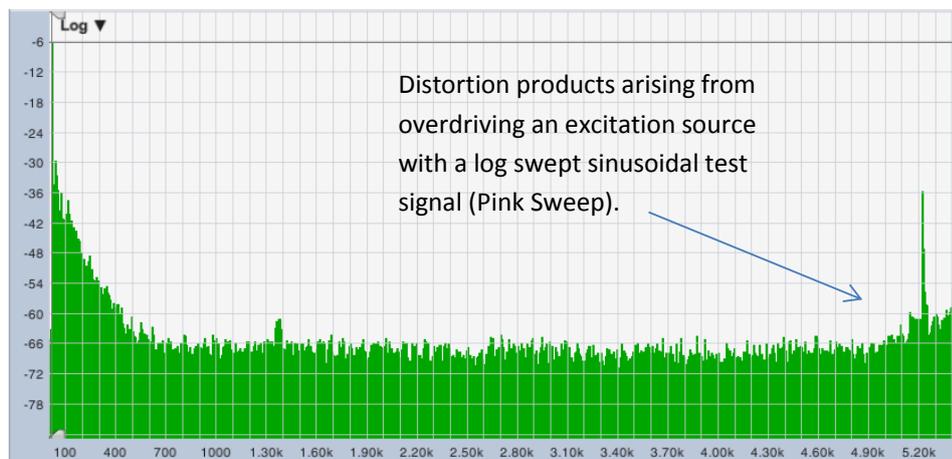


Figure 139: An impulse response measured using a log swept sine signal (Pink Sweep), showing harmonic distortion products from the excitation loudspeaker piled up at the end of the time record. In this case, the speaker being used to excite the room was overdriven and the distortion component was quite significant.

As regards STI measurement, IEC 60628 states that, "When using a sine sweep technique, the distortion components that are inherent within the method shall be edited out or removed from the IR before calculation of the STI can be undertaken." It is our opinion however, that this requirement argues for the use of period-matched noise rather than sweeps for STI measurement, since the masking effect of high levels of distortion in lower-spec announcement systems can significantly affect speech intelligibility, and properly should be included in the measurement.

Dual Channel IR Measurement Using Random Stimulus Signals

An excitation signal that is not completely contained within or, if continuous, has its periodicity precisely matched to the time constant of a discrete Fourier transform being used for analysis is effectively random as far as the DFT is concerned. In Smaart's signal generator, the Random pink noise option or any pseudorandom cycle length with periodicity longer than the FFT size used are effectively random. (Periods shorter than the FFT size should never be used because they won't contain energy at all FFT bin frequencies.) Other examples would include music or noise signals with arbitrary periodicity from sources other than Smaart.

Perhaps the best argument in favor of using random signals for IR measurement is, because you can. If you want to make a measurement with music instead of noise, you can. If it's easier to generate pink noise from a mixing board or in a processor than it would be to inject a test signal into the signal chain from Smaart, that will work. The only absolute requirements are that the measurement system needs to capture an exact copy of the signal going into the SUT and that signal must contain enough energy at all frequencies of interest to you to make a solid measurement.

The main caveats associated with random stimulus signals are a relatively high level of noise, meaning that you have to measure over a longer period than you would with a purpose-built signal to get comparable results. It is left up to the operator to decide how much averaging or how long a time window to use and the actual dynamic range of the SUT is ambiguous. This can be a critical factor in speech intelligibility measurements.

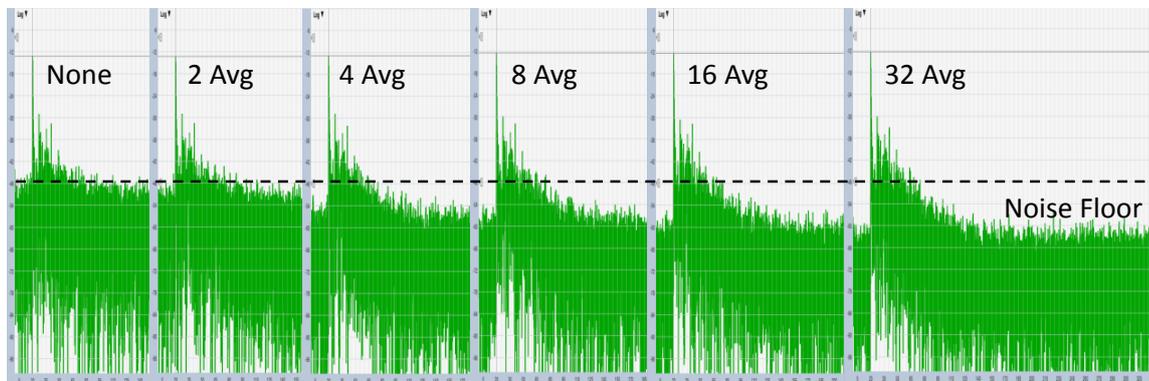


Figure 140: The effect of averaging on an IR measurement made using a random stimulus signal. In theory, each doubling of the number of averages increases signal-to noise ratio by 3 dB.

Reducing Noise in IR Measurements Made Using Random Stimulus Signals

There are three basic things you can do to improve the dynamic range of measurements made using random test signals. The first is to delay the reference signal to match the timing of the measurement signal, so that the data windows line up. You should always do this when measuring with random signals. The second is to evaluate the system over a longer period of time by increasing the DFT size or by averaging multiple measurements (or both). The third is to simply measure louder, which also applies to deterministic and direct IR measurements – in that case you're increasing the signal-to-noise ratio of the measurement by increasing the level of the actual signal, rather than statistically.

Averaging works by inducing regression to the mean in random components of the IR (that is to say, the noisy part). Let's say you take a signal – any signal, maybe an impulse response – and mix it with random noise. Obviously, you get a noisy signal. There is no way to tell just by looking which part was signal and which was noise. But if you take several copies of the same signal and mix each one with different noise, then average all of them together; the noise component of each noisy signal (being random, and different in each case) should start to average toward zero – the theoretical arithmetic mean for random audio noise – while the signal parts (being the same in every case) should average out to themselves.

Of course, all of this depends on the assumption that the signal part of the signal is the same in every case. When working indoors that should generally be a safe assumption. After all, we are working with what we assume to be linear, time-invariant systems in a fairly controlled environment where the worst that could probably happen from one pass to the next is a blast of hot or cold air from an HVAC system causing a slight change to the speed of sound. It might be a larger concern if you needed to make an IR measurement outdoors under windy conditions for some reason. In any scenario where there might be a possibility of any significant time variance during the measurement period, you would probably be better off increasing the measurement time window and/or using a period-matched stimulus signal rather than upping the number of averages.

In theory, averaging two IR measurements or doubling the FFT size used for a single IR measurement should improve signal-to-noise ratio of the measurement by 3 dB. Note that both result in doubling the measurement time, which is really the key to the whole thing. Each additional doubling (2, 4, 8, 16...) of the measurement duration should theoretically get you another 3 dB, although in practice you might reach a point of diminishing returns at some point.

Selecting Excitation Sources and Positions

Excitation source positions should be places that sound would normally emanate from when the system under test is in service. If the loudspeakers you are using to excite the room are the places that sound normally comes from then you've got that part covered. Otherwise an omnidirectional sound source of some kind should be placed on the stage, podium, lectern, pulpit or whatever location(s) that would best simulate normal use of the room/system, and at an appropriate height.

Directional Loudspeakers and Reverberation Time

For the specific purposes of reverberation time measurement, a potential complicating factor may arise if an installed sound system is to be used to excite the room. Impulse response measurements made with directional loudspeakers typically have higher direct to reverberant ratios than IR measurements made by other methods, which could affect reverberation times in the upper octaves. ISO 3382-1 unequivocally states that “the sound source shall be as close to omnidirectional as possible” and provides criteria for assessing the omnidirectionality of a prospective source. ISO 3382-2 specifies measurement procedures for three levels of accuracy in reverberation time measurements: Survey (quick and dirty), Engineering (pretty good) and Precision (very good). For the Precision method, the requirements for the excitation source are identical to those specified in 3382-1 but 3382-2 goes on to say that “For the survey and engineering measurements, there are no specific requirements for the directivity.”

Chapter 10: Measuring an Acoustical Impulse Response

Clearly, the spirit of the law is that omnidirectional sources are preferred for reverberation time measurement, but the standard does leave a little wiggle room when measuring in “ordinary rooms” (as opposed to formal musical performance spaces). Given a choice between a measurement made with a directional speaker or not making a measurement at all, a less than ideal measurement is usually better than none. However if any potential errors or subjectivity in evaluating reverberation time are a source of great concern for you, it might be necessary to bring in an omnidirectional measurement speaker and do it by the book – or at least record a few balloon pops, just to have a second opinion.

Reverberation times aside, it is worth mentioning that for most other purposes, IR measurements made using an installed sound system that is actually used for amplified performances in the space you are measuring will be more representative of actual use of the system than measurements made by any other means. In some cases, in order to get everything you need, you might need to make one set of measurements using an omnidirectional source positioned on the stage, another using the installed sound system and perhaps even a third using the house paging system, to estimate its intelligibility.

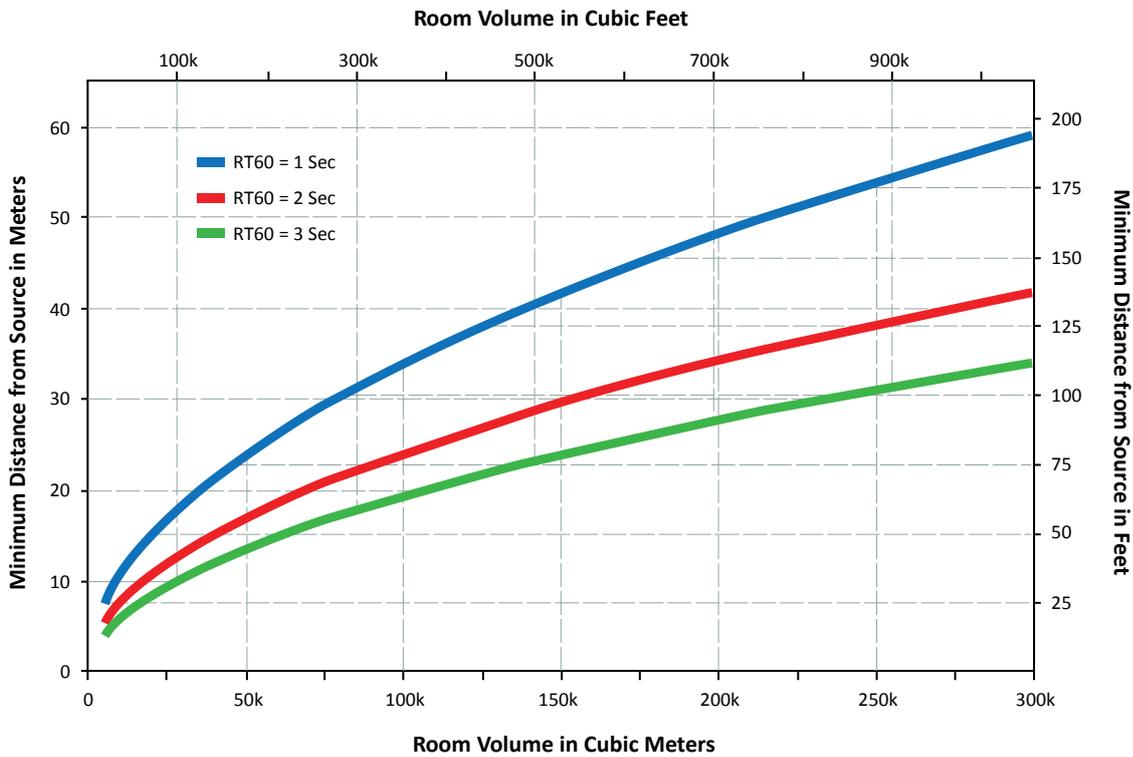


Figure 141: Minimum distance to any measurement position from the excitation source (e.g., a loudspeaker) used for reverberation time measurements. The minimum distance is a function of room volume, estimated reverberation time and the speed of sound, as described by Eq. 1 above. This example uses speed of sound at 20° C (68° F); i.e., 343.6 meters/sec or 1127.4 fps.

Minimum Distance from Sound Sources

Another general requirement for measurement positions used for room IR measurements is that they need to be located far enough away from loudspeakers or other sound sources being used to excite the room to ensure that the measurement is not unduly dominated by direct sound. ISO 3382-2 provides the formula shown in the equation below for calculating the minimum distance (d_{\min}) for any measurement position from an excitation source. Figure 141 provides a graph of this relationship.

$$d_{\min} = 2\sqrt{V/c\hat{T}}$$

where

V is the volume of the room in cubic meters

c is the speed of sound in meters/sec

\hat{T} is estimated reverberation time in seconds

Selection of Measurement Positions

The first and most obvious rule for selecting measurement positions is that you generally want to measure from places where one would expect to find listeners when the system under test is in service. If a tree falls in the forest and no one is there to hear it, who really cares if it makes a sound? You might also want to give special attention to any areas where you think there could be problems. Other than that, it is kind of like taking an opinion poll. If we measure from a single position, we have one “opinion” of what the room sounds like. If we sample from a several different locations, we might reasonably expect to see some consensus emerge as to the most common characteristics of the system response, and for position-dependent differences to begin to average out. The more measurement positions, the lower the theoretical margin of error, assuming the positions are chosen so as to be statistically valid.

For the Survey method in ISO 3382-2, a single stimulus source location is measured from at least two measurement locations, providing a theoretical margin of error of $\pm 10\%$ for octave bands. The Engineering method calls for at least two stimulus source positions and six independent source-microphone combinations for a nominal accuracy $\pm 5\%$ for octave bands or $\pm 10\%$ in 1/3-octave bands. The precision method calls for 12 independent source-microphone combinations using at least two different stimulus source locations and reduces measurement uncertainty to no more than $\pm 2.5\%$ for octave bands and $\pm 5\%$ for 1/3-octave bands.

ISO 3382-2 specifies that all measurement positions should be at least one-half wavelength apart and at least one-quarter wavelength from any reflecting surface including the floor. For example, if we wanted to measure as low as the 125 Hz octave band, the lower band edge is at ~ 90 Hz. At 68° F (20° C), the speed of sound in air is 1127.4 feet per second (343.6 mps) and so one wavelength at 90 Hz would be about 12.5 ft (3.8 m). From that, we could conclude that no two mic positions should be less than 6.25 ft (1.9 m) apart and all microphones should be at least 3.13 ft (0.95 m) above the floor and at least that far from any wall or other reflecting surface. For the 63 Hz band you would need to double those distances.

Of course ISO 3382-1 applies specifically to measurement of reverberation time in rooms, so what about acoustical measurements made for other purposes. Two other standards we could look at as a guide to

Chapter 10: Measuring an Acoustical Impulse Response

microphone placement are ANSI S1.2, Criteria for Evaluating Room Noise, and SMPTE 202M, the current standard for calibrating cinema sound systems. ANSI S1.2 has this to say about measurement positions:

“Sound measurements for rating room noise under this standard shall be made at locations that are near the average normal standing or seated height of human ears in the space: 5’-6” for standing and 4’-0” for seated adults – 3’-6” standing and 2’-6” for seated children. The microphone shall be no closer than 2’-0” from any sound reflecting surface or 4’-0” from the intersection of two intersecting reflecting surfaces, or 8’-0” from the intersection of three intersecting reflecting surfaces.”

SMPTE 202M recommends that microphones be placed:

“In indoor theaters, at position S [...] and position R [...] should it exist, and at a sufficient number of other positions to reduce the standard deviation of measured position-to-position response to less than 3 dB, which will typically be achieved with four positions. [...] It is recommended that measurements be made at a normal seated ear height between 1.0 m and 1.2 m (3.3 ft and 4.0 ft), but not closer than 150 mm (6 in) from the top of a seat, and not closer than 1.5 m (4.9 ft) to any wall and 5.0 m (16.4 ft) from the loudspeaker(s).”

(Position “S” generally works out to be a little left or right of the approximate center of the room on the main floor. Position R is for balconies.)

We can see that these are all in general agreement (depending to some extent on frequency) even though one talking about reverberation time, another is for background noise surveys, and the third is for RTA measurements of cinema sound system. They are probably also not out of line with positions you would intuitively choose for frequency-domain transfer function measurements of a sound system.

Selecting Measurement Parameters

Once you have done all the groundwork of determining your source and measurement positions and figuring out which measurement technique to use, the part of the measurement procedure that directly involves Smaart is actually pretty easy. Basically, you just need to select your measurement parameters, turn on the signal generator (or other stimulus signal source) and kick off the measurement. The two main things you need to concern yourself with at that point are the stimulus level and the measurement duration, which will be some combination of the FFT size and the number of averages.

Input source

If you already have one or more Transfer Function measurements configured and will be using one of those to make your measurement, use the *Group* and *TF Pair* selectors shown in to select the one you want. To create a new TF pair, click the little hammer and wrench button next to the *Group* selector to open the *Measurement Config* window, then click the *New TF Measurement* button (see Figure 142). This pops up another dialog where you can select the input device and channels that you want to use and give the new measurement pair a name. If you are unfamiliar with how to set up your measurement system for transfer function and dual-channel measurements, *Appendix E* has example setup diagrams.

Excitation Level

The rule of thumb for setting the excitation level for IR measurements is that you would like to be able to get at least 40-50 dB above the background noise level. In reverberation time measurements, we evaluate reverberant decay over a range starting 5 dB down from the arrival of direct sound (normally the highest peak in the IR) and extending down another 20 or 30 dB from the start point.

A 30 dB range is preferred but 20 dB is OK if you can't get 30. Either way, the lower end of the range needs to be at least 10 dB above the noise floor of the IR measurement. When you add that all up, you're looking for a minimum of 45 dB of dynamic range for a 30 dB evaluation range, and at least 35 dB for a 20 dB range – and that's in a perfect world, with no noise artifacts from the measurement process itself. In the real world, adding another 5 to 10 dB on top of that would be a definite nice-to-have (unless of course that would drive the system into distortion or blow something up).

To figure out how loud you need to be, you can simply measure the background noise level. We are looking for a relative relationship so you don't even really need to be calibrated for SPL (unless perhaps you plan on doing an STI measurement). Just set the sound level meter in Smaart to *Slow SPL* and watch the meter for ten or twenty seconds with no output signal running to get a feel for the baseline noise level, then start the signal generator at a low level and gradually increase the gain until you reach the target excitation level (or as close to it as you can reasonably get).

Input Levels

Once you have your output levels nailed down, adjust your input levels (by whatever means available) until both the measurement and reference signal levels (labeled M and R on the Control Bar) are roughly even and running at a reasonable level. The yellow segment of the meters in Smaart runs from -12 dB to -6 dB full scale and that's the target zone. The meters are peak reading and we hard limit peaks for noise signals in Smaart, but you have to also allow for fluctuations due to background noise in acoustical measurements and if you use noise from another source you may see wider variations in the peak levels. With sinusoidal sweeps you can run the levels a little higher if you like, due to the lower crest factor of the signal but you always want to keep the levels out of the red. If you are doing a single channel measurement, you will probably have to waste a few balloons or fire off a few blank cartridges while adjusting the measurement channel gain to get a good solid signal level with no clipping on the input level meter.



Figure 142: The measurement signal level (M) is running at a comfortable level. The reference (R) channel is clipping

Measurement Duration (Time Window)

For dual-channel IR measurements, the time window is a function of FFT size (see Figure 143). If you are only concerned with measuring delay times then a good rule of thumb for how long the measurement needs to be is 3 times the longest delay time that you want to measure. If you want to measure reverberation time and early to late energy ratios, then the 60 dB decay time (RT60) of the system is a good target. This is kind of a functional requirement for period-matched dual-channel measurements but it's also a pretty good practical target regardless of how you are measuring. Ideally you would like to

Chapter 10: Measuring an Acoustical Impulse Response

measure 30 dB of reverberant decay and the lower end of the evaluation range should be at least 10 dB above the noise floor, so that's 40 dB which of course is two thirds of 60. By the time you factor in propagation delay, early decay and maybe enough of a noise tail to see the dynamic range of the measurement, chances are you have eaten well into the remaining third.

Of course, both of these rules require either knowing the delay time or RT60 before you've measured them. That typically means you have to guess, then measure, then possibly adjust your guess and measure again. For delay times you can use the distance to the source divided by the speed of sound as a starting point. For "guesstimating" purposes, you can use 1130 feet or 345 meters per second at typical room temperatures – the speed of sound increases with temperature so if it is very hot where you are working you might adjust your estimate upward a little, or downward if it's cold.

For reverberation times, one to two seconds should at least get you in the ballpark for most theaters and auditoriums. Stadiums and other large structures can have much longer reverb times. There's never any harm in measuring over too long a period, so you may want to err to the high side. If you make a preliminary measurement and you are happy with the results you might even be done. If not, you can adjust accordingly and measure again. Note that as a rule, lower frequencies tend to decay more slowly than highs, meaning that the limiting factor may be the reverberation times in the lowest octaves that your stimulus source can excite. So be sure to check the lower bands when estimating reverb times.

FFT Size (Time window)

For dual-channel measurements the duration of the measurement is determined by the FFT time constant – that is, the time required to record enough samples for a given FFT size at whatever sampling rate you are using. In Impulse mode, Smart gives you the time constant in milliseconds, along with the frame size in samples for each available FFT/DFT size.

Averaging and Overlap

Averaging, as we discussed earlier in this chapter, is primarily something you concern yourself with when using effectively random stimulus signals. With random or effectively random stimulus signals, deciding how much is averaging is enough is kind of judgment call but typical settings are in the 4-16 range. In very noisy environments, you may want to use a larger value and/or consider using a period matched signal. When measuring with period-matched noise or sweeps, averaging is normally set to "None" or 2, although it is still possible that a higher setting could prove helpful if measuring in an extremely noisy environment.

Another factor that affects how averaging works is the Overlap setting found in Impulse Response options (Options menu > Impulse Response). When overlap is set to 0% each FFT is calculated from unique data, giving you the maximum amount of noise reduction that you can get from a given number

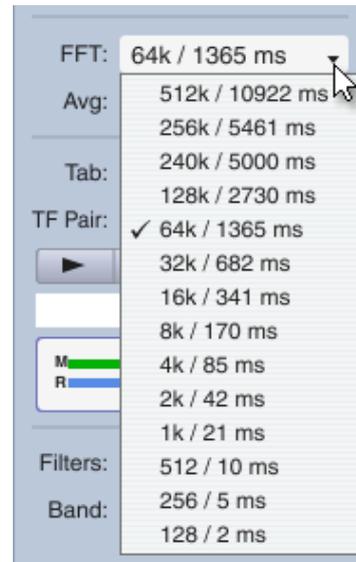


Figure 143: IR Mode FFT size selector showing the time constant in milliseconds for each FFT size

of averages. When you set the measurement *Overlap* to a non-zero value then successive FFTs share some data in common (see Figure 131 on page 158) – remember that *measurement* overlap and *spectrograph* overlap are two different things, but the principle is the same. If measurement overlap is set to 50%, it only takes a little longer to record 16 averages than it would for 8 at 0% overlap. You don't get the full benefit of averaging 16 unique FFTs in that case and processing time increases but you should see at least a little better signal-to-noise than you would get 8 with some net time savings.

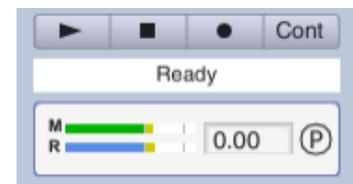
Delay Compensation

When making IR measurements with random signal sources you will get much better results if you compensate for the delay time through the system under test. So plan on making the measurement twice if you don't already know the delay time; once to find the delay and a second time for a keeper. The button labeled “@” (for Peak) that appears next to the input level meters in IR mode (see Figure 142) sets the reference signal delay to the highest peak in the impulse response.

Pushing the Button and Making the Measurement

Having nailed down which measurement technique will get you the results you need, chosen your excitation sources and measurement positions, set your input and output levels, and selected your FFT length and number of averages (if applicable), all that's left to do is push the button(s). For a dual-channel measurement, start your excitation signal (unless you're using a triggered sweep in Smaart, in which case it will start by itself) and click the start (▶) button in the Control Bar on the right side of the main window. Smaart will take it from there and display the measurement results when it has finished.

For a single-channel (direct IR) measurement, click the record (●) button, then click the start (▶) button, pop your balloon or fire your blank pistol (or whatever), give the system a few seconds to ring out, and then click the stop button (■) to end the recording and display the results.



Saving Your Work

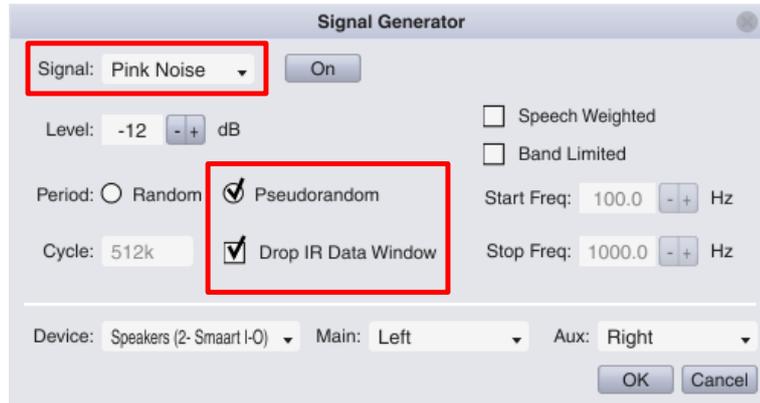
When you measure an IR in Smaart does not automatically assume that you want to save it. Sometimes it takes a few tries to get it right and we don't need a lot of old baggage piling up in the process. When you have a measurement that you are happy with and want to save, click in the File menu and select Save Impulse Response then select the directory where you want to store the file and give it a name. Note that if you have cropped the file for display purposes using the Crop function, only the displayed portion of the time record is written to file.

Recap: Common Settings for Dual-channel IR Measurements

The following are some common “go-to” default settings for IR measurements in Smaart that should typically work well for most rooms.

Signal Type

If you are able to use Smaart’s signal generator as your stimulus signal source then period matched pseudorandom noise is a good all-around choice for signal type. To turn on this option, open the signal generator control panel, select Pink Noise as the signal type, then tick the boxes labeled *Pseudorandom* and *Drop IR Data Window*.



FFT size and Averaging

128K makes a good default setting for FFT size. At 48k sampling rate, that gives you almost 3 seconds of time window. You would generally have to be measuring in a pretty huge space to need more than that, but it’s not ridiculously long for smaller venues. For averages, a setting of 2 is a good when using period-matched noise. If you have to use a random signal source for some reason, up the number of averages to 8, or maybe 16 if you are working in a noisier environment. We are assuming 0% overlap for averaging.



Excitation Level

If you need to measure reverberation time, then your excitation level needs to be a minimum of 45 dB above the background noise level for T30 (preferred) or at least 35 dB above to get T20. For most other purposes, any excitation level that is comfortably above the background level should be fine.

Input Levels

When using random or pseudorandom noise signals, -12 to -15 dB or so is the preferred input level for any kind of measurement in Smaart including IR measurements. -12 dB is the point where the input levels in Smaart turn yellow.

Delay Time

When using period-matched noise as your excitation signal you can set the delay time to 0 if you don’t already have it set for your selected signal pair. If you do, then there’s no harm in leaving it alone. If you are using a random noise source and don’t already know the delay time for your signal pair, run the IR measurement once to find it, then click the “@” button to set it, then run the measurement again.