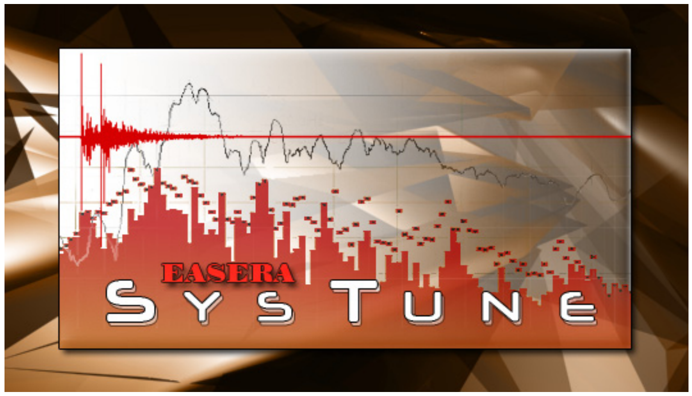

EASERA SysTune

-

Developed by

AFMG Ahnert Feistel Media Group

The creators of EASE and EASERA

www.afmg.eu

Software Manual, Rev. 1, December 2007

Contents

Contents 2
Preface..... 5
SysTune's capabilities include: 6
Equipment Requirements..... 6
Software Support 7
EASERA SysTune Installation and Licencing 8
1. Installation Instructions..... 8
1.1 Microsoft .NET Framework 2.0..... 8
1.2 EASERA SysTune Startup 8
EASERA SysTune..... 8
AFMG Licence Manager 8
1.3 EASERA SysTune User Files..... 8
1.4 Licencing the Software 8
2. Licencing Instructions..... 9
2.1 Online Licencing..... 9
2.2 Licencing by File 9
Reference File 9
Licence File..... 9
2.3 AFMG Licence Manager Program 10
Licence Status 10
Online Licencing..... 11
Licencing by File – Licence Tab 12
Licencing by File – Terminate Tab..... 13
Licencing by File – Import / Export Tab 14
Proxy Server Configuration Window 15
Program Tutorial..... 16
1. Introduction..... 16
Preface..... 16
Starting the Software..... 16
Screen Layout 17
2. Measurements with a Single Input Channel 19
Selecting the Input Channel..... 19
2.1. Time Signal..... 20
Calibrating an Input Channel..... 21
Adjusting the View Limits..... 25
Summary..... 27
2.2. Input Spectrum..... 28
Choosing the FFT Size..... 32
Averaging over Time 33
Checking the Calibration 34
Summary..... 36
2.3. Spectrogram..... 36
Summary..... 39
3. Measurements with an Excitation Signal..... 40
3.1. Excitation Signal..... 40

Choosing a Stimulus Signal 40
Activating the Output Channel 43
Frequency Response Measurements 43
Summary 44
3.2. Capturing and Comparing Measurements 44
Overlay Properties 46
Active Curve 47
Saving, Removing and Loading Overlays 48
Averaging Measurements 49
Adding Cursors 52
Summary 56
4. Dual-FFT Measurements 57
Setup with Internal Reference 57
Setup with External Reference 58
4.1. Transfer Function Measurements 59
Measuring Principles 59
Computation of the Transfer Function 60
Real-Time Deconvolution (RTD™) 61
Transfer Function in EASERA SysTune 62
Example 63
Summary 65
4.2. Impulse Response 65
Trouble-Shooting 67
Impulse Response Measurements 68
Time Alignment of Loudspeakers 71
Saving Impulse Responses 73
Level Meters 73
Summary 75
4.3. ETC 75
Analyzing Reflections 76
Summary 77
4.4. Magnitude 78
Relationship between Transfer Function and Impulse Response 78
Magnitude of the Transfer Function 78
Coherence and IR Stability 81
Adjusting the Gain Offset 83
Averaging Transfer Functions 84
Exporting Graph Data 86
Introduction to Windowing 87
Windowing in SysTune 88
Windowing Options and Options Window 92
TFC Window™ 96
Summary 98
4.5. Phase 98
Meaning of Phase Data 98
Using the Phase Graph For Time Alignment 100
Smoothing and Wrapping Phase 104
5. Further Measurements 106

5.1 Level Measurements 106
Noise Criteria 106
SPL and LEQ Measurements 107
SPL and LEQ Monitor 107
Histo Graph 109
Summary 111
5.2 Measurements using Multiple Signal Channels 111
Changing Soundcard and Driver 112
Status Bar 113
Multi-Channel Measurements 114
5.3 Reverberation Time and Speech Intelligibility 115
Reverberation Time 116
Speech Intelligibility 120
Summary 122
5.4 Measurements using Speech and Music Signals 122
Noise Suppression Tools 124
6. Additional Topics 126
6.1 Menu Structure 126
6.2 Short Cuts 128
6.3 View Limits 129
6.4 Trouble Shooting 130
Graph Reference 131
Graphs [All] 131
Buttons 131
Mouse 131
Graphs [Input] 132
Time Signal 132
Spectrum 133
Spectrogram 134
Graphs [Levels] 135
NC {Noise Criteria} 135
Histo {Histogram} 136
SPL {Sound Pressure Level} 136
LEQ {Level Equivalent} 136
Graphs [Transfer Function] 137
IR {Impulse Response} 137
ETC {Energy Time Curve} 138
Mag {Magnitude} 139
Phs {Phase} 140
Graphs [Results] 141
RT {Reverberation Time} 141
STI {Speech Transmission Index } 142

Preface

EASERA SysTune™ is one of the newest products from AFMG. Being more than just the little sister of the room-acoustic measuring software EASERA, EASERA SysTune complements and expands the feature set of EASERA into the world of real-time analysis.

In particular, SysTune is aimed at all people involved with acoustic measurements and system tuning especially in live sound applications. EASERA SysTune offers patent-pending, novel and exciting features. While frequency displays for input spectrum and transfer function have become an accepted standard, SysTune sets new benchmarks with its real-time impulse response displays and analysis tools.

The worldwide unique Real-Time Deconvolution RTD™ engine of SysTune is most advanced technology. Being able to process IR data of 5 to 10 s length in real-time and at high refresh rates, EASERA SysTune opens the door to perform simple room-acoustic measurements. In fact, SysTune's function to look at reverberation times and speech intelligibility in occupied venues just using a reference signal and a measuring microphone is like a dream come true for many acousticians.

On the other hand, a newly developed TFC™ window (Time-Frequency-Constant window) and simultaneous views of windowed transfer function and impulse response facilitate the investigation of any part of the measured system response. Delay-alignment of sound systems in the (real-)time domain and their equalization in the frequency domain was never easier than with SysTune.

EASERA SysTune also represents a remarkable step forward in its capability to process up to 8 (eight) input channels at the same time in order to offer spatially averaged spectrum and transfer function displays. Cutting-edge Intel libraries, high performance multi-threading technology and hand-crafted compiler optimizations enable SysTune to provide real-time refresh rates even for complex applications.

We hope you will enjoy working with this new measurement tool.

The Team at AFMG.

Ahnert Feistel Media Group
Arkonastr. 45-49
13189 Berlin
Germany

Web: www.afmg.eu, www.easerasystune.com
Email: support@afmg.eu

SysTune's capabilities include:

- 8-channel, 8 kHz to 192 kHz sampling rates
- Real-Time data acquisition & display in both time and frequency domains at high refresh rates using live sound, pink noise, sweeps or other stimulus signals
- Real-Time Deconvolution (RTD™) for analysis of impulse response and complete frequency response based on a signal channel and a reference channel (Dual-Channel FFT)
- Real-Time Impulse Response, Magnitude and Phase displays. Newly developed time-frequency-constant (TFC™) window to investigate early energy arrivals in detail
- Precise real-time spectrogram display for feedback analysis
- Input spectrum and frequency response of up to 8-channels can be averaged (Multi-Channel-FFT)
- Measured data (Impulse Response) results can be easily exported to EASERA and EASERA Pro for additional post processing and in depth analysis
- Live RT and STI calculations instantly
- SPL, LEQ and NC measurements; Level histograms.
- Coherence and IR stability displays allow quick and easy time alignment of loudspeakers using real-time impulse response data
- Cursors and overlays for easier comparison of captured curves
- Integrated signal generator for log-sweep and pink noise stimuli of standard FFT time lengths
- Windows Direct Sound, Wave/MME, ASIO audio drivers; interface to EASERA Gateway; Multi-threaded, full support for multi-processor computers

Equipment Requirements

EASERA SysTune runs under Windows 2000, Windows XP, and Windows Vista operating systems on PC's with a minimum graphics resolution of 960 x 720; 1024 x 768 resolution is preferred. Windows 95, 98 , NT and ME (Millennium) are not supported.

CPU should be at least 1 GHz with support for the Intel SSE instruction set, working memory (RAM) should be at least 256 MB and at least 1 GB or more of free hard disk space should be available.

A soundcard is required. EASERA SysTune supports all common soundcards with up to 8 input channels, bit-resolutions up to 32 Bit and sampling rates of up to 192 kHz. Windows, DirectSound, Wave and ASIO drivers are supported, If more than two input channels will be used, ASIO drivers are required. For one or two input channels Direct Sound (MSDirectX) can be used as well as Wave/MME drivers (MS Windows Audio-API). See also the Audio Device Information Viewer on www.EASERASysTune.com .

For precision measurements an EASERA GATEWAY high performance AD/DA converter/preamp is recommended.

Software Support

If you have questions about operating the software, please search this document and refer to the textbooks and papers listed in the respective chapters. In addition please visit our dedicated EASERA SysTune website www.EASERASysTune.com and the AFMG internet forum www.afmg-network.com as well as the website of your EASERA SysTune distributor:

- Worldwide Distribution by Renkus-Heinz, Inc.: www.renkus-heinz.com
- Distribution in Germany by AudioOne GmbH: www.audioone.de
- Educational version through ADA-Foundation gGmbH: www.ada-foundation.com
- Copyright/Manufacturer: SDA Software Design Ahnert GmbH: www.sda.de

EASERA SysTune Installation and Licencing

1. Installation Instructions

1.1 Microsoft .NET Framework 2.0

Please note that you need to have .NET Framework 2.0 installed before installing EASERA SysTune:
http://www.microsoft.com/downloads/details.aspx?FamilyID=0856eacb-4362-4b0d-8edd-aab15c5e04f5&displaylang=en

1.2 EASERA SysTune Startup

Insert the EASERA SysTune CD into the computer. This will automatically run the Setup file to begin the EASERA SysTune Startup application. Follow the instructions on the screen to install both EASERA SysTune and the AFMG Licence Manager.

EASERA SysTune

This will create a folder (`C:\Program Files\AFMG\EASERA SysTune`) for the program and place an EASERA SysTune icon on the Windows Desktop.

The installer will also automatically create two sample file subdirectories; the directory `\Signals\` contains a selection of excitation signals, the directory `\IRs\` contains a set of impulse responses.

AFMG Licence Manager

This will create a folder (`C:\Program Files\AFMG\AFMG Licence Manager`) for the program and place an AFMG Licence Manager icon on the Windows Desktop.

1.3 EASERA SysTune User Files

Finally, insert the user-specific EASERA SysTune User Files CD into the computer to automatically run the Setup file. Select `Typical` to install with the preferred settings. This will make the software licence available to all users of this computer. Select `Customize` to change this preset to a different location. The installer will then create a folder for the user files and allow the AFMG Licence Manager installed above to register a licence for the software.

1.4 Licencing the Software

Double-Click the AFMG Licence Manager icon on the Windows Desktop to register the program.

2. Licencing Instructions

2.1 Online Licencing

To further improve your comfort we have created a licencing web application to run on our web server. It allows you to easily download an EASERA SysTune licence via internet (being online with your EASERA SysTune computer assumed). This means the software sends the computer's reference information to our web application, which creates licence information on our server. This information is then automatically downloaded and installed. So with a single button push on `Download Licence` you can unlock EASERA SysTune.

By subscribing to the licence agreement you are entitled to install the program on two computers. After those two installations, additional licences must be purchased. Please see your distributor for prices.

If you intend to uninstall EASERA SysTune from one or both of the original computers then please upload the licence information from that computer by clicking on Upload Licence. This will allow you to download this licence again and then unlock EASERA SysTune on a different computer.

2.2 Licencing by File

You should only use this option if you are not able to use the online licencing functions.

Reference File
The *`Reference File`* is a file generated by the AFMG Licence Manager program and placed in the EASERA SysTune *`LicenceFiles`* folder. This file is different for each installation. If you have more than one computer each will have its own *`Reference File`*. To order a licence you must send the *`Reference File`* to SDA (support@afmg.eu) by E-Mail .

Licence File
The *`Licence File`* is a file generated by SDA, which is linked to the *`Reference File`*. The *`Licence File`* is supplied to you by E-Mail. Loading this file into the AFMG Licence Manager with `Install Licence` unlocks the particular EASERA SysTune version purchased.

If you intend to uninstall EASERA SysTune from one or both of the original computers then remove the licence information from that computer before by using the `Termination`. `Terminate Licence` creates a *`Termination File`* which you must send to SDA by E-Mail. This will allow you to order a new licence for the terminated one and unlock EASERA SysTune on a different computer or computers. There can be only two operational programs at the same time without additional licences!

2.3 AFMG Licence Manager Program

This program allows the licencing of EASERA SysTune.

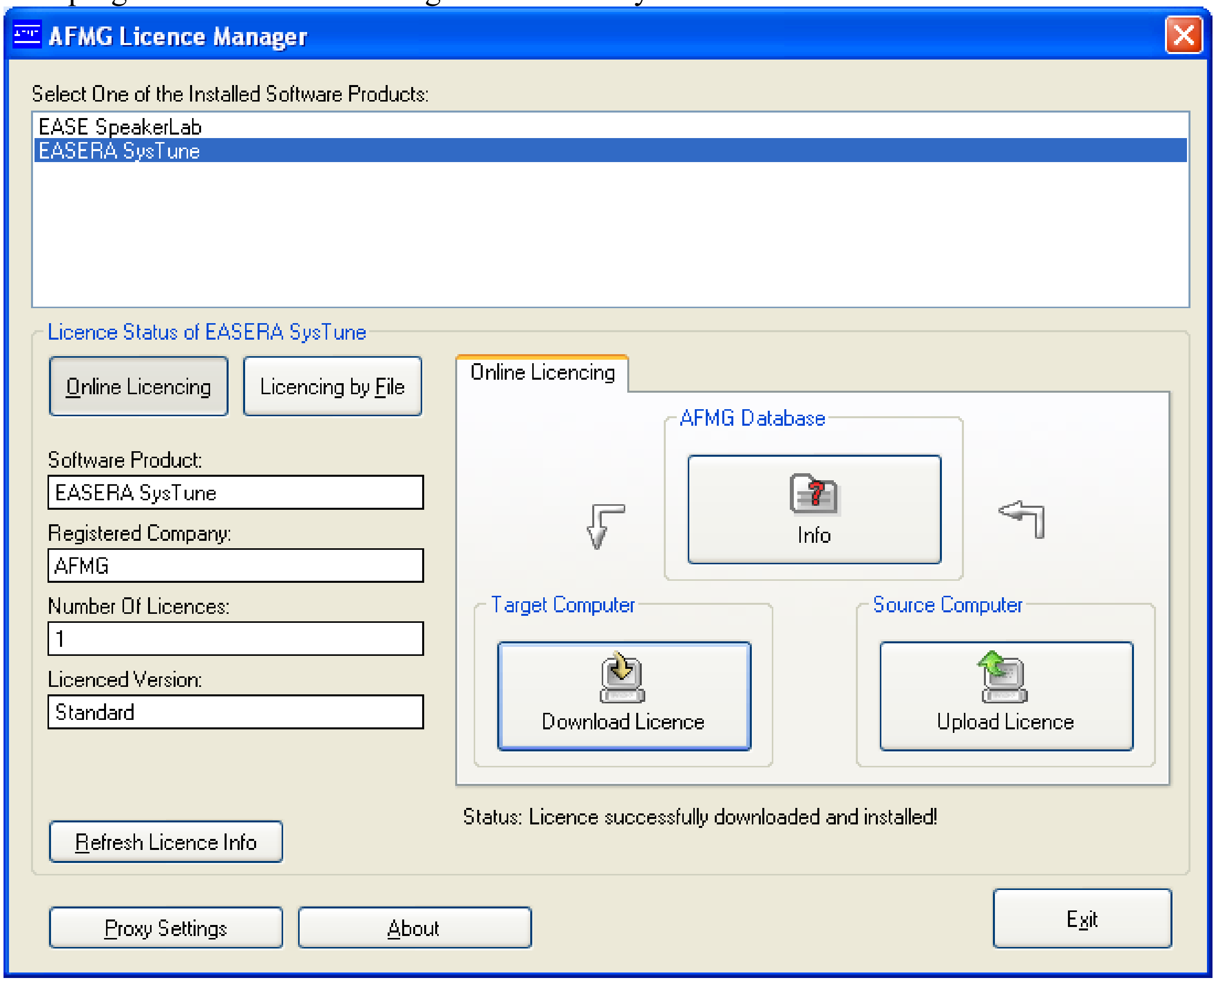

- `Proxy Settings:` Opens the proxy server configuration window. It allows adjusting proxy server settings for online licencing through a proxy server. see also:
  Proxy Server Configuration Window
- `About:` Shows information about the currently installed AFMG Licence Manager.
- `Exit:` Closes the AFMG Licence Manager window.

Licence Status

- `Online Licencing:` Shows the Online Licencing Tab in the right part of the AFMG Licence Manager.
- `Licencing by File:` If it is not possible for you to be online with your computer an EASERA SysTune licence can be ordered via email instead. To do that, generate a reference file and send this file to SDA (support@afmg.eu). As a response you will receive a licence file from SDA which is needed to unlock EASERA SysTune.

  This button enables three tabs in the right part of the AFMG Licence Manager labeled `Licence`, `Terminate`, and `Import/Export`. See below for instructions on how to use these tabs.

- `Software Product:` This field shows the AFMG software product selected for licencing.
- `Registered Company:` This field shows the company name for which the installed licence is registered.
- `Number of Licences:` This field shows how many licences are available on this computer.
- `Licenced Version:` This field shows which software version is unlocked.
- `Refresh Licence Info:` Reloads the licence information.

Online Licencing

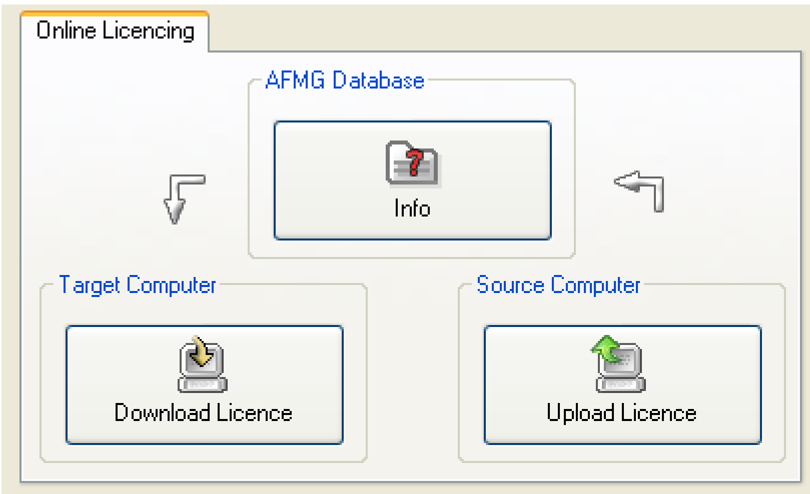

- `Info:` Checks the EASERA SysTune licence database on the SDA web server and downloads information about the user registration, the purchased version and the number of licences still available.
- `Download Licence:` Downloads one licence from the SDA web server. The number in the `Number of Licences` field will be increased by one.
- `Upload Licence:` Terminates all available licences. The licences will be uploaded to the SDA web server and will be available for new downloads later. If you intend hard disk manipulations or to buy a new computer you should use this procedure to prevent a licence being lost. It just means a licence backup for a certain time. After this procedure EASERA SysTune will reset to an unlicenced mode.

Licencing by File – Licence Tab

This tab allows a licence to be installed on this computer. You should only use this option if this computer is not online and a licence download is not possible.

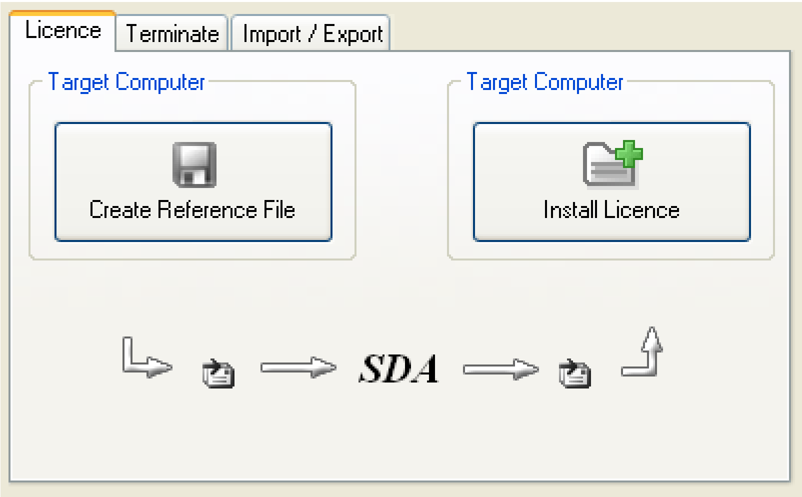

- `Create Reference File`: Creates a *`Reference File`* (*.erf format) to send to SDA. To use the command:
    1. Click on `Create Reference File`.
    2. This opens a `Save Reference File` window after a confirmation message.
    3. Use the `Save In` portion of the window to select the folder where you would like to save the *`Reference File`*.
    4. Click on the `Save` button.

  After saving the file, a "Send Email now?" prompt appears. Click on `Yes` to automatically send a licence order with the reference file as an attachment using an installed email client (e.g. MS Outlook or MS Outlook Express) to SDA. It is also possible to save the file and to mail it later attached to a licence order to SDA.
- `Install Licence`: Click to load a licence file (*.elf format) and install a licence for EASERA SysTune. To use this command: Click on `Install Licence`.
    1. This opens an `Open Licence File` window.
    2. Use the `Look In` portion of the window to find the folder containing the licence file.
    3. Click on the *`Licence File`* name.
    4. Click on the `Open` button.

  If the reference signature in the *`Licence File`* matches this computer, the licence will be installed. The licence information and parameters will be shown in the Licence Status frame.

Licencing by File – Terminate Tab

This tab allows a licence to be uninstalled or removed from this computer. You should only use this option if this computer is not online and a licence upload is not possible.

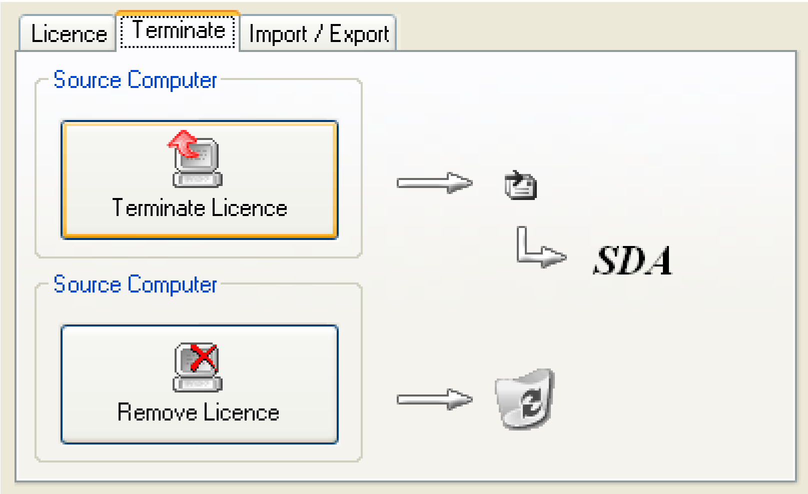

- `Terminate Licence:` Creates a *`Termination File`* (*.etf format) to send to SDA. To use the command:
    1. Click on `Terminate Licence.`
    2. This opens a `Save Termination File` window after a confirmation message.
    3. Use the `Save In` portion of the window to select the folder where you would like to save the *`Termination File`*.
    4. Click on the `Save` button.

  After saving the file, a "Send Email now?" prompt appears. Click on `Yes` to automatically send the termination file as an attachment using installed email client (e.g. MS Outlook or MS Outlook Express) to SDA. It is also possible to save the file and to mail it later to SDA. There it will be verified and if it is correct you can order a replacement licence any time for the terminated one. After this procedure EASERA SysTune will be reset to an unlicenced mode.
- `Remove Licence:` Click to remove all traces of the licence on this computer.
  Note: Be very careful with this button! All licence information will be deleted from the computer. This option should only be used in case of general licencing problems due to software or hardware errors. Please contact SDA before using this command or your EASERA SysTune licence may be lost completely.

Licencing by File – Import / Export Tab

This tab allows the licence transfer between two computers – source and target. The source computer is the computer from where a licence will be exported, which will be imported into the target computer later.

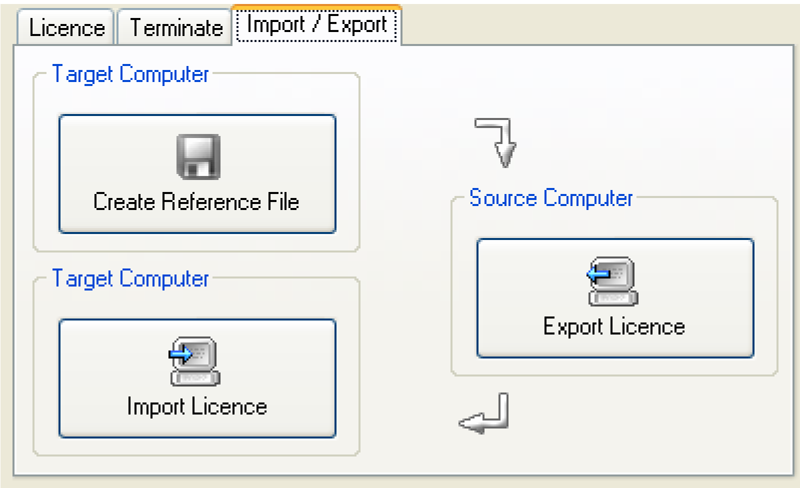

`Target Computer`

- `Create Reference File`: Creates the target computers *`Import Reference File`* (*.eif format). To use the command:

    1. Click on `Create Reference File`.
    2. This opens a `Save Import Reference File` window after a confirmation message.
    3. Use the `Save In` portion of the window to select the folder where you would like to save the *`Reference File`*.
    4. Click on the `Save` button.

  After saving the file you need to copy the *`Reference File`* from the target computer to the source computer (see `Export Licence` below).

- `Import Licence`: Loads the *`Export File`* (*.exf format) from the source computer. To use this command:

    1. Click on `Import Licence`.
    2. This opens the `Open Export File` window.
    3. Use the `Look In` portion of the window to find the folder containing the *`Export File`*.
    4. Click on the *`Export File`* name.
    5. Click on the `Open` button.

  If the reference signature from the *`Export File`* matches this computer, then the licence will be installed. The licence information and parameters will be shown in the Licence Status frame.

`Source Computer`

- `Export Licence`: Loads the *`Import Reference File`* (*.eif format) from the target computer. To use this command:

    1. Click on `Export Licence`.

2. This will open the `Open Import Reference File` window after a confirmation message.
3. Use the `Look In` portion of the window to find the folder containing the *`Import Reference File`*.
4. Click on the *`Import Reference File`* name.
5. Click on the `Open` button.

If the reference signature matches the licence profile on this computer an EASERA SysTune *`Export File`* (*.exf format) will be generated and a `Save Export File` window will be opened:

6. Use the `Save In` portion of the window to select the folder where you would like to save the *`Export File`*.
7. Click on the `Save` button.

After saving you need to copy the *`Export File`* to the target computer (see `Import Licence` above)

Proxy Server Configuration Window

If you are running EASERA SysTune in a secured intranet, access to the internet for online licencing may be blocked by a firewall, depending on your network's configuration. Configuring AFMG Licence Manager to use the local proxy server on your network may resolve this issue. If you are unsure of the appropriate proxy server settings, please consult your network administrator.

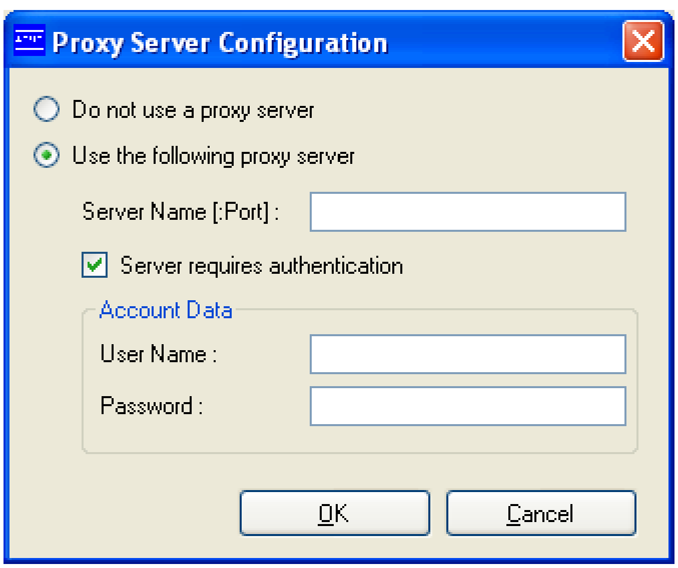

`Use the following proxy server`

- `Server Name [:Port] :` Proxy server name or IP address and port number for internet access.
- `Server requires Authentication:` Check if the proxy server needs an additional authentication.
- `Account Data:` User name and password to authenticate on the proxy server.
- `OK:` Accepts configuration settings and closes the window.
- `Cancel:` Discards configuration settings and closes the window.

Program Tutorial

1. Introduction

Preface

This part of the EASERA SysTune Tutorial is a guide that explains step-by-step all of the important functions of the software and their background in acoustic measurements. It is recommended that you work through this guide at least once if you are a beginner with EASERA SysTune or with measuring software in general. Advanced users can also gain new insights from the following exercises, because the software represents a new approach to making live-sound measurements in several respects.

The next sections will assume a fresh installation and the software in its default configuration. If you have already worked with the software before, make sure you reset all of the configuration data first by selecting EASERA SysTune (use default settings) from the Windows Start Menu under AFMG ▶ EASERA SysTune ▶. Otherwise some displays and calculation results may look different from the ones presented here.

For simpler printing, our explanations use screen shots based on the standard system colors scheme, rather than the default color scheme with bright colors on a black background. To use the same colors as we do, go to the menu labeled Configure, the sub menu Color Scheme and then select the menu item System Colors.

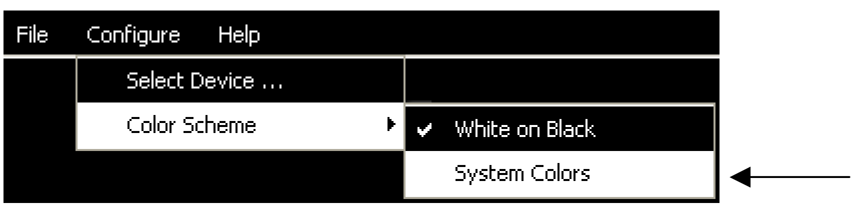

At this time let us agree about a convention for references to the graphic user interface of the software: We will describe items by the sequence of labels in the hierarchy of the controls in the user interface. In the above case, the full reference would be Configure | Color Scheme | System Colors. In simpler cases we may just refer, for example, to the OK and Cancel button of a window in the same manner.

Starting the Software

To start the software click (or double-click) on the EASERA SysTune icon on your desktop.

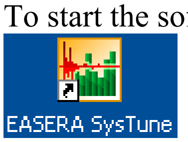

Upon start-up SysTune will show you a screen like the following:

Note that the right figure already shows the program window in system colors. We switched to this alternative, printer-friendly color scheme using the menu command CONFIGURE | COLOR SCHEME | SYSTEM COLORS.

Screen Layout

Let us first have a look at the general screen layout of SysTune. The program window consists of several areas, each with its own purpose:

- At the very top (1) you will find the window caption including the program name and the program version number which is often helpful when you need software support. The menu is located in the same area and gives access to all general functions and parameters, like file saving and loading or program options.
- The control area (2) is located on the left; here is where you select input and output channels, excitation signals and other measurement parameters.
- The right part of the screen (3) is split vertically into two functionally equivalent areas. Each of them shows a graph, a selection menu above the graph as well as a panel for display and calculation options (3a) to the right of the graph.
- In between the two graphs, right in the middle of the window, there is a bar (4) that displays the current coordinates of the mouse and gives access to the mouse modes as well. We will call this area the mouse bar.
- The status bar (5) is located at the bottom of the window. It shows details about the current measurement setup.

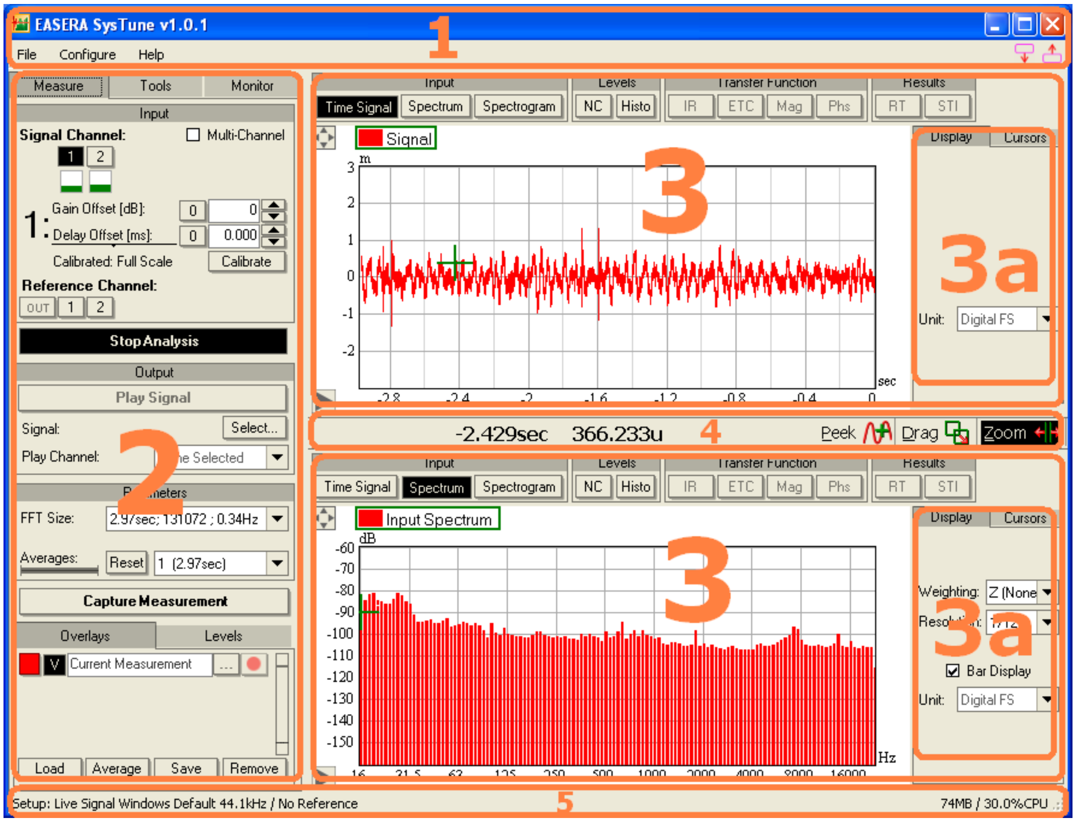
EASERA SysTune v1.0.1
File
Configure
Help
1
Measure
Tools
Monitor
Input
Signal Channel:
Multi-Channel
Gain Offset [dB]:
Delay Offset [ms]:
Calibrated: Full Scale
Calibrate
Reference Channel:
Stop Analysis
Output
Play Signal
Signal:
Select...
Play Channel:
2
FFT Size:
Averages:
Reset
Capture Measurement
Overlays
Levels
Current Measurement
Load
Average
Save
Remove
Time Signal
Spectrum
Spectrogram
NC
Histo
IR
ETC
Mag
Phs
RT
STI
Transfer Function
Results
Signal
3
3a
Display
Cursors
Unit:
Digital FS
-2.429sec 366.233u
4
Peek
Drag
Zoom
Input Spectrum
Weighting:
Bar Display
Setup: Live Signal Windows Default 44.1kHz / No Reference
5
74MB / 30.0%CPU

2. Measurements with a Single Input Channel

One of the most basic measurements one can do with a PC and a soundcard is monitoring an input channel. This is what we want to do first; we will look at the signal at the input of the measuring system in both time and frequency domain views. To pick up the acoustic signal, convert it into the digital domain and make it available to the software, a setup is needed as shown here:

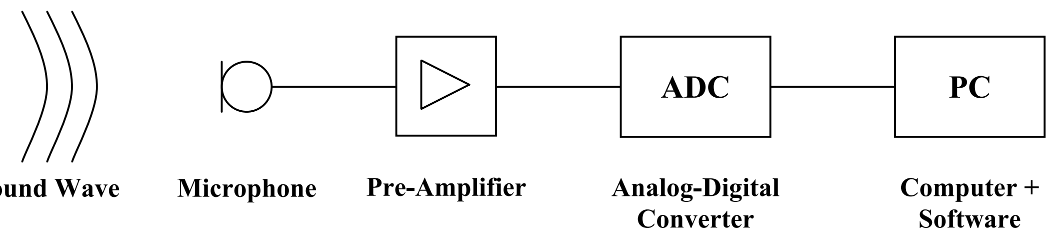

A measurement microphone is used to record the acoustic signal of the sound pressure field and convert it into an electric signal. For most microphones a pre-amplifier is needed to achieve sufficient signal-to-noise ratio. After that the analog voltage signal is transformed by the analog/digital converter into a digital stream of bits that can be received by the driver of the soundcard and then displayed in the software domain.

Selecting the Input Channel

To do such an analysis in SysTune you only need to have a soundcard connected to or built inside your computer. When the software is started it will automatically choose the audio device that is already selected as the default device for audio playback and recording in Windows (please see chapter 5.2 for details about how to change the current audio driver in SysTune). Also by default, SysTune will select the first input channel of the soundcard.

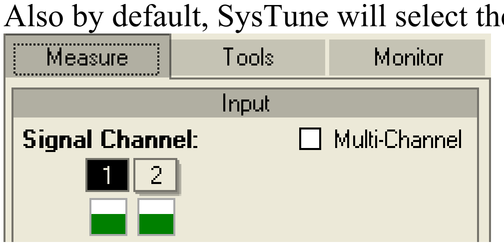

The current input channel is indicated by a highlighted button in the row of buttons located in the control panel on the left below the label SIGNAL CHANNEL. These buttons are labeled with numbers according to the associated input channels. Depending on the connected soundcard you may see up to 8 input channels here. The minimum number of input channels is always 2, so there will always be at least two buttons.

Note: If the hardware supports only a single input channel, Windows will automatically image this channel and create a quasi-stereo configuration.

You can change the current input channel by left-clicking on the button with the corresponding number.

Below the row of input buttons there is a row of small level displays, each related to the input directly above. These so-called mini-meters show the current signal level at the input. They are particularly useful to monitor the status of all connected inputs simultaneously. The mini-meter shows a vertical green bar of varying height when levels are in the normal range.

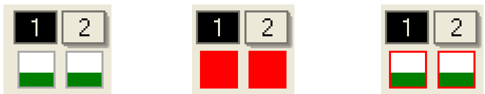

However, when the signal at the port of the A/D converter is greater than the maximum that is possible for the electronics, the input signal is clipped upon conversion. The mini-meters indicate proximity to clip level by yellow color, for –6 to –1 dB below clip level, and by red color, for signals of –1 dB below clipping and higher. To ensure the data validity of your measuring system, make sure that you do not exceed clip level at any time. It is good practice to adjust the gain control in such a way that the peaks of the signal are maximally in the yellow range.

Hint: Because clipping can happen for only very short periods of time, such as during signal peaks, you may not always be able to catch the red bar lighting up with your eyes. For this reason, the frames of the mini-meters remember the last clip state in the order of green-yellow-red until reset with a mouse click directly on the mini-meter.

2.1. Time Signal

Now let us look at the input signal as it arrives in the software domain. In the default configuration, the software starts with the TIME SIGNAL button selected for the top graph. Along with the graphs SPECTRUM and SPECTROGRAM in the same group, this display can be used immediately for any kind of INPUT signal without adjustment of additional parameters.

If there is no signal except for noise at the input, the graph will look similar to the following picture and it will be continuously moving from the right to the left.

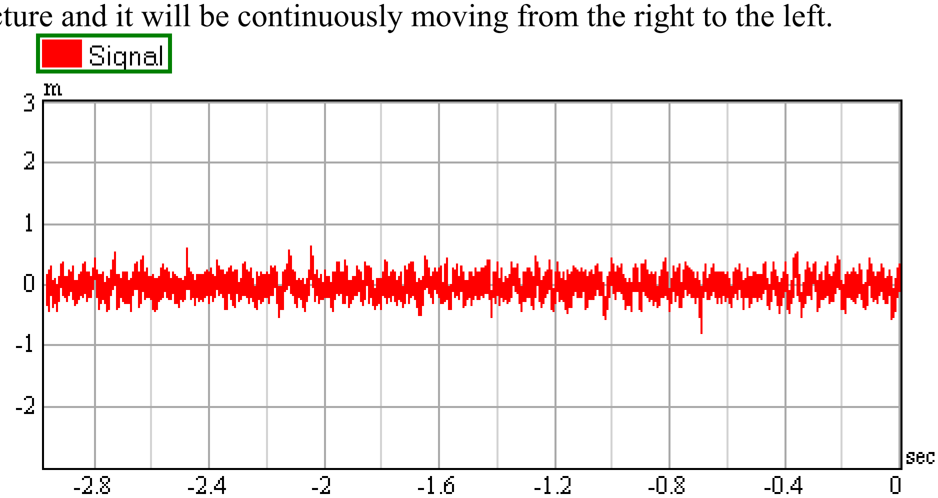

If the graph does not change with time, make sure you have started the real-time analysis, it is on by default for the very first program start. To do that look at the control panel to the left, right below the INPUT section. If the first large button is labeled START ANALYSIS, the real-time functions of SysTune are currently suspended. Left-click on the button to restart the analysis. If the button is already labeled STOP ANALYSIS and it is highlighted, the TIME SIGNAL graph should be updating continuously. If this is not the case, please refer to the trouble shooting section at the end of the tutorial.

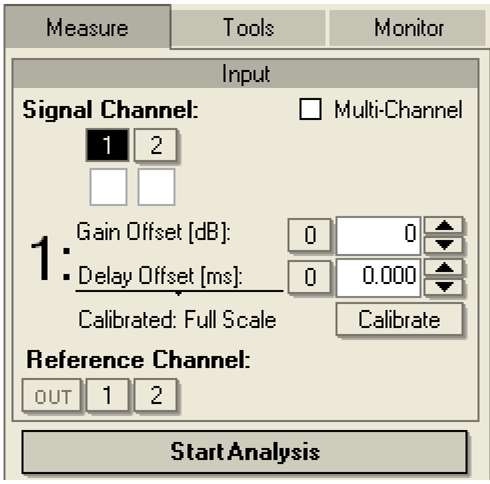

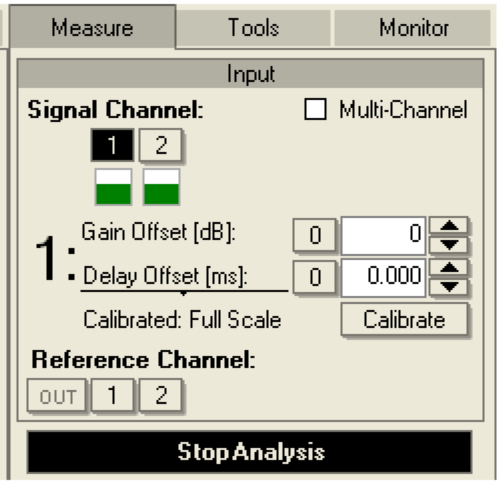

Back to the TIME SIGNAL graph. The horizontal axis shows the time passed by, maximally for the period of the current FFT block size; we will come back to that a little bit later. The vertical axis shows the signal amplitude in digital units, also called full-scale (FS). Because generally EASERA SysTune does not know which hardware is being used, it cannot display real-world numbers like Pascals (Pa) or Volts (V) directly. But the software can be calibrated very easily. Calibration here means giving the software a relationship between digital values, which is the only thing the software really knows about, and physical values, as they can be measured in the real-world.

Calibrating an Input Channel

Now we would like to calibrate the first input channel. If you have switched to a different channel in the meantime, click on the button labeled 1 to activate the first channel again. Right below the mini-meters, the software shows an area related to the properties of the input that is currently selected. The fields GAIN and DELAY correspond to parameters that we will investigate in a little while, we will first focus on the calibration. It can be started by a left-click on the CALIBRATE button below the DELAY text field. Note that the current calibration status is always shown to the left of the CALIBRATE button. If the input has not yet been calibrated the label will show FULL SCALE.

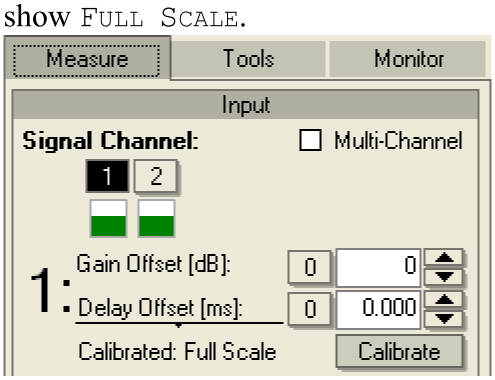

After pressing the CALIBRATE button the CALIBRATION window will open. It shows the ordinal number of the input channel selected for calibration in the window caption.

Since we will be performing an acoustic measurement with a microphone, we need to tell SysTune the relationship between pressure units, Pascals or dBSPL, and digital units, full-scale or dBFS. This is exactly what happens on the first tab USE PRESSURE, the second tab USE VOLTAGE can be used for electrical measurements.

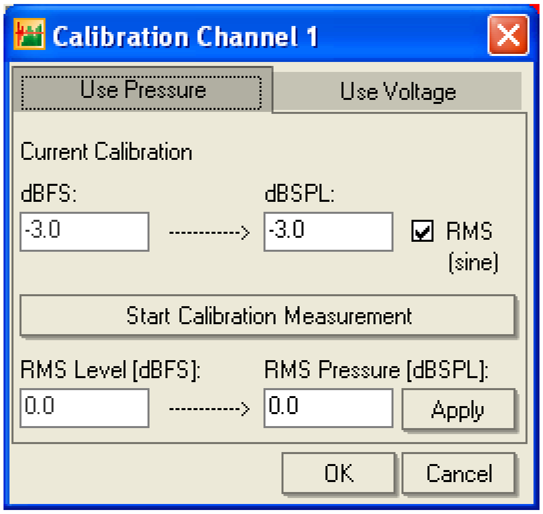

The upper part of the CALIBRATION window displays the state of the CURRENT CALIBRATION. The sound pressure level that is equivalent to a defined full scale level is shown in the dBSPL text field. Later on, if you know the calibration for an input channel you can enter it directly here. The check box labeled RMS (SINE) allows you to toggle between the display of the peak or the RMS numbers for a sinusoidal signal.

For a full acoustic calibration do the following:

- put a microphone calibrator on the microphone,
- click on the START CALIBRATION MEASUREMENT button,
- wait until the value in the RMS LEVEL [dBFS] text field settles and then press the STOP CALIBRATION MEASUREMENT button,

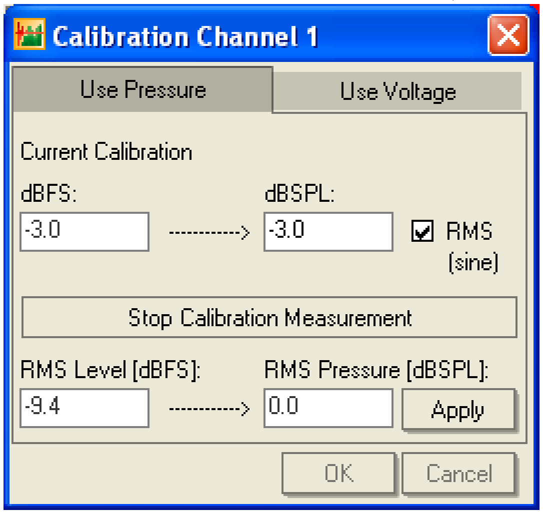

- finally enter the corresponding pressure value in the field labeled RMS PRESSURE [dBSPL], such as 114, and click on APPLY.

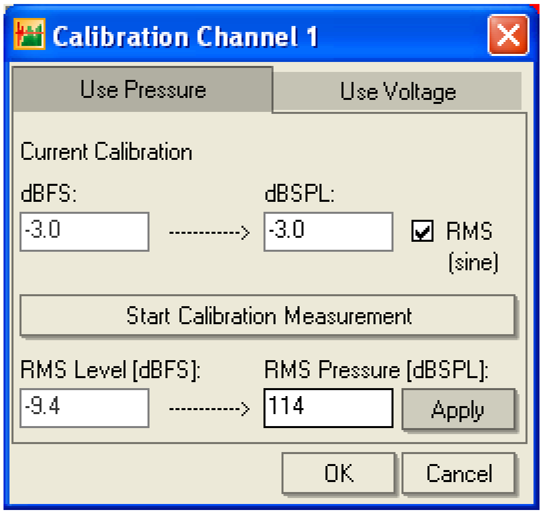

The CURRENT CALIBRATION will be updated immediately as shown below. To confirm the calibration and close the window press OK.

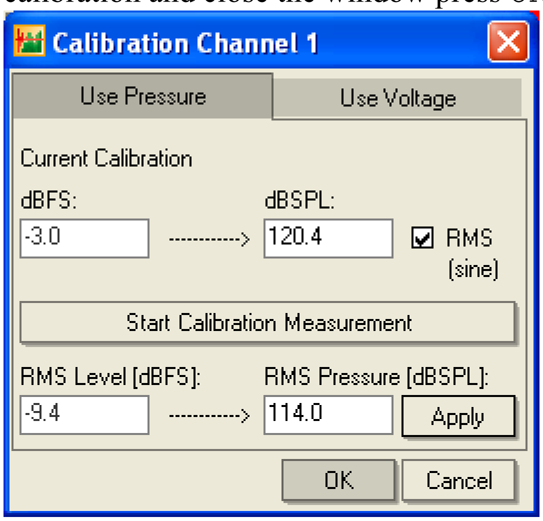

Your measurement setup is now calibrated and the control panel will show the new calibration state PRESSURE.

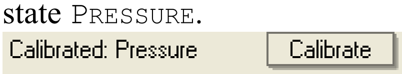

Hint: If you do not have a calibrator at hand while going through this tutorial, try to whistle into the microphone and enter a value of 80 dBSPL. This will likely be right within an error of +/-20 dB and will allow you to follow the subsequent steps for calibrated data.

Having calibrated our measurement system successfully, we can now switch the TIME SIGNAL display to a physical unit. To do that go to the right panel and change the current setting for the unit from DIGITAL FS to PHYSICAL.

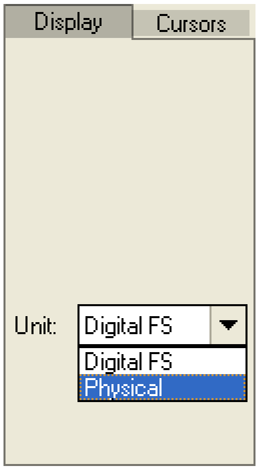

The TIME SIGNAL graph will immediately reflect that change by showing Pa or mPa for the vertical scale.

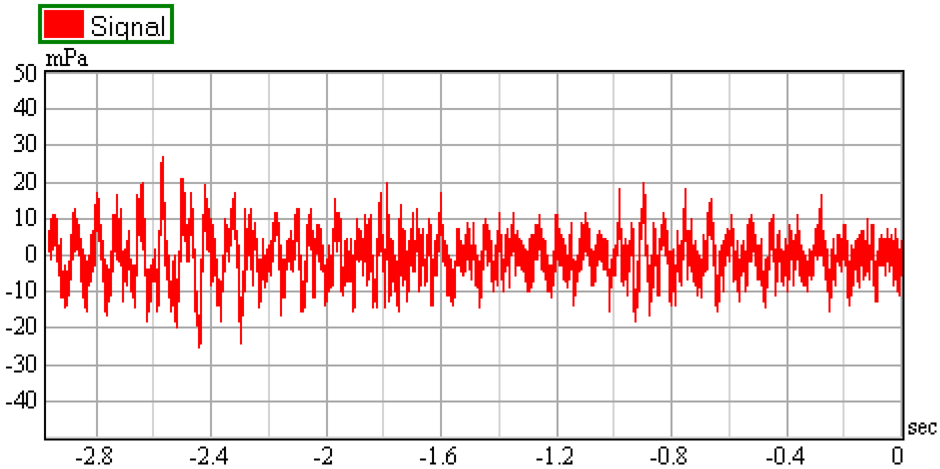

You may clap your hands close to the microphone or generate some other impulse-like noise to see the effect in the graph. You may even let the input clip just for this moment. Have a short look at the mini-meter and how it reacts to the changing input signal.

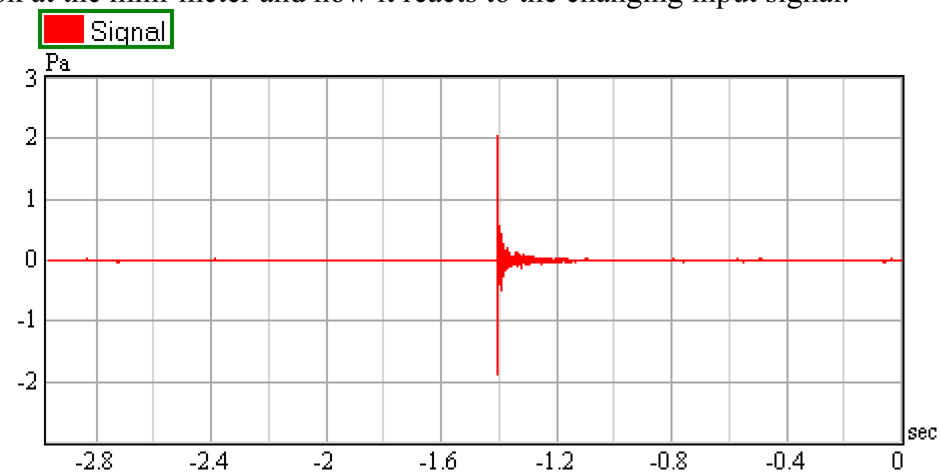

In the TIME SIGNAL graph you will notice that the vertical axis is scaled automatically to a larger section to show all of the data that arrives. However, it does not automatically collapse to the original range after the spike has left the displayed period of time. That is because the program expects more signals of that order of magnitude and therefore remembers the maximum view limits. To reset the view just double-click on the graph.

Adjusting the View Limits
Now it is time to look more closely at the scaling of the diagram. By default, the current mouse mode is the ZOOM mode, as indicated in the mouse bar.

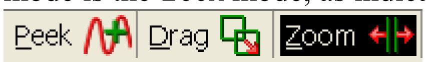

In this mode you can use the left mouse button to zoom into the graph with respect to the horizontal axis. To do that, left-click on the start value for the new view limits and keep the mouse button pressed while dragging the mouse pointer to the stop value. While dragging, the program will indicate the current start and stop values by vertical lines or zoom markers. Finally, release the mouse button to confirm the new view window. Also in ZOOM mouse mode, you can use the right mouse button to select the view limits in a similar way for the vertical axis.

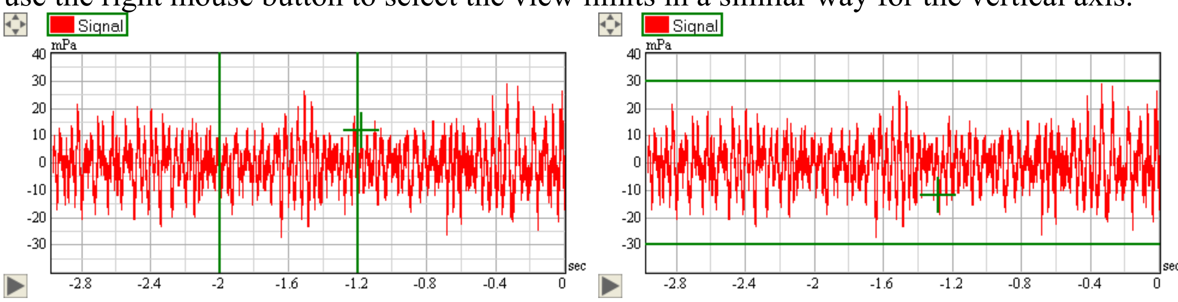

Hint: The zoom markers snap to the lines of the graph. To freely select the zoom area, keep the Alt key pressed while dragging the mouse.

Let us try out the DRAG mouse mode, too. Select this mouse mode by first clicking on the button labeled DRAG in the Mouse bar, then left-click on the graph and keep the left mouse button pressed. When you now move the mouse you can shift the current view port freely. The PEEK mouse mode is the third mouse mode available for all graphs, but we will come back to it at a later point of time, when there is more meaningful data to peek at.

At any time you may return to the full view limits by double-clicking in the graph area or by a left click on the auto-scale button in the upper left corner of the graph.

In addition to changing the view limits with the mouse, you can also enter them directly as numerical values. To do this we need to open the view limits section with a left click on the triangle button in the lower left corner of the graph.

This command slightly rescales the graph vertically to create some space for a new panel; the view limits section, shown in the next picture.

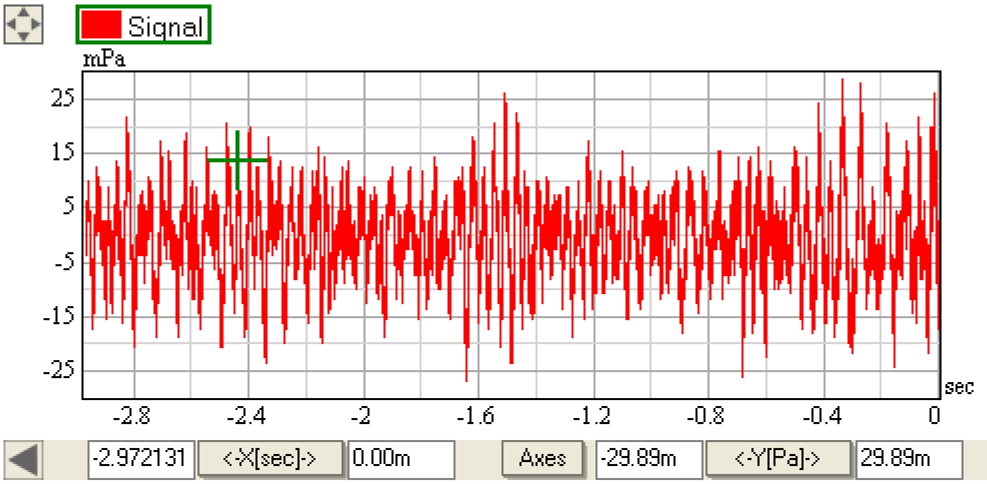

The left two text fields, located on either side of the button <-X[SEC]->, define the start and end point for the horizontal or X-axis. The right two text fields, located on either side of the button <-Y[PA]->, denote the view limits for the vertical or Y-axis. The box to the left of each of the buttons always contains the start and the box to the right contains the end value. For the moment, let us select a time period of 2 seconds for X and a range of +/-50 mPa for Y to exercise.

Hint: You can enter numerical values using exponent prefixes, such as "m". A value of 50 mPa can thus be entered either as 0.05 or as 50 m into the text field.

The result of this change is shown in the picture below. Depending on your input signal and calibration it will look a bit different, of course. Now let us go back to the full view limits. While you can use a double-click to achieve that, you can also use the buttons <-X[SEC]-> and <-Y[PA]-> to individually return the view limits to their default setting for the current data set. Left click on <-X[SEC]-> to scale the horizontal axis to contain all data points, left click on <-Y[PA]-> to do the same for the vertical axis.

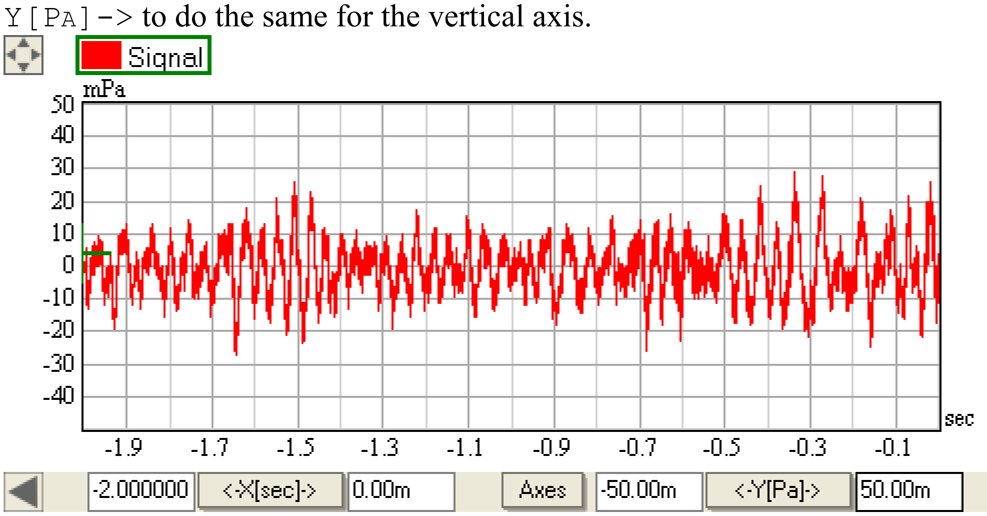

Now that you have learned all of this you should be able to pick up a signal with the microphone, stop the analysis for a while, zoom into the sound event to look at it as a function of time, zoom back to the full view and start the real-time analysis again.

Tech-Note:
The Time Signal graph shows the raw data as it is generated by the A/D converter. As in all data in the software domain, it is discretized. The continuous voltage signal at the input is sampled with the selected sample rate. This means that every sample displayed in the software domain is actually an average over a small period of time. This period is exactly the inverse of the sample rate, for 48 kHz this is about 20 µs or 0.02 ms. Shorter time events cannot be resolved.

Also with regard to the magnitude, data is discretized. Depending on the A/D converter available the signal amplitude is rastered with nominally 16 to 32 bits. For acoustic measurements 16 bits are usually sufficient, this corresponds to about 32,000 values between 0 and 1 full-scale and represents a dynamic range of 96 dB. For electronic measurements a higher resolution is often desirable, although soundcards in the normal price range will supply seldom more than 20 bits effectively which represents a dynamic range of 120 dB. The bit resolution determines how accurately small values and small changes in the input voltage can be represented in the software domain.

Summary

In this section we have made our first simple measurements with SysTune. We have selected an input channel, calibrated it and looked at the incoming signal in the time domain. We are now also able to navigate through the displayed area of the TIME SIGNAL graph with functions that work the same for all graphs in SysTune.

2.2. Input Spectrum

We just looked at the input data in the time domain; now let us take a look at the same data in the frequency domain. In the default configuration, EASERA SysTune starts with SPECTRUM selected for the bottom graph. This view shows the frequency data that corresponds to the time data in the upper graph TIME SIGNAL.

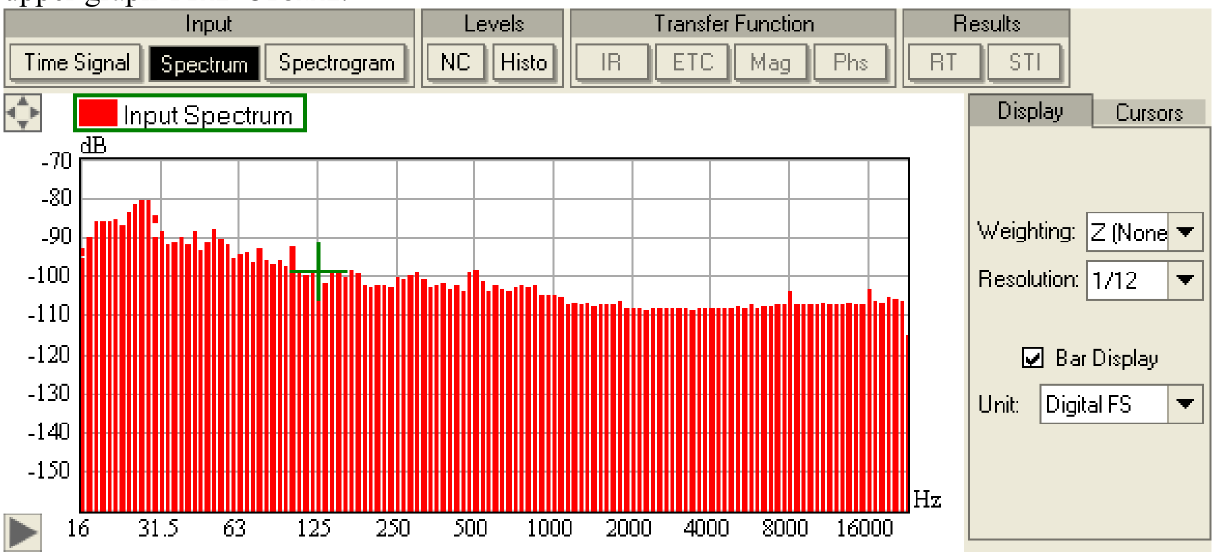

By default, the diagram is shown as a bar display in 1/12th octave resolution and with no weighting applied. Depending on the dynamics of the input signal you will also see a second curve, namely the peak hold curve. It shows the short time history of the spectrum curve.

In the right panel on the DISPLAY tab you can select settings different from that and we will go through them now briefly. The first selection that can be made is the WEIGHTING applied to the frequency data. The three weightings A, B, C superimpose different correction curves to take into account the characteristics of human hearing.

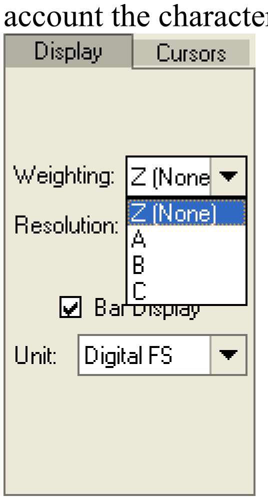

The human ear is less sensitive to signals at low and high frequencies compared to the mid range around 500 Hz to 2 kHz. Therefore, sound level measurements do not correspond directly to the perceived loudness of a signal. A weighted display of the input spectrum accounts for this effect as it shows the levels as they would be perceived according to the A, B or C weighting standards

(ANSI S1.4 (A, B, C) or IEC 61672-1 (A, C, Z)). In fact, the different weighting functions have their background in different types of signals, like pure tones or noise, which are – again – perceived differently by the human hearing system. The following picture shows the A, B and C weighting filters as a function of frequency.

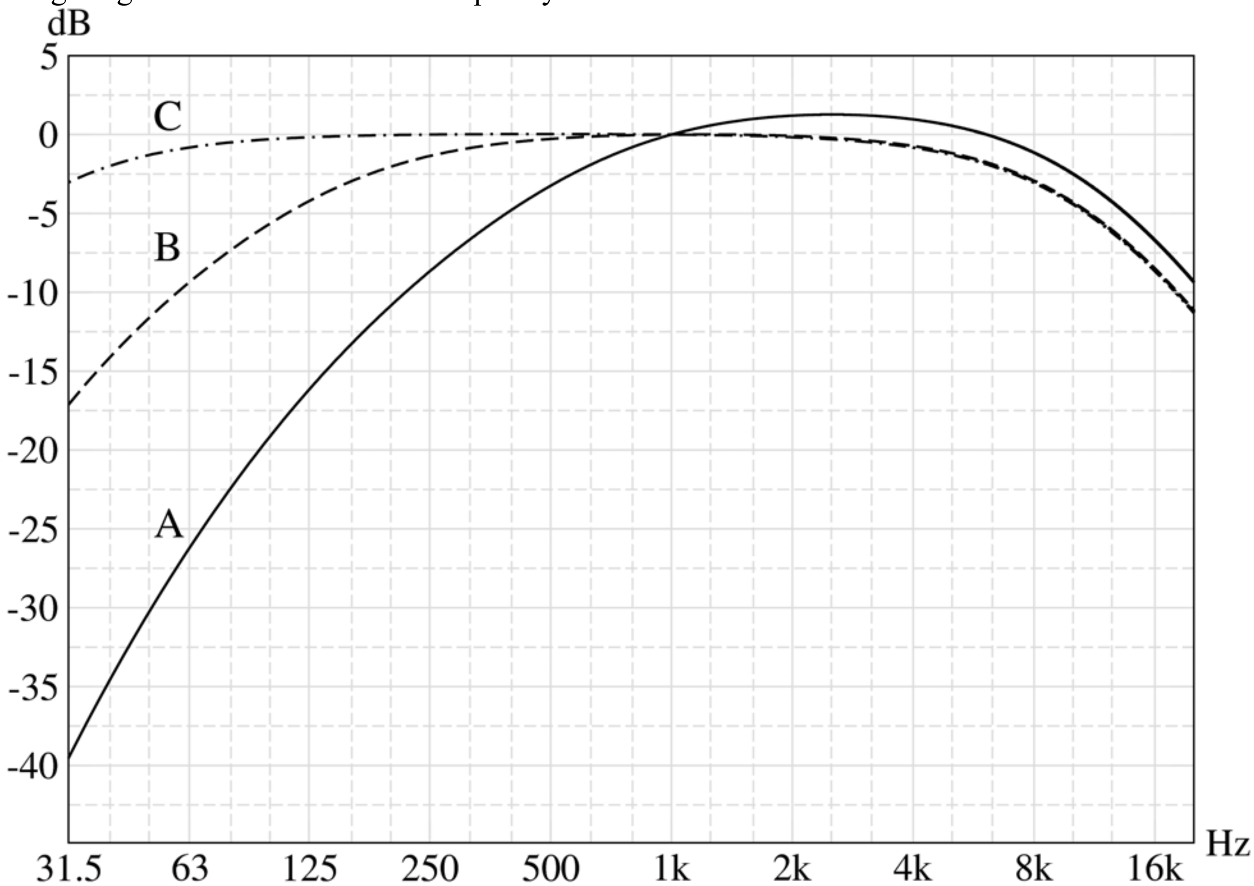

Switch between the different weighting functions to see their effect on the SPECTRUM data. The setting Z (NONE) will always take you back to the unweighted (also called zero, flat or linear) graph.

The second parameter in the DISPLAY panel is the Resolution. It controls how wide the bands are over which the frequency data is combined.

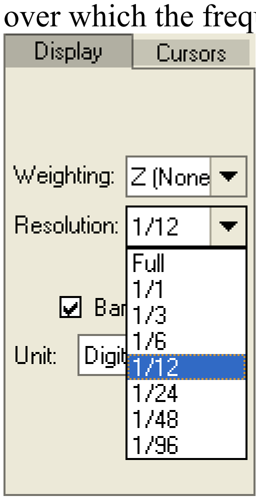

The result of applying an FFT to the time data is frequency data with linear spacing, which means equal spacing between adjacent frequency data points. To display this data in fractional octave bands like `1/1` or `1/12` all of the data points lying in one band are summed to a single value to obtain the level for that frequency band. The FULL resolution is the only resolution where the data is displayed raw, although it is seldom used.

Tech-Note:
Pink noise has the characteristic property that it is a flat curve when shown in summed fractional octave bands. This type of view corresponds to the power contents of the signal. The same holds true for other pink signals, which are signals with a 3 dB drop of power density per octave band, like a log-sweep. In contrast, signals with constant power density over frequency, like White noise, show levels increasing with frequency in such a summation graph.

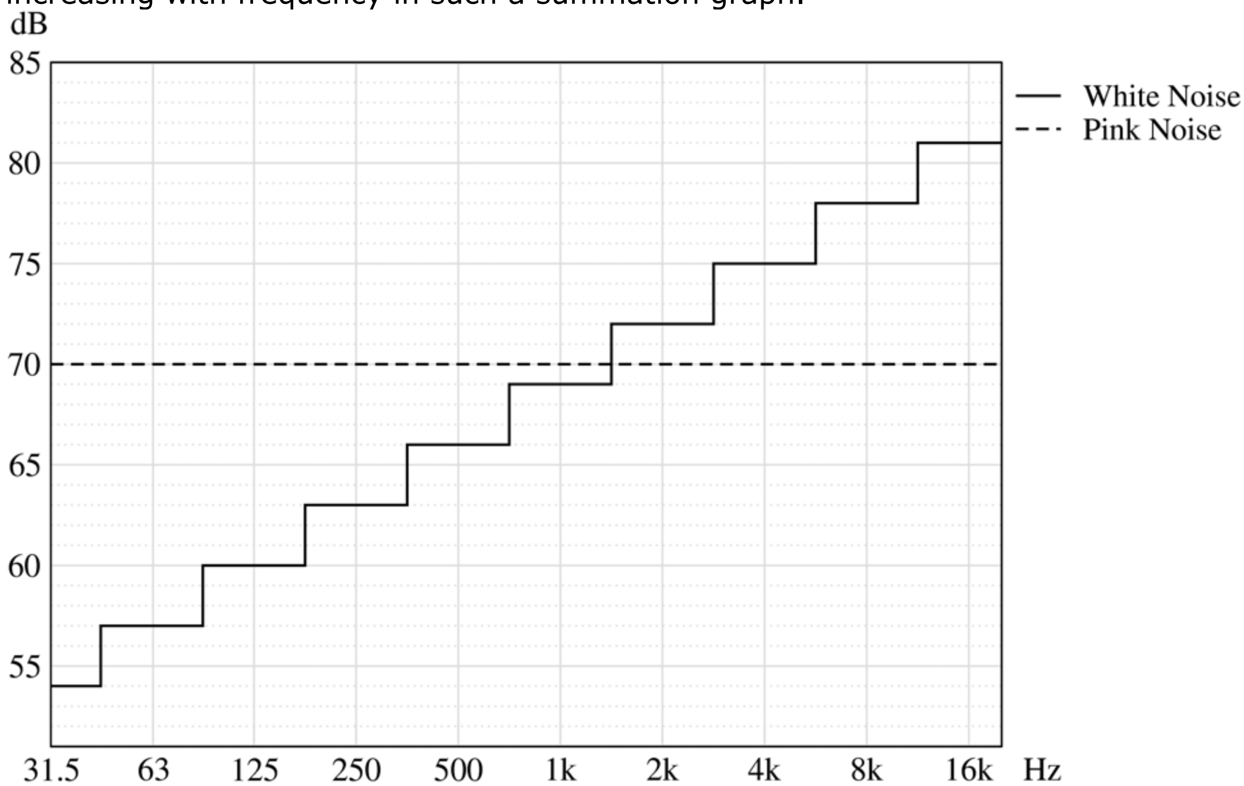

When using the Full resolution graph the behavior will change, because now the program displays power densities instead of powers, that is summed power densities. In this kind of graph White noise is represented by a flat function of frequency and Pink noise as a curve decreasing in level by 3 dB per octave.

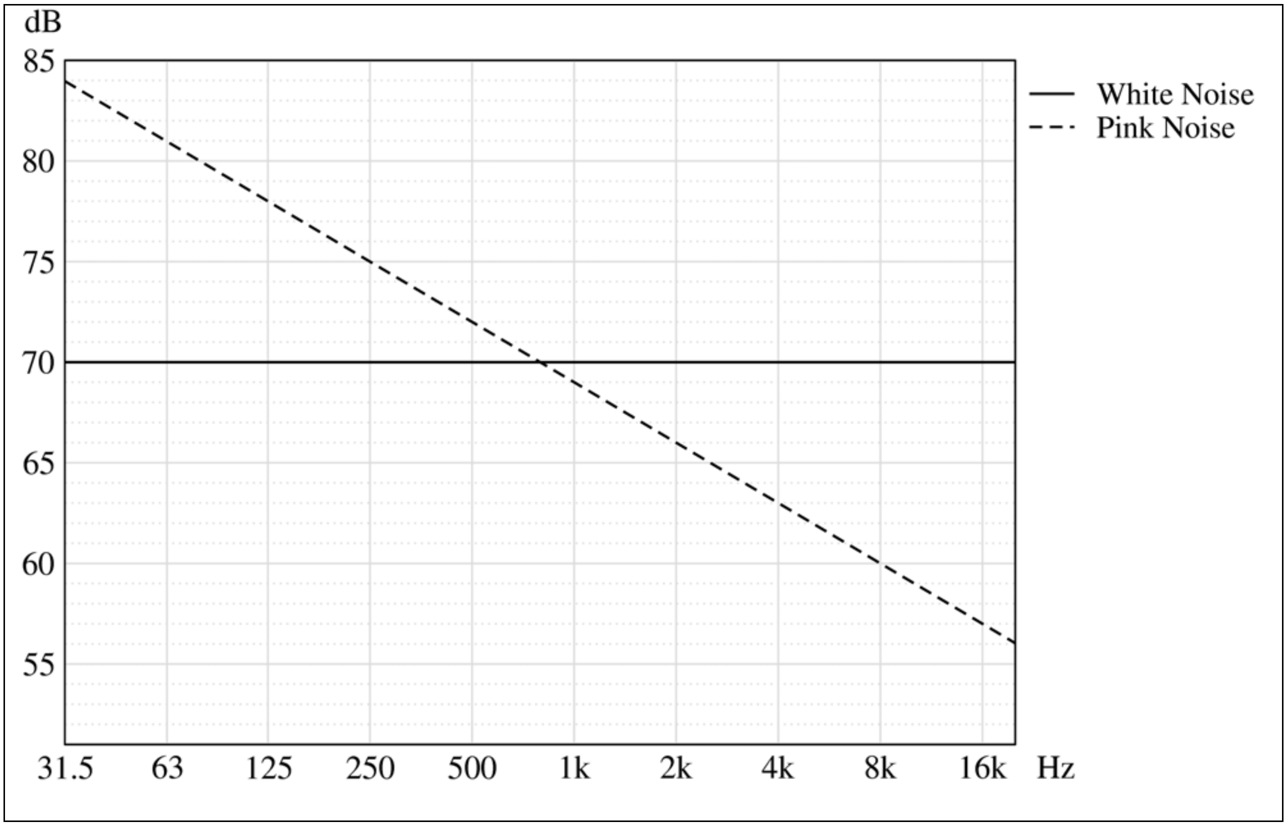

Finally, the SPECTRUM display can be shown in two ways, either as a bar graph or as a curve graph. Use the check box BAR DISPLAY to toggle between them. While it is more common to use a bar graph for fractional octave diagrams, it is often more difficult to use this kind of view for analysis purposes. Especially looking at overlaid curves which we will discuss in detail below is much easier for lines only. Note that there is no BAR DISPLAY for the FULL resolution setting.

For now, let us choose no weighting, 1/3rd octave bands and a curve display. As we have also already calibrated the input channel, we may also select PHYSICAL as the UNIT. After making these selections you should see a screen like the one below.

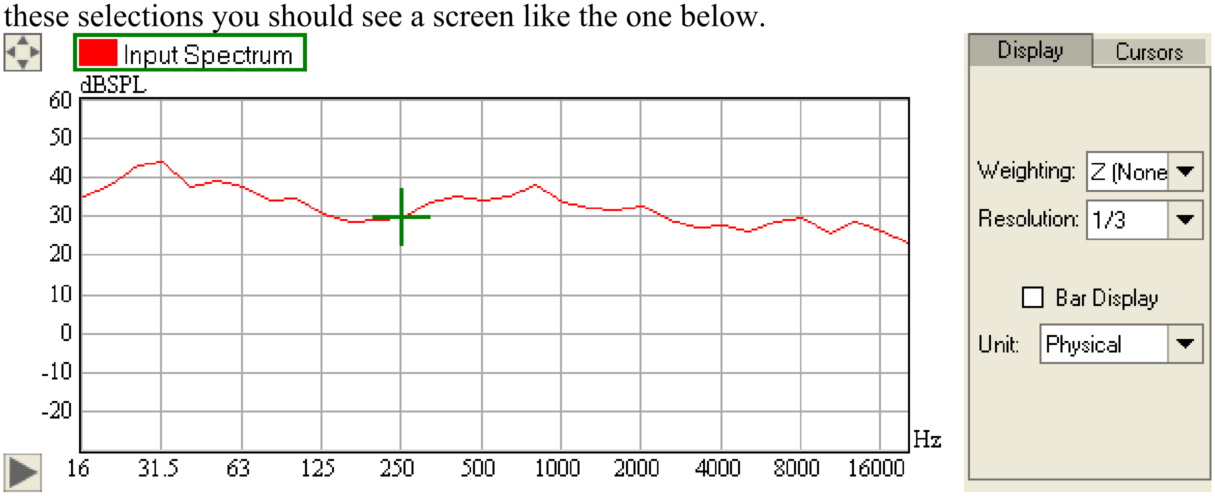

Also this display is still based on the raw data from the input, just with different calculation parameters for the post processing. The original time domain data is continuously retrieved by the program, transformed into the frequency domain and displayed as a spectrum, which is level as a function of frequency.

Choosing the FFT Size

The FFT size is an important parameter with regard to the frequency resolution. It can be selected using the drop down list labeled `FFT SIZE` in the control panel on the left. Each item of the list shows the FFT block size or time length in seconds, the corresponding number of samples according to the current sample rate as well as the frequency resolution, for example `2.97SEC; 131072; 0.34HZ`, if using the sample rate of 44.1 kHz.

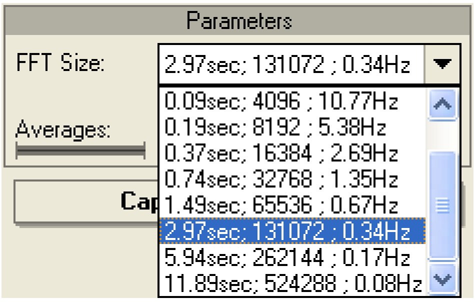

Note that for shorter time lengths, e.g. higher time resolution, the frequency resolution decreases. This means that the spectrum display can only resolve short time events by compromising the resolution in the frequency domain. Vice versa, a high spectral resolution, for example desirable to identify resonances, will require a long FFT time length and thus it will have a very long time dependency.

Tech-Note:
As in the real world, in software the time and frequency domain are also strongly interrelated. The spectrum as displayed in SysTune is derived by means of a Fast Fourier Transform (FFT). This transform creates a set frequency samples from a given amount of time samples. The more time samples are used for the transform, the higher is the density of data points in the frequency spectrum and thus the resolution. Sample length Δt and frequency resolution Δf for the FFT are related by the equation $\Delta f = 1 / \Delta t$. (See for example: Oppenheim, Schafer: Discrete-Time Signal Processing, 1999, Prentice-Hall, Inc., New Jersey)

In SysTune a set of useful time lengths for the FFT is predefined. There is no sense in very short block sizes especially, like only 4 or 8 samples, because then the frequency resolution becomes far too low. Very long block sizes like several minutes are also not available, because the measuring times become impractical.

Since each frequency data point is derived from all time samples of the given block, the resulting data must be understood as an average over the full time length used. For signals varying quickly over the time period of an FFT the resulting spectrum will be the time-average of that signal over that period.

Another important point with respect to the FFT is that generally some windowing must be applied to the FFT block. Because a cyclic FFT is used for the transform from the time domain, any abrupt changes between the start and the end of the block will cause disturbing artifacts in the frequency domain. A flat-top window applied to the FFT block helps to smoothen this transition and is particularly needed for signals that are smooth and periodic, like a sine wave, but with a period different from the FFT block length. (See for example: Fredric J. Harris: On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform, Proceedings of the IEEE, Vol. 66, No. 1, January 1978)

Let us have a look at the effect of changing the FFT size. Select a time length of 3 seconds or so and create a sharp, loud impulse at the microphone. You will see that the spectrum increases immediately, stays elevated for the time of the FFT block length and then drops to its original state. Now switch to a short FFT block length, such as about 0.2 seconds. Again, create an impulse and watch the spectrum rise and decay. In practice, there will seldom be a need for FFT sizes beyond these lengths.
Also, have a look at the frequency resolution at this time. For an FFT size of 0.2 s the spacing between frequency points is about 5 Hz. You will recognize that at the low end of the spectrum, the graph looks stepped and very rough. This roughness in frequency is due to the fine resolution in time. However, in practice you will seldom need frequency resolutions much higher than 5 Hz.

Averaging over Time
The second important parameter for the calculation of the Input Spectrum is the number of Averages. You will find the drop down list right below the FFT Size selection in the Parameters section. This list shows the number of FFT blocks to be averaged and the corresponding overall time length.

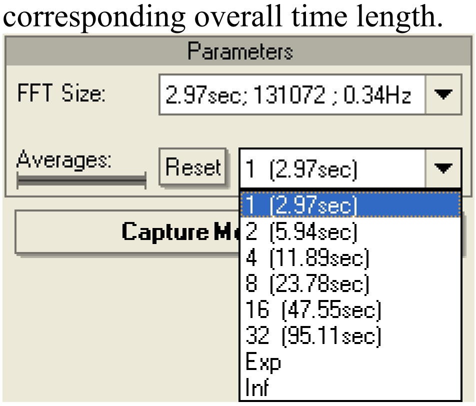

The number of Averages defines how many FFT blocks are transformed into the frequency domain and then averaged to yield the Spectrum. The longer you average the data, the less significant will singular time events affect the overall Spectrum. Also, the signal-to-noise ratio is increased by 3 dB for every doubling of the number of averages. On the other hand, just like the FFT block size, a long averaging time reduces the temporal resolution. When you average over 20 seconds of time, you will not be able to identify a peak of a few milliseconds length.

Tech-Note:
The overall time is what counts for the acquired spectrum data. The main reason to split the time length into an FFT block size and into a number of averages is to keep the performance requirements practical, because they can become too high for very large FFT block sizes.

So, for the same time length, half the FFT size and twice the number of averages will yield about the same result. Note that this will change the frequency resolution and some of the data as well because a different number of FFT windows – one per average - are applied.

Below the label AVERAGES there is a small horizontal meter that shows the time that has elapsed since the measurement was last started and it is shown relative to the overall averaging time. It indicates how much of the data in the current display actually is valid data. At any time you may hit the RESET button next to the meter to restart the measuring process.

Hint: Assume you are performing a spectrum analysis in a venue and you are using a long averaging time. Now, unexpectedly, the measurement is disturbed by someone shutting a door, then just push the Reset button to restart the data gathering process.

At the bottom of the list of AVERAGES you will find additional selections, EXP and INF:

- The setting INF lets the averaging process simply continue forever instead of covering only a limited period of time. This option may be helpful when the maximum number of Averages does not provide enough signal-to-noise ratio.
- The EXP item also runs infinitely but it applies exponential weighting to the averaging process. This makes time blocks further in the past less important than recent ones. The slope of the weighting function can be selected in the OPTIONS window. This function is useful if you would like to monitor average levels with a smooth roll-off of high peaks over time.

Remember in this respect that the regular averaging settings provide a hard cut-off, which makes peaks disappear abruptly when they leave the time period selected for averaging.

Checking the Calibration

So far we have only been looking at a random noise signal at the input. It is certainly just as interesting to see how a sinusoidal signal appears in SysTune. The simplest way to do that is to take the same microphone calibrator we used a little bit earlier and put it on the microphone. Also, switch to an FFT SIZE of about 1 second length and select 1 for the AVERAGES. The frequency RESOLUTION should still be 1/3 and no WEIGHTING should be used, return to a BAR DISPLAY for a moment.

After switching on the signal, you should see a fairly large horizontal bar in the TIME SIGNAL graph and a distinguished peak above some noise floor in the SPECTRUM. You may have to double-click into each drawing to get back to the full view. If you have calibrated the input as described before, you can also switch the current UNIT to PHYSICAL for both top and bottom graphs. They will then look approximately like the following.

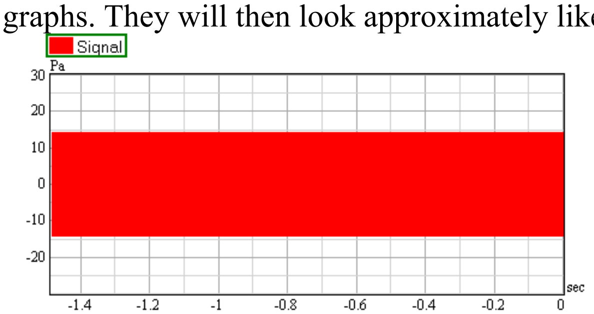

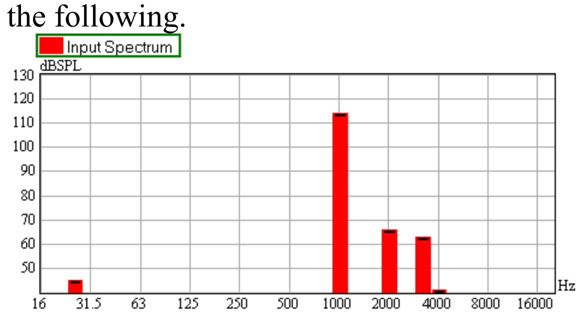

If the input was calibrated correctly, then the Spectrum should now show a peak at the calibration frequency (here 1000 Hz). For that frequency band, the sound pressure level of the bar should be the same as the level you calibrated to (here 114 dBSPL), maybe it will be off by a tenth of a dB. You can verify this by zooming into the area of interest, but there is an easier way as well. You may have already noticed that whenever you move the mouse over the graph a green cross is following the mouse, tracking precisely the current curve. The values that correspond to the location of the tracking cross on the horizontal and vertical axis can be viewed on the mouse bar; it is centered between the top and the bottom graph. Note that the read-outs are always related to the graph where the mouse is hovering. Carefully move your mouse close to the location of the peak. Let the tracking cross settle on the top of the peak and look at the mouse bar. It should now be showing the frequency and level of the peak.

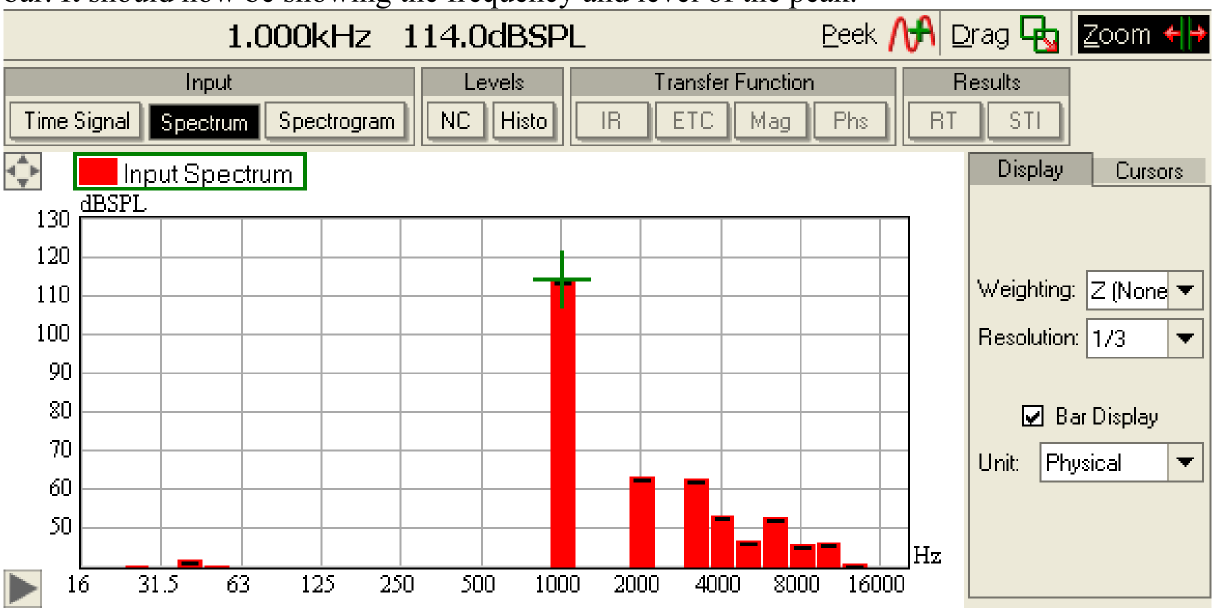

For further analysis, stop the measurement for a while (STOP ANALYSIS on the left side) as well as the calibrator and zoom into the upper graph, the TIME SIGNAL. If you zoom in (drag using the left mouse button), you will immediately recognize that the horizontal bar is actually a compressed sine wave signal. Again, use the mouse tracking cursor, find the maximum and read it off the mouse bar.

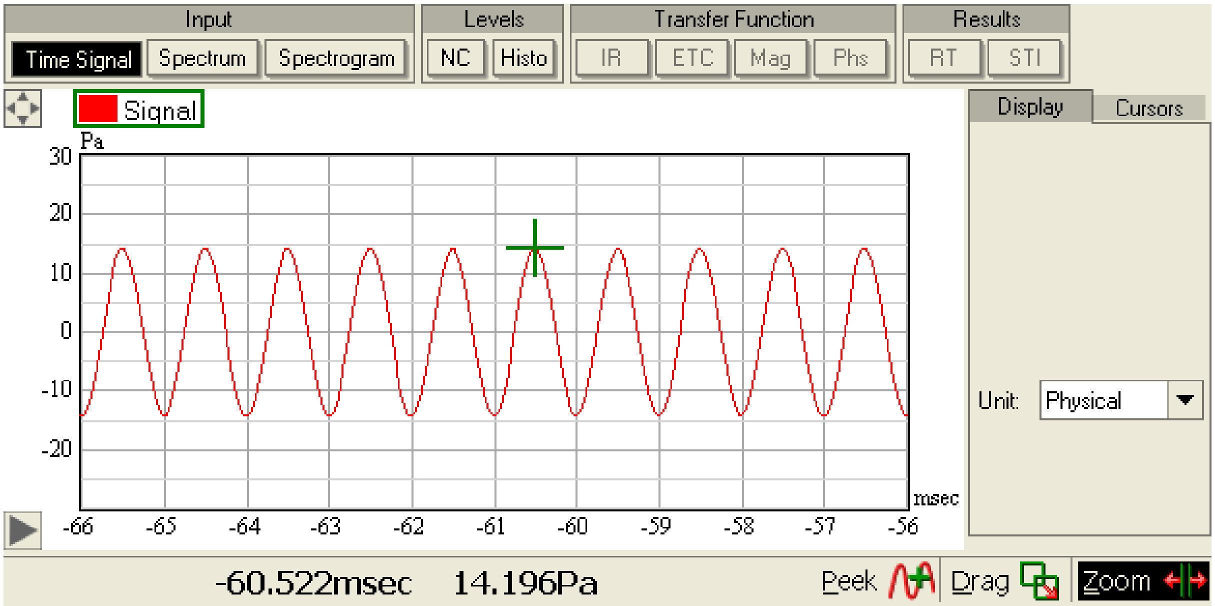

In our case we find 14.2 Pa which equals 117 dBSPL. At first glance, this seems wrong as we measured a 114 dBSPL before in the frequency domain. However, here we need to distinguish between peak values, like the maximum shown in the TIME SIGNAL, and RMS (root-mean-square) average values, like shown in the SPECTRUM. For a sinusoidal wave the peak value is 3 dB higher than the RMS value and that is exactly what we found.

Hint: Repeat this relationship in your mind for a moment, because we will encounter this more often while working with the software. It is important to understand that some parts of the program show peak values, while others will use RMS values. It depends on the purpose which type of value is needed or used. Regarding measuring platforms it can happen that errors sneak into reports or analyses when it is not clearly distinguished between peak and RMS. If you talk about the level of a signal, make sure you always mention what kind of level you refer to.

Summary

We have introduced the input spectrum as another way to look at the data recorded by the microphone. We have talked about various display parameters like the weighting curve and the resolution. We have also introduced the measuring parameters FFT size and number of averages. In this respect we discussed about the trade-off between time resolution and frequency resolution, this relationship will become even clearer in the next section when we look at the spectrogram. Finally, we looked at the signal of an external sine wave generator in both time and frequency and verified the calibration.

2.3. Spectrogram

The SPECTROGRAM is the third type of graph available for input data. Switch the upper graph from TIME SIGNAL to SPECTROGRAM to see it. Rather than a line or bar graph the display now shows a color map that is continuously moving from the bottom to the top. But only at first glance it looks much different from the Spectrum graph that we investigated a little earlier.

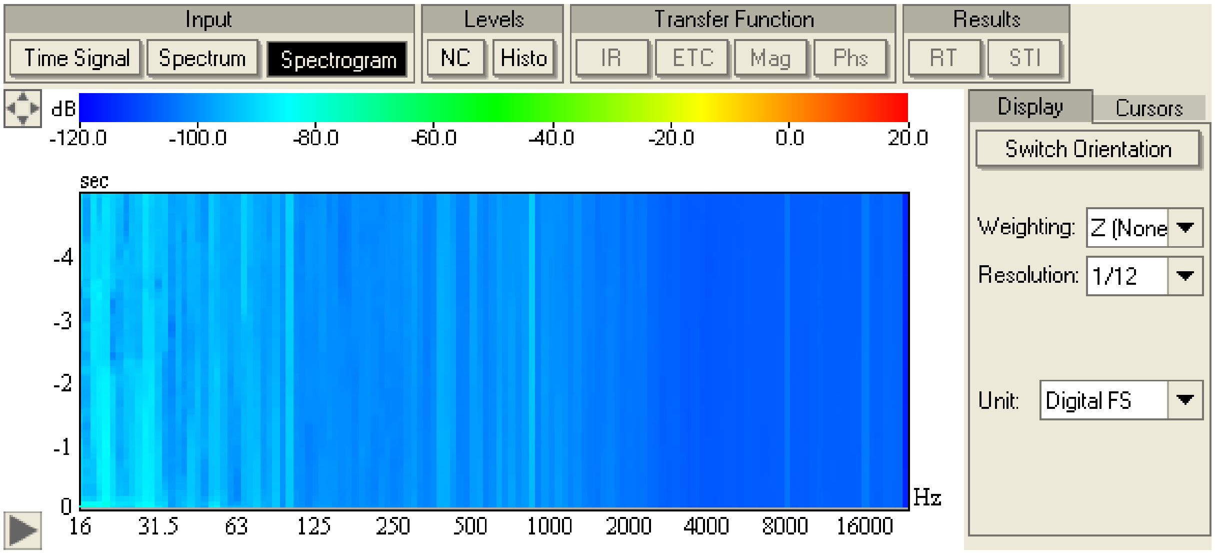

You will notice that on the right hand we also have the selections for WEIGHTING, RESOLUTION and UNIT. For the moment please choose the same settings as displayed here: NONE for WEIGHTING, 1/12 for RESOLUTION and DIGITAL FS for UNIT. Also return the measuring parameters to their original state, which was an FFT SIZE of about 3 seconds and just a single AVERAGE.

In the graph, the horizontal axis shows the frequency range just like the SPECTRUM. However, the vertical axis no longer measures level but rather the time that has passed by, from now to a defined point in the past. The level is mapped to a color scale and you will find that scale directly above the spectrogram itself. In contrast to the two previous graphs we looked at, the SPECTROGRAM can be considered to be a three-dimensional graph because it shows level as a function of frequency and of time. In fact, you can imagine it as nothing else than many SPECTRUM graphs in a chain and arranged depending on the time when they were recorded. Equivalently, you may consider the vertical axis as the time axis defined in the TIME SIGNAL plot, but now you can see the whole spectrum for each point of time.

But still, our reasoning from the previous section is true. There is no such thing as a high resolution of events in time and in frequency simultaneously. To make this clear, we will run a short test. First switch the displayed RESOLUTION to 1/96. After that, create some impulse-like signal at the microphone, for example, clap your hands.
What you will see as a result in the SPECTROGRAM is a signal that is precisely defined in the frequency domain and that is strongly smeared out in the time domain - note the very thin long vertical lines of constant color. These resolutions correspond directly to the FFT SIZE selected before.

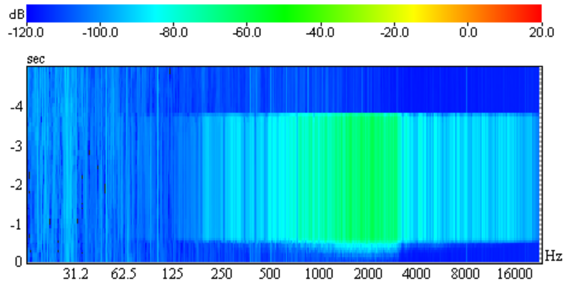

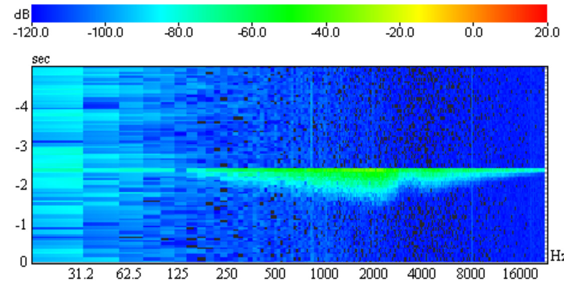

Now, change the FFT SIZE to something like 0.05 seconds and clap your hands again. You will see a signal that is well defined in time but has a rough frequency resolution. This is especially visible for the low frequencies where you can see long horizontal lines of the same color.

You may not see exactly the same results as printed here, naturally the actual levels depend on the sensitivity of the microphone, on your soundcard etc.. However, the default color scale is configured to be very wide, covering a range from –120 dBFS to 20 dBFS. To obtain a greater color resolution check the SPECTRUM graph at the bottom for the approximate minimum and maximum levels you would like to see. This might be –120 dBFS to –60 dBFS for 1/12th octave bands, the levels depend on the resolution of the frequency bands. To adjust the colors for that range, double-click on the color scale to open a window like the one below:

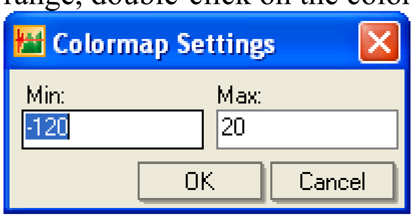

Now change the limits as desired and press OK. This will immediately rescale the spectrogram plot and use the new colors.

Hint: The ability to change the color range is very useful in practice. It allows you to "zoom in" to a range of levels and recognize differences more clearly. Similarly, you may select wide limits to suppress details and get a fast overview over the dynamic range of the input signal.

With respect to the frequency axis, all of the functions to select the view port can be used: mouse functions, manual entry, full view etc.. With regard to the time axis, the SPECTROGRAM is limited in so far, as it does not allow for a starting time other than zero. Consequently the right mouse button cannot be used to select a specific time range. Only the overall time length for the spectrogram can be entered manually, but not its starting point. It is noteworthy, that due to performance reasons the spectrogram data is not kept explicitly. As a result, if you zoom in into a frequency region you will lose the data outside of that range. Upon return to the full view blank areas will be shown for those frequencies where the history data was not recorded. Remember, that you can stop the SPECTROGRAM at any time using the button STOP ANALYSIS in the control panel on the left.

Only for the SPECTROGRAM there is a command called SWITCH DIRECTION which allows you to swap the vertical and horizontal axes. Depending on the type of investigation and the available screen area you might be interested in more screen space for the frequency domain or for the time domain. The default configuration is with the horizontal axis showing frequency. This makes it easier to align the view to the SPECTRUM when that graph is shown as a second display above or below the SPECTROGRAM.

As we outlined earlier, the main purpose of the SPECTROGRAM is to have a history of the input spectrum over time. On the one hand, this is very useful to obtain an impression of the spectral distribution of a sound event after it has already disappeared in the SPECTRUM graph. On the other hand, the SPECTROGRAM can be employed in the frequency domain to search and investigate resonances and find feedback situations.

Summary

This section explained the third graph that is dedicated to the signal at the input, namely the Spectrogram. Being a 3D graph it uses a color map to represent levels as a function of frequency and time. We have seen how strongly interrelated time and frequency resolution are.

3. Measurements with an Excitation Signal

So far we have only been concerned with rather "passive" measurements, which means we did not use a real excitation signal other than the sinusoidal from the microphone calibrator. In this part we will introduce and employ the stimulus signals included with EASERA SysTune to measure the frequency response of a system.

But before we start making measurements, let us look at the general setup. In addition to recording the acoustic signal, we now want to generate it as well. To convert a signal that is given in the digital domain into the acoustic domain, an output chain must be established that is similar to the input chain introduced earlier.

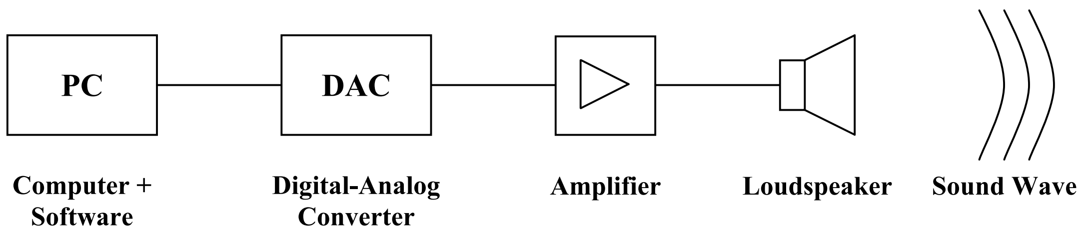

The software on the computer provides the driver of the soundcard with the signal to be played. Then the digital/analog converter of the soundcard creates a voltage signal from that digital data stream. The electric signal is amplified and passed to a loudspeaker which radiates a sound wave, finally.

For the next steps you will need to connect a loudspeaker to the output of your soundcard, preferably to the first channel.

Important: Make sure the gain of the amplifier or loudspeaker is turned down before using the software to play signals through the new output chain.

3.1. Excitation Signal

Choosing a Stimulus Signal

At first, we have to choose a stimulus signal in EASERA SysTune. To do that, go to the OUTPUT frame on the left and press the SELECT button right below the disabled PLAY SIGNAL button. (You will not be able to play anything unless you have selected the signal and the play channel as well.)

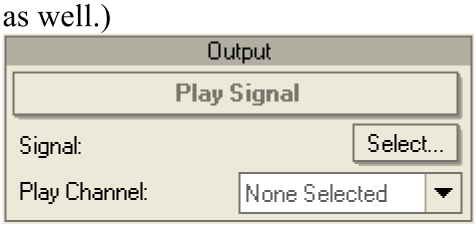

This will bring up the SIGNAL SELECTION window for EASERA SysTune.

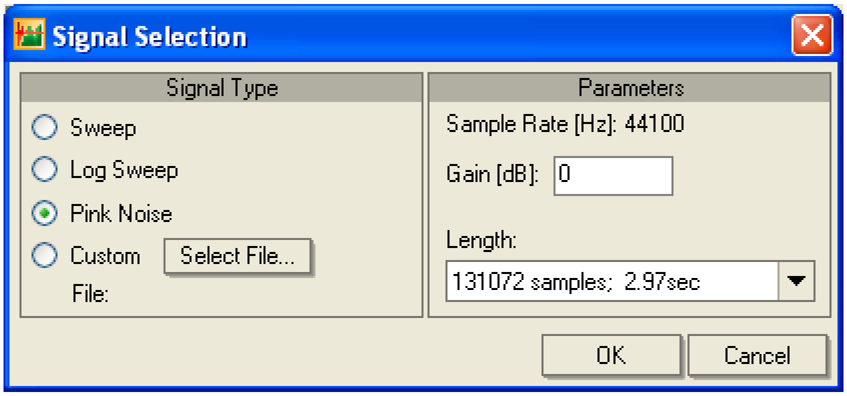

On the left, a Signal Type can be selected: You have the choice between a linear Sweep, a Log Sweep, Pink Noise or a Custom audio file. Here is a short overview:

- The linear sweep is a very common signal; sometimes it is also called a chirp or time-stretched pulse (TSP). Like all sweeps, this signal consists of a sinusoidal signal that changes its frequency with time. More precisely, the frequency of the linear sweep increases linearly with sweep time, this corresponds to a constant sweep rate (swept Hertz per second).
- In contrast to that, the logarithmic sweep signal spends the same time in each octave band. The advantage of the linear sweep, namely high signal power and thus a high signal-to-noise ratio for the higher frequencies is its disadvantage at the same time, because it does not work very well for more sensitive high-frequency loudspeakers. The power contents of the log sweep are significantly lower in the high frequency range, which is the reason why it is often preferred over the linear sweep.
- However, many times sweep-like signals cannot be employed since they are considered quite annoying due to their characteristic sound structure. Therefore, pink noise has become a common choice as well. It provides decreased signal levels for the higher frequencies and due to its continuous, evenly distributed level it finds more acceptance among occasional listeners and bystanders as well.
- As an alternative to these three options you may also load your own, user-defined stimulus signal as an audio file. To do that press the Select File... button and open a file from your hard disk. Note that due to memory limitations such files cannot be longer than 524,288 samples. At a 48 kHz sample rate this number of samples corresponds to about 11 seconds of time. If a signal file longer than that is selected, EASERA SysTune will automatically cut it to the maximum size possible upon loading.

Technically, you can use any kind of custom signal here, but you have to consider its purpose and the measurements you want to perform. The selected signal must provide enough contents in the frequency range of interest to achieve sufficient signal-to-noise ratio. With band-limited or even tonal signals, like some speech and music samples, it is likely that you will not be able to obtain a usable broadband frequency response (or impulse response) of the system under test. Although the predefined signals shipping with SysTune may not be ideal to present to an audience, they are highly effective for measurements. Please also see chapter 5.4 where we discuss the use of speech and music as stimulus signals.

Hint: If you are not certain about the spectrum of a custom signal, look at its frequency response at a high resolution. (We will discuss a bit later how to do that in SysTune.) If you find deep gaps of significant width within the frequency range of interest, do not use that signal for the measurement.

If you do not have the desired excitation signal available as an audio file, but only as a track on a CD or if you want to utilize any other signals external to the software program, you can play them right away and watch the result at the microphone input in SysTune. For transfer function measurements with external signals please refer to the next chapter of the tutorial.

Tech-Note:
The measurement world is rich with tools and many of them use their own, custom-built excitation signals. Besides the linear and logarithmic sweeps, there are sweep signals with other weightings too. Time Delay Spectrometry (TDS) uses a sort of linear (or white) sweep. For loudspeaker measurements, so-called weighted sweeps have been introduced by EASERA. These sweeps are attenuated by about 20 dB for the high frequencies, but they do not drop as far down as the log sweep does.
Also, maximum length sequences (MLS) are well-known stimulus signals. They are comparable to random white noise but have a deterministic character that allows them to be processed very quickly. This was especially important some years ago when computing power was more limited.

For real-time measurements the signals that are most often useful provide a kind of self-similarity. That means, if you look at a section of the signal that is much shorter than the signal itself, the spectrum looks similar. This is not the case for any kind of sweep signal, but is the case for all typical noise signals. This property of self-similarity allows you to choose time frames, that is, FFT sizes, for investigation times shorter than the time duration of the signal.
But sweeps have advantages too, specifically when measuring slightly unstable systems. We will come back to the choice of the right signal, when we talk about the measurement of transfer functions.

On the right hand side of the Signal Selection window you find additional Parameters to be selected. At the top, the current Sample Rate is shown. Below that you may adjust the output Gain in dB for any type of signal. Because SysTune does not provide any level controls in the software, this parameter is useful if you do not want to touch the gain settings of the soundcard or amplifier. At the bottom of the Parameters frame, you can also select the Length of the signal in seconds. This option is not available for Custom signals, because their length is read directly from the audio-file.

For our purposes here, select Pink Noise, enter a Gain of –6 dB and choose a Length of about 3 seconds from the drop down list. Press OK to confirm these settings. After return to the main window your choice is immediately reflected by the software. It now shows the type and length of the signal close to the Signal label.

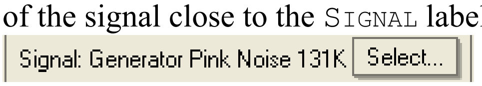

If you still have the short `FFT Size` selected that we used in the `Spectrogram` section, the program will show a warning like the one below:

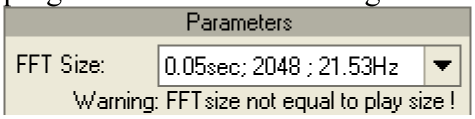

Generally, your `FFT Size` should never be larger than the time period of the excitation signal, if the signal is periodic. The `FFT Size` can be shorter than the signal periodicity if a smaller section of the signal still covers the frequency range of interest. For example, pink noise can be used with FFT sizes shorter than the signal length, but sweeps cannot. For the next steps we will switch back to an `FFT Size` of about 3 seconds, it should have the same length as the pink noise signal we just selected.

Activating the Output Channel

Secondly, we need to choose the output port for the stimulus. This happens using the `Play Channel` drop down list; it is located directly below the area where we have just selected the excitation signal. This list shows `None Selected` as a default, we have to change that to the desired output channel, in our case to `1`. If you are not certain which channel is actually connected or if you want to feed several loudspeakers with the same signal output, choose `All Channels` instead.

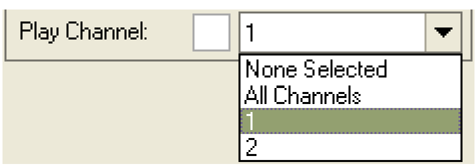

Later on, you might be seeing more than just the ordinal numbers in the list of output channels. For some drivers, SysTune will also show the name for each channel if available.

You may have noticed that there is a small box just between the `Play Channel` label and the selection list. It is another mini-meter, similar to the ones we introduced for the input ports. But this one monitors the output signal. It will warn you in the same manner if the output signal gets too close to the clipping thresholds and might be distorted.

Frequency Response Measurements

Now, we are all set to make some measurements. Once more check your level settings and press `Play Signal` to start the output.

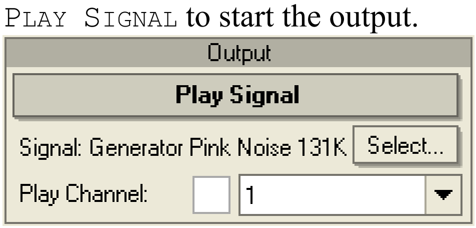

When SysTune plays a signal to the output, the mini-meter will awake and show the current level. Like the input, the meter shows yellow when the signal level is above –6 dBFS and it shows red when you are close to clipping, namely at -1 dBFS and higher. Again, the frame remembers the last threshold passed and can be reset by a mouse click on the meter.

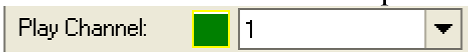

Hint: Make sure you are always safely in the green area. It depends on the sound hardware and the quality of the digital/analog converters how high the digital signal amplitude can be. For some configurations, you may need to be below –6 dBFS even with the peaks to ensure that the conversion process does not create any distortion at the output. This is especially important for noise-like signals.

Switch the top graph to show the Time Signal and the bottom graph to Spectrum for the following discussion. If you are calibrated you should also see the level at the microphone displayed in Pa and dBSPL, as shown below.

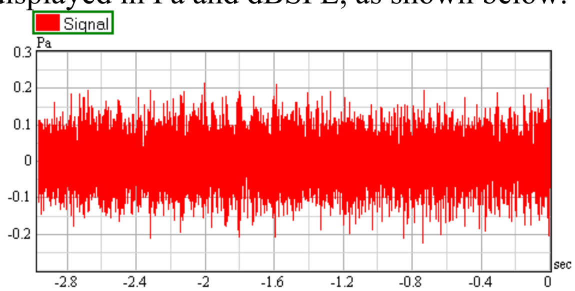

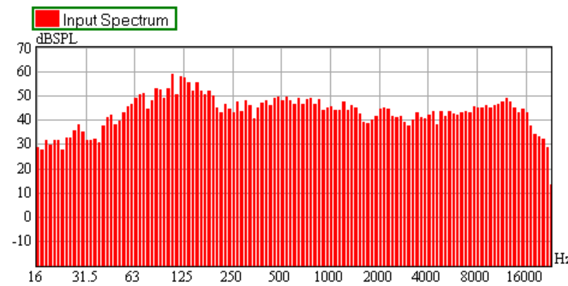

Naturally, the graphs on your screen will look a bit different from these pictures, depending on your exact measurement setup. But in any case the Spectrum shows the actual frequency response of the system under test. As mentioned before, this is one of the characteristic properties of pink noise, when it is summed in fractional octave bands like in the view here. Because the signal itself produces a flat spectrum, what you see is caused by the system under test. For magnitude-only evaluation you do not need a reference channel.

You may move the microphone or loudspeaker a bit to see the effect on the measurement. You can also change delay, EQ or gain settings for the loudspeaker, but remember that changing any gain on the input side will invalidate your calibration.

Summary

In this section we have taken a first look at the output side of the measurement chain. We selected the stimulus signal and the output channel to excite the system under test. Finally, we briefly reviewed the graphs showing the signal that is currently received at the input.

3.2. Capturing and Comparing Measurements

Having created a working setup for measuring the system under test, we would like to make several measurements, store them and compare them. For easier viewing, disable the Bar Display and maximize the Spectrum graph to full screen. If the Spectrum graph is at the bottom, so use the following tool button from the menu bar:

The tool button just to the left of it can be used to maximize the top graph:

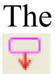

When you maximize a graph the window will change to a layout like the following:

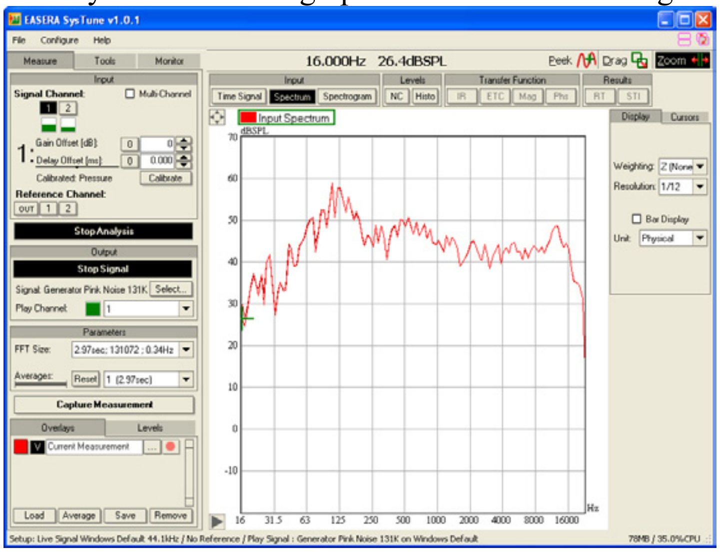

The visible graph may be toggled using this tool button:

To later get back to the split view use the tool button:

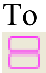

Now, let us capture the current measurement. Press CAPTURE MEASUREMENT once - you will find the button in the left control panel below the PARAMETERS section. Immediately, the software creates a copy of the current measurement data and adds it to the list of OVERLAYS directly below the button.

By default, the new item is named “OV1” and a color is assigned to it automatically. Afterwards the overlay list and the graph will look approximately like this:

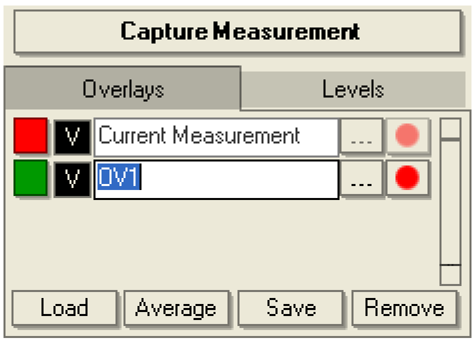

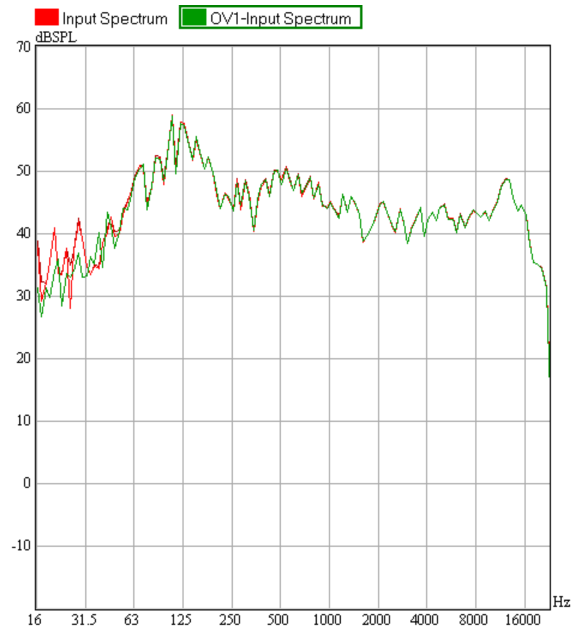

The captured data is shown in the same graph, overlaid on the running measurement. Depending on the circumstances you may also see the peak hold curve for the CURRENT MEASUREMENT. Feel free to stop the analysis for a moment and start it later again when we need to make some more measurements.

Overlay Properties

To change the name of the measurement, click in the text field and enter a new label. Because we are going to capture measurements for a few different microphone positions let us call it “Mic Pos 1”.

To the left of the text field there is a button labeled V for visible. Press the button to toggle the visibility of a particular overlay curve. You can even hide the current (running) measurement.

At the very left, there is a second button that shows the color of the overlay curve, you can click on it to change the curve color in a small dialog window.

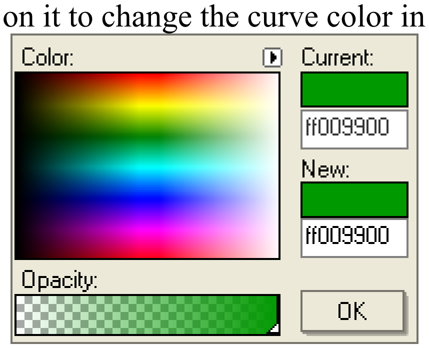

In addition to picking the COLOR directly from a map you may also enter a hexadecimal value under NEW and define the OPACITY of the curve as well. To close the dialog window, press OK.

To the right of the text field there are two more buttons. The first one, labeled ..., opens the OVERLAY PROPERTIES window. It consists of two tabs. The first tab DISPLAY allows you to add some Comments, to change the Name of the overlay curve, to enter an additional Gain value in dB as well as to select the primary Color, which we already discussed in the paragraph above. The secondary color button labeled Color Sec is the curve color used for displays that include the reference channel. We will come back to that a bit later.

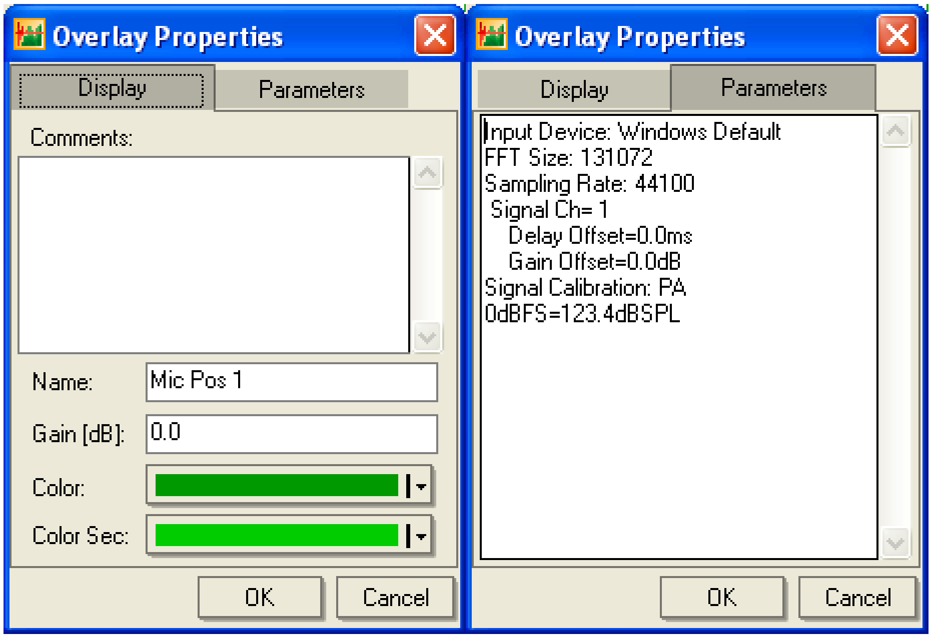

The second tab Parameters shows the software configuration from the time when the measurement was captured. Right now, we do not want to change anything here in this window, press OK to close it.

Back in the overlay list, the rightmost button with a red record symbol allows you to overwrite the particular measurement. When you press this button, a new data set is captured from the running analysis. This command is similar to Capture Measurement, but it stores the data directly in the selected overlay instead of creating a new list item.

Hint: This function is particularly useful if you have multiple measurement locations. After a first series of measurements, you can update entries individually if the acquired data was erroneous. You may even prepare a measurement sequence by capturing empty (nonsense) overlays first, label them and then replace them step by step with real data.

Active Curve

You may have noticed already that the graph now includes a new curve, and its legend was also extended to display the name of the additional overlay. Furthermore, there is a green frame around the label. It indicates which curve is active and thus supports several special functions. At first, the active curve is always drawn in the foreground. You can switch between curves by simply clicking on another label in the legend. Also, the tracking cursor refers to the active curve, which means that the coordinates displayed in the mouse bar are always related to the active overlay.

There are a few more means to switch the active curve, for example, by a click in the corresponding text field in the Overlays list. Alternatively, you can use the + and - keys to switch between curves. Note that for this last function the graph where you want to change the active curve must be active, only then can it accept key strokes specific to the graph. To activate a graph, simply click on it, it will then display a red vertical bar at the very left to show that it can receive key commands.

The following two pictures show the SPECTRUM graph with the CURRENT MEASUREMENT active and with the captured measurement “Mic Pos 1” active. For the latter, the graph has been activated also, so that key strokes for that graph can be used. Finally, you can see that the tracking cross is always located on the active curve.

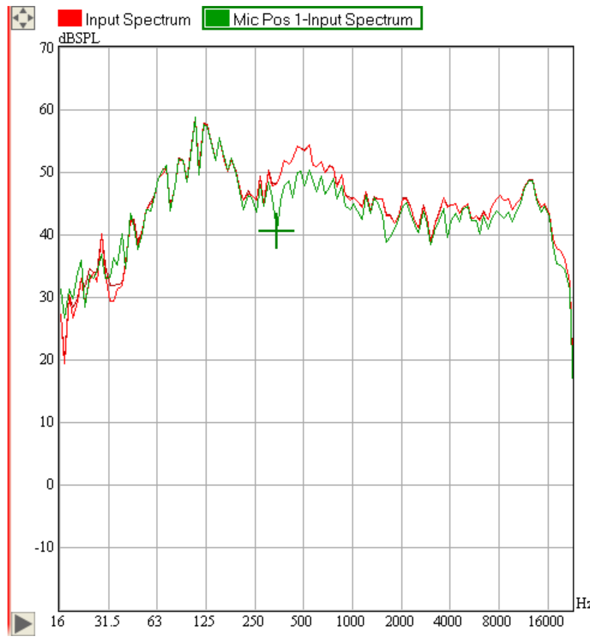

Hint: Sometimes you will be working with a set of measurements in a crowded display. In such a case, a particular curve might be hidden by some other overlaid curves. By activating each of the curves one by one you can browse through them and investigate them more closely.

Saving, Removing and Loading Overlays

In many cases, captured measurements must be saved for further evaluation, processing or comparison. Below the list of overlays, you find a row of buttons labeled LOAD, AVERAGE, SAVE and REMOVE. To save a measurement from the overlay list, just press SAVE and select the overlay that you want to write to a file. All curves that have not yet been saved will be automatically selected when the window is opened. Then press OK.

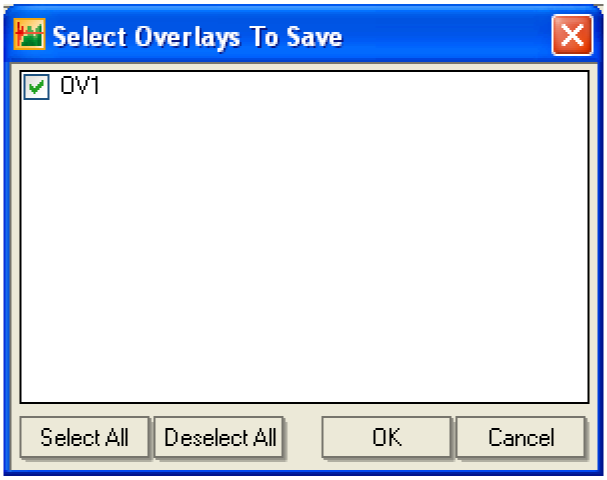

To clean up the list of overlays, you can do a similar thing. Press REMOVE to open a list of measurements to remove from the internal memory. All curves that are not visible will be automatically selected when the window is opened. Then press OK. Before that, make sure that you have saved the measurements if they contain valuable information.

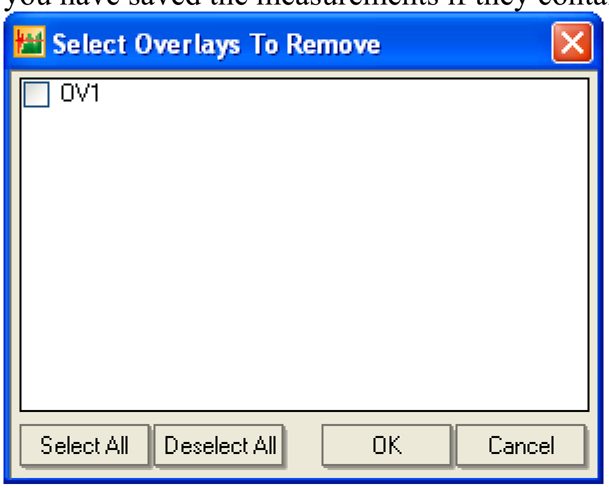

To load measurements into the overlay list is straight forward, too, press LOAD and select a file from your hard disk.

Averaging Measurements

The AVERAGE button allows you to calculate an average from multiple overlays. To use this command in a meaningful way, we would like to have a few more measurements. So, change the measurement setup a little bit, for example move the microphone to another location, capture a second overlay and label it "Mic Pos 2". Repeat the same and create an overlay with the name "Mic Pos 3". Also, assign colors to the curves as you like. You will finally end up with a display similar to the one shown below.

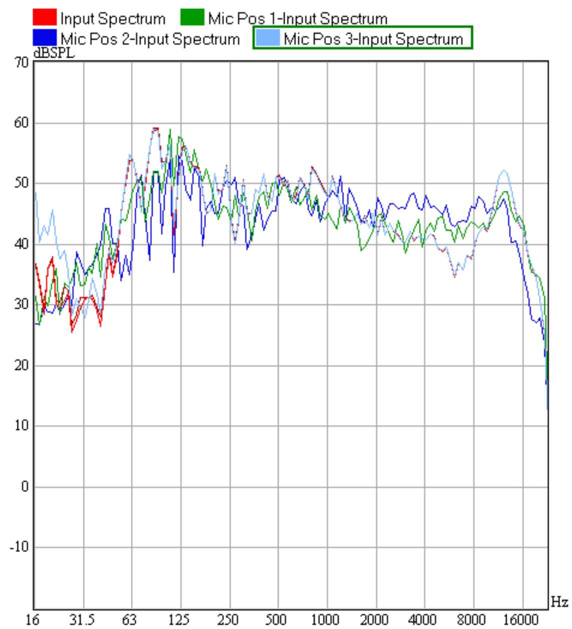

Now we would like to create a representative frequency response for the covered microphone positions. Press AVERAGE to open the dialog for that and select all three overlays for processing. After you have confirmed with OK the program will calculate the average spectrum for the selected measurements and add the result as a new item to the OVERLAYS list.

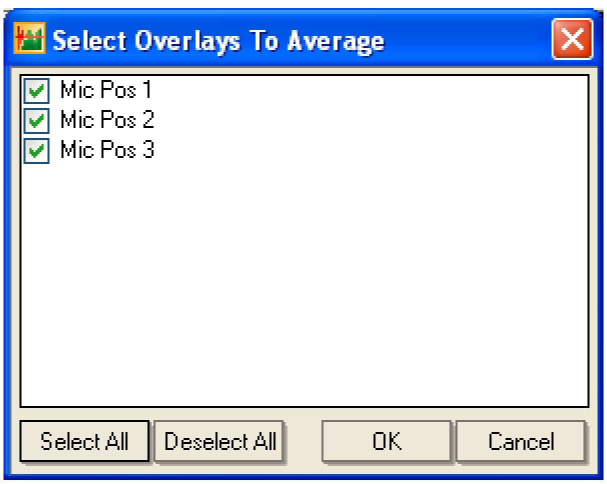

After that, the SPECTRUM graph will look similar to the picture below. Have a closer look at the average curve and how it relates to the other overlays. You can give it a name and a different color, if you like. To store it for later use, save it to a file, just like we have done before.

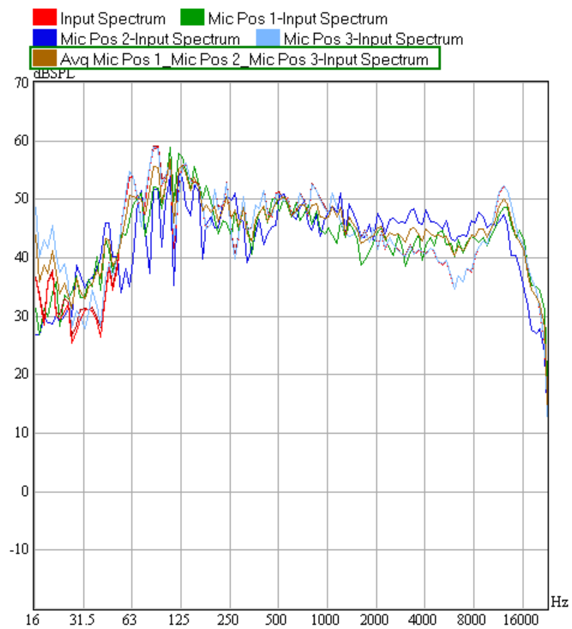

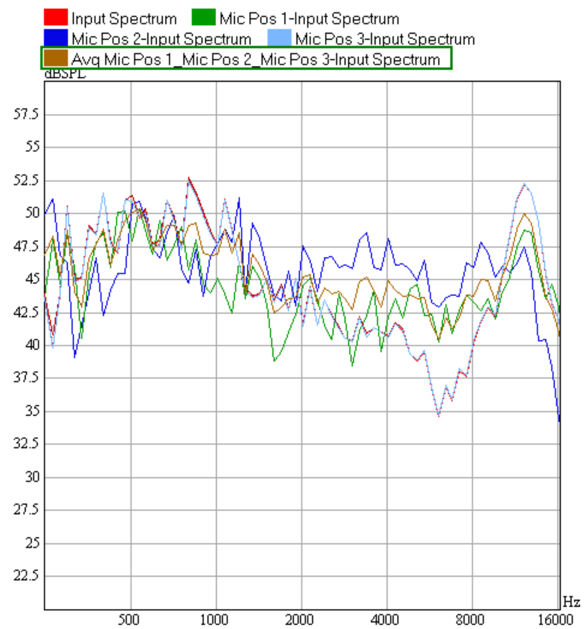

Tech-Note:
The average that we have just computed is a power average, which means that the power spectra of the individual measurements have been averaged. In a logarithmic display such an average will be dominated by the maximum curve for any given frequency data point, so that the resulting average curve is not exactly centered among all curves.
The power average is the most common way to average data in our application domain. In contrast, complex (or vector) averaging as well as magnitude-only averaging assume coherent signals and cannot be used to derive a representative frequency response of a room or listening area. They find more usage when electronic signal chains are involved.

Capturing, loading and averaging overlays all require PC memory for storage. When using overlays extensively, at some point in time you will encounter the following message:

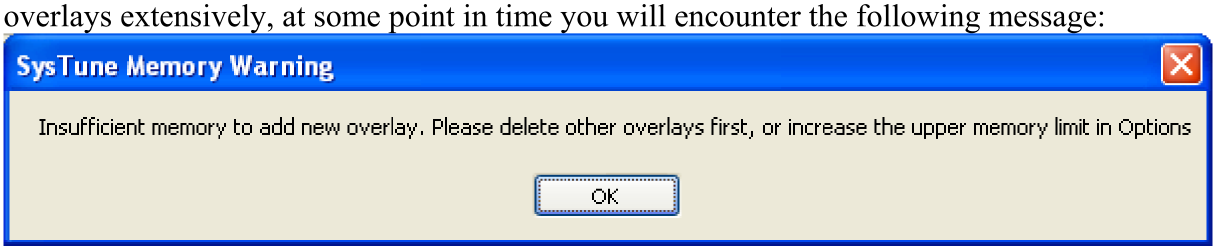

This happens when the next step is going to exceed the maximum memory that is available for overlay curves. In this case, it is recommended that you save and remove some of the measurements from the list. The maximum memory is set to 80% of the physical RAM available; this default can be changed in the OPTIONS window, which we will look at more closely in one of the following sections.

For the next function, we should clean up the overlay list a little bit. So, remove all of the curves you have added during this exercise and leave just one data set and the CURRENT MEASUREMENT.

Adding Cursors

When we introduced the mouse bar - the bar that shows the coordinates of the tracking cross and the tool buttons to access the mouse modes ZOOM and DRAG - we left out the third mouse mode PEEK. At this time, we want to make up for that and introduce you to the powerful functions associated with PEEK. To explain all of these features clearly, you should have only a single overlay in addition to the running measurement, displayed in a maximized SPECTRUM graph (If you forgot how to get there just follow the first part of the previous section once more). We also want to look at a band resolution of 1/3rd octaves with no weighting. You should see a picture like the following:

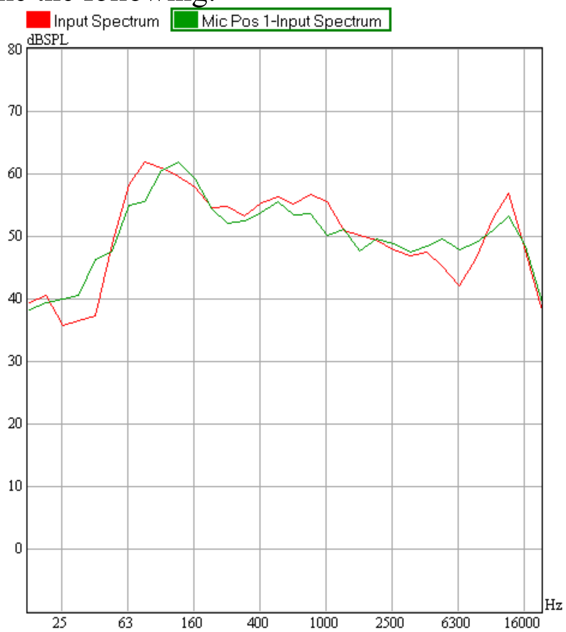

Now switch to the mouse mode PEEK using the button located in the mouse bar. You may have noticed that for a maximized graph the mouse bar is placed right below the menu.

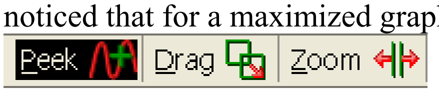

As in all mouse modes, you have the tracking cross that shows the mouse coordinates projected on the active curve. In PEEK mode, when you drag the mouse while keeping the left mouse button pressed, a display will appear that shows the differential between the starting point and the current cursor location for both vertical and horizontal axis.

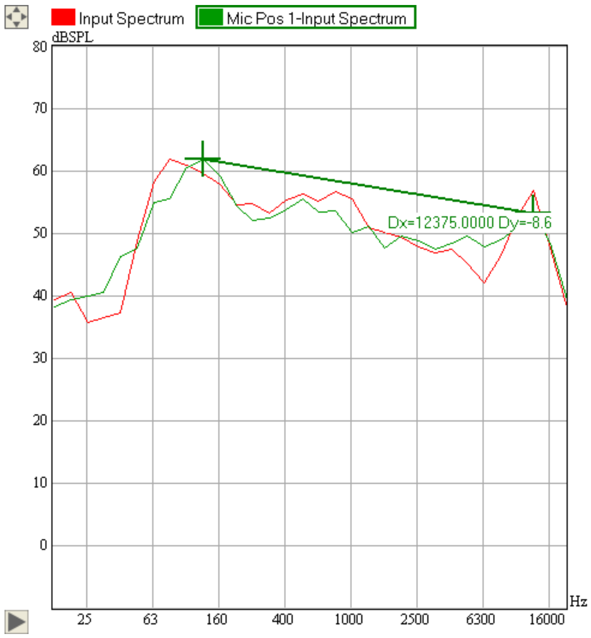

Hint: With the Peek function you can even read out level differences between two curves. Use the +/- buttons to toggle the active curve while you hold the left mouse button. The tracking cross will automatically switch from one curve to the next. Note, that these short cuts only apply to the active graph, which is indicated by a red vertical line.

To evaluate and compare multiple locations, use the right mouse button. With each click a fixed cursor is placed in the graph. It shows the position of the cursor on the horizontal axis as well as the magnitude value of the vertical axis for each overlaid curve. The color of a numerical value corresponds to the color of the associated overlay curve. You can insert cursors at any time, with the analysis stopped or running.

To practice this, try to insert 3 cursors at 125 Hz, at 1000 Hz and at 8 kHz. The final graph with cursors will look approximately like in the following picture.

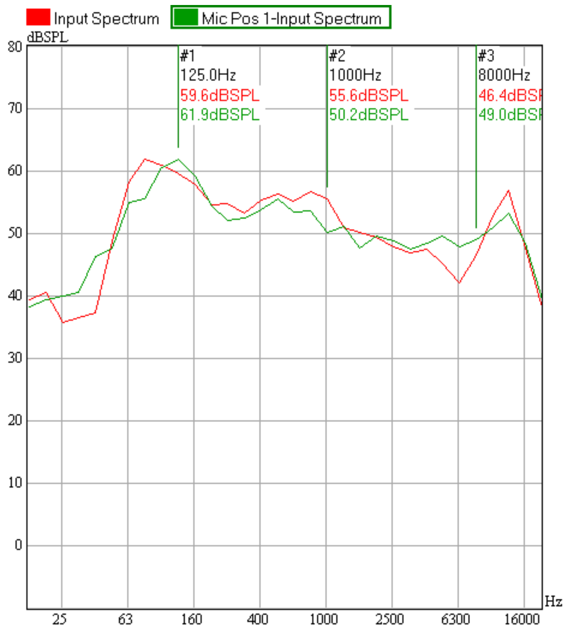

If you have not captured exactly the right frequencies, it is easy to start over again. For that, look at the CURSORS list on the right, it is hidden on the tab behind the DISPLAY tab. It shows all of the cursors inserted into the graph:

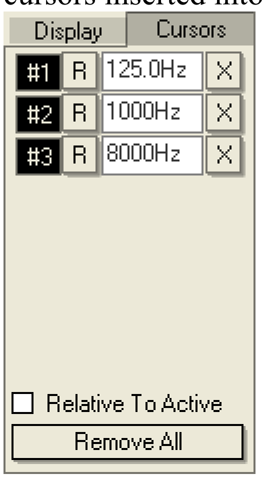

To clear the list, simply press REMOVE ALL and the cursors will be deleted. Individual cursors may be removed using the X button on the right. You can toggle visibility of a particular cursor without deleting it by clicking on the #NUMBER button on the very left.

But the most powerful function hides behind the R button. It allows you to select one cursor to be the reference for all others. What that means is that with a reference cursor all other cursors display their magnitude values relative to that reference. For a moment, make the first cursor the reference and you will see a graph like this:

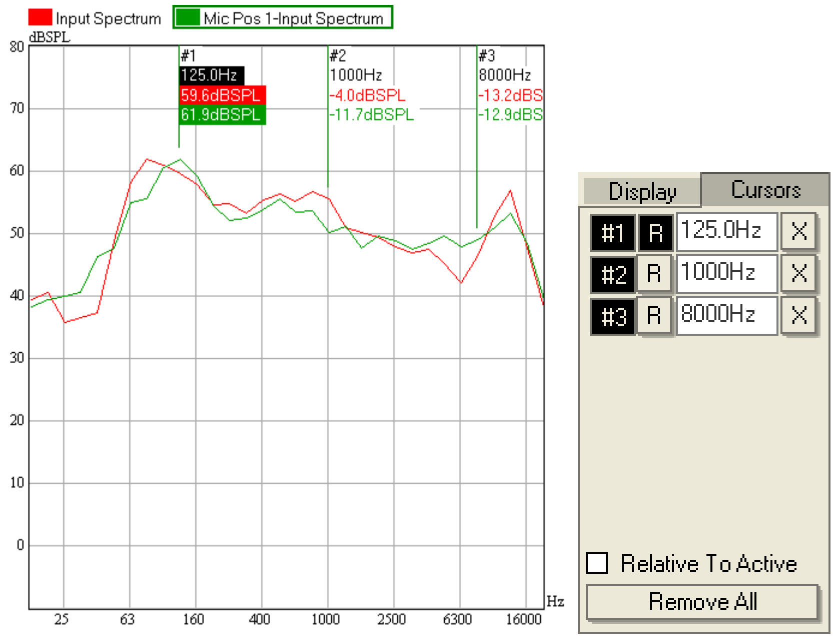

The second and the third cursor have switched from absolute levels to relative levels. Note that all relative values always show a sign, either + or – in front of the numerical value. In the graph, the current reference cursor can be identified by the solid background.

Hint: The differential cursor display is very helpful to determine level differences along the horizontal axis. For example, in the frequency domain you may easily investigate the relative level of the harmonics when you feed a sine wave signal to the system. In the time domain, you can check exactly how pressure or voltage decrease or increase over time.

Another important function can be activated by enabling the checkmark RELATIVE TO ACTIVE. While the reference cursor function displays magnitude values relative to a selected cursor, that is a point on the horizontal axis, this option calculates the values relative to the active curve. A typical display with the first overlay as the reference curve is shown below.

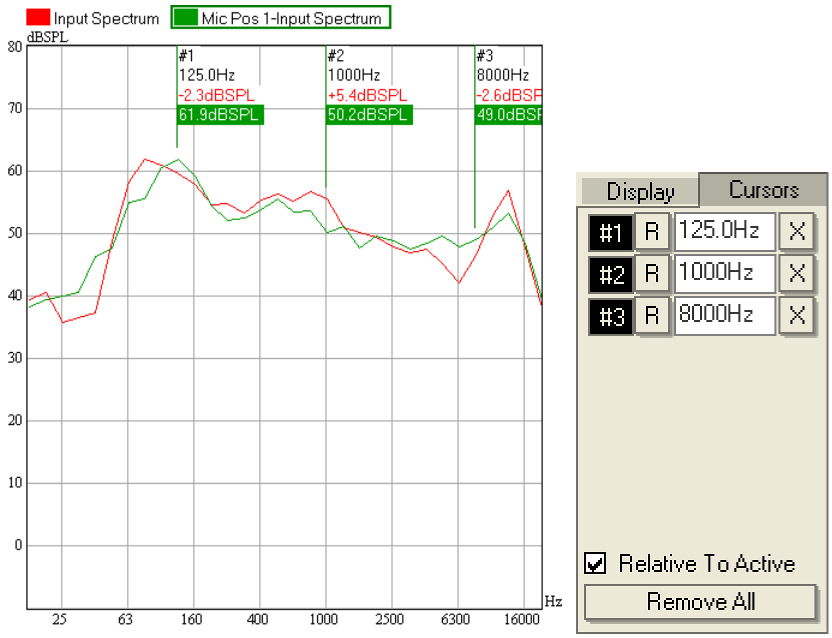

As a result, instead of absolute levels the red curve labels show levels relative to the green curve. Also here, the magnitude values that are used as a reference have a solid background. Click on a different curve in the legend to change the active curve and thus the reference curve.

Hint: Displaying cursors relative to the active curve is particularly useful if you would like to compare levels of several measurements or monitor the running measurement relative to a loaded file. You may even use both differential displays, reference cursor and reference curve at the same time.

At the end of this part, we would like to emphasize that the overlay and cursor functions are available for all 2D graphs in EASERA SysTune. Repeat part of the above steps or all of them for another graph if you would like to practice a bit.

Summary

We have visited and discussed some of the very core functions of EASERA SysTune in this section. We learned how to capture measurements, label them and manipulate them. The concept of overlay graphs was introduced and the concept of the active curve as well. We have seen how stored measurements can be saved to and loaded from the hard disk and how they can be removed from the program's memory. Also, we have calculated the power average from a set of measurements to obtain the representative frequency response. Finally, we used the Peek mouse mode to compare curves and to add cursors to a graph.

4. Dual-FFT Measurements

This chapter is concerned with the main measuring functions of EASERA SysTune. While the previous sections covered basic measurements with just a single input signal, we will now look at measurements using two channels. One channel will be the signal channel and the second channel will be called the reference channel. Note that the reference channel can be established in two ways:

- as a virtual loopback in the software domain, that is, using the signal played to the output as the reference signal, or
- as an external signal, provided to the software through a second input port.

Let us shortly review these two configurations, before we discuss how an impulse response and transfer function can be derived from that.

Setup with Internal Reference

With an internal reference, we will use the software as a signal generator and feed the signal back to the software after it passed through the system under test. The reference channel is purely virtual, you can imagine it as a loopback linking the software output to the input. For this type of measurement the reference signal is exactly the same signal as the one sent to the physical output of the soundcard.

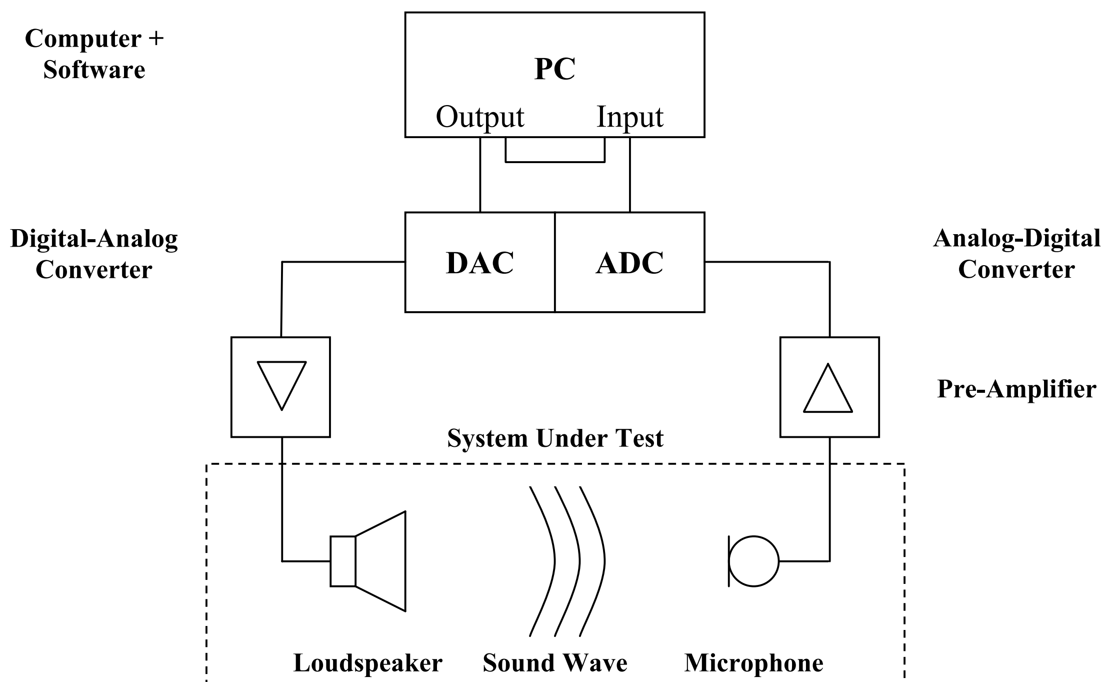

This type of configuration is most useful for transfer function measurements when there is only one physical input and output channel available. If the soundcard (DAC, ADC) allows for more channels, it is generally recommended to realize the loopback electronically, which means with a direct cable connection between another input and output. This has the advantage that the effects of the DAC and ADC, such as latency or frequency dependence, are automatically removed from the measurement of the transfer function. Of course, this only makes sense if you do not consider the converters as a part of the system under test.

Setup with External Reference

With an external reference, we use a second physical input to provide the reference signal. This can either be an electric signal, for example from a console, or it can be an acoustic signal, from another microphone in the room. It can also be the physical output of the soundcard. If set up in this way, the software can still be used as a signal generator, but it will not directly take into account the signal played. It will only use the reference supplied externally.

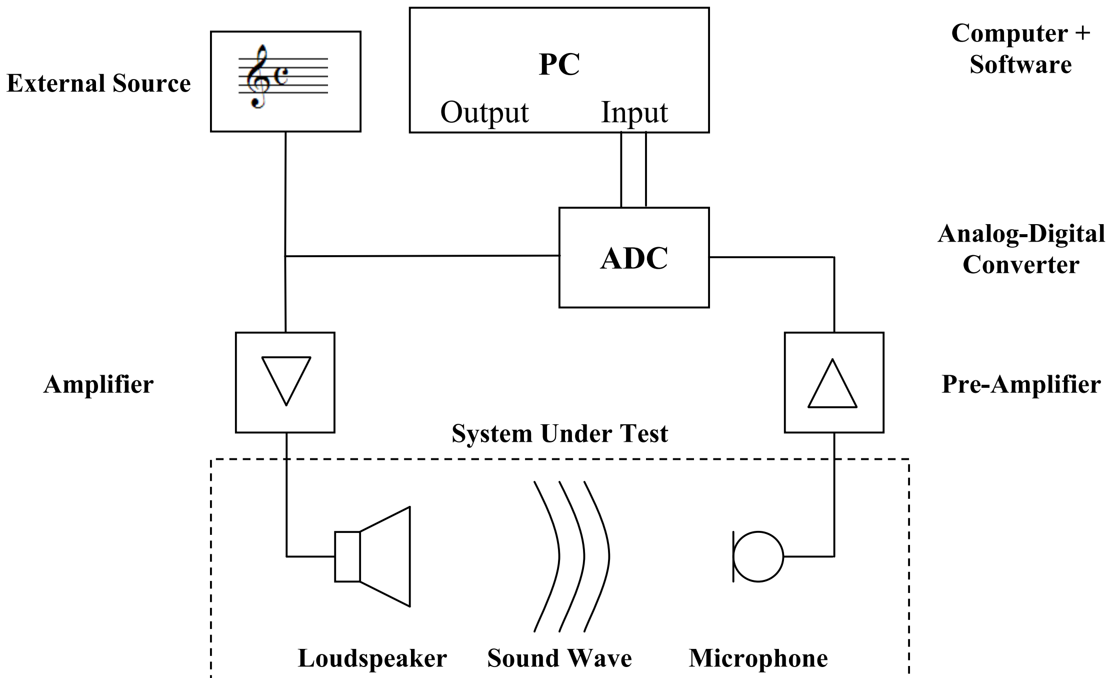

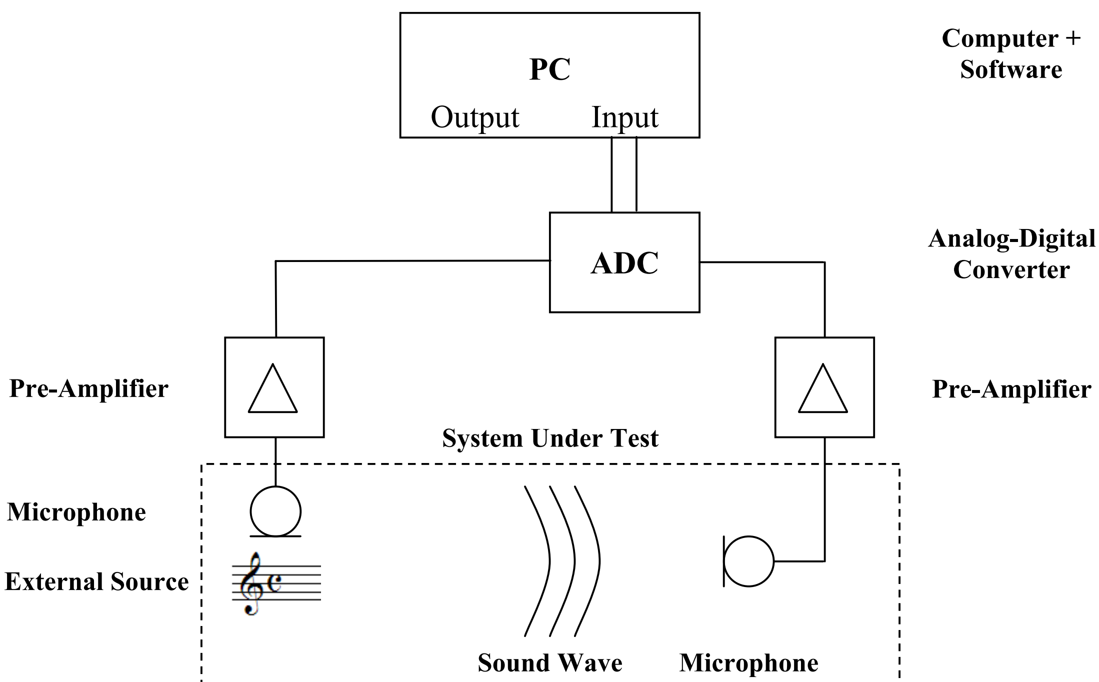

4.1. Transfer Function Measurements

Given a measurement setup that provides a signal channel and a reference channel we may now derive the transfer function for the system under test. But before we go into the details, let us once more clarify the term transfer function. It is generally used in two different ways and to prevent confusion the following should be emphasized:

- In general system theory, the transfer function is a fairly abstract entity. It can be any kind of mathematical expression that describes the output of a system in response to an input. This specifically includes all kinds of domains, like frequency and time, etc..
- In our field, the transfer function is used mainly with a more precise meaning, namely that it is a complex frequency response which describes the spectral response of the system under test to an ideal, impulse-like input. The impulse response is considered as the equivalent of the transfer function in the time domain.

In the following we shall only refer to the second definition, it is more commonly used and it will not lead to misunderstandings.

Measuring Principles

To explain measurements with either of the two setups above, we want to repeat briefly what an impulse response and a transfer function is in our sense here. First of all – what are we measuring? So far we have spoken rather mysteriously about the system under test (SUT).

We need to define more clearly that we always mean linear, time-invariant systems. In this respect, linearity can be simply understood to be, that a linear system responds to a signal at its input with a signal at its output that is proportional to that input.

The output can be higher in level, equal or lower, but the relationship is always linear:

Output = Input x Factor

This basically means that the relationship or the factor between output and input does not depend on the level of the input. In the same way, time-invariance means that this factor does not change over time as well. It has exactly the same value whenever a signal is applied.

For better understanding, here are a few examples:

- A system clipping the signal at a threshold is not a linear system
- A system with a feedback loop is not a linear system
- A system with a high noise floor is not a linear system
- A system with dynamic EQ is not a time-invariant system
- A venue with a sound reinforcement system and a strong wind blowing is not a time-invariant system
- An electronic enhancement system (reverberator) with continuously changing filters is not a time-invariant system

Another point to add here is that in acoustics, a system can not usually be described only by a single factor of proportionality. In fact, the factor above is normally a function of frequency. This function is what we call the transfer function of the system. It should also be emphasized that this mentioned factor can be a complex number, consisting of a real and imaginary part or of magnitude and phase. The equivalent of the transfer function in the time domain is called the impulse response.

Tech-Note:
The process that the linear, time-invariant system applies to the input signal is called convolution in system theory. The input signal is convolved with the transfer function (or impulse response) of the system to yield the output signal. In this sense, the convolution is equivalent to a complex multiplication in the frequency domain and to the numerically more difficult evaluation of the convolution integral in the time domain.

At this time, we remark that in practice, systems can only be approximately linear and time-invariant. You should be aware of the fact that the accuracy of the measurements you make with SysTune also depends on how close the system under test is to a linear, time-invariant system.

Computation of the Transfer Function
All that being said, how can we actually derive the transfer function from a measurement? The principle idea is very simple; we take the equation defined above and simply resolve it for the factor – for our transfer function:

Factor = Output / Input

What this means is that we just need to feed the system under test with an input signal and measure the output signal. If we know both of them we can immediately calculate the transfer function of the system. It is also obvious, that we cannot derive a transfer function for frequencies where there is no input to the system under test, or simply put, we cannot divide by zero.

We denote that for all this seemingly simple math we need to be in the frequency domain. The direct derivation of the impulse response (transfer function in the time domain) is possible, but more difficult and costly with respect to performance.

Tech-Note:
The process of deriving the transfer function (or impulse response) of a system from the input and output signal is called deconvolution in system theory. It is equivalent to a complex division in the frequency domain.

Measurement platforms based on the deconvolution principle can only derive the transfer function of the system for frequencies that are excited by the input signal. If the stimulus is band limited or contains only singular tones, some frequencies are not included. For these frequencies the deconvolution would correspond to a division by zero which is not defined.

In EASERA SysTune this problem is resolved by applying an adapted WIENER algorithm. It basically puts a threshold on the input signal and disregards any data for frequencies too low in level. Of course, this algorithm only overcomes formal problems associated with those undefined frequency bands, it still can not derive the system transfer function for them.

Because of the complexity involved with the time domain, EASERA SysTune calculates the transfer function in the frequency domain. We have already seen in the first chapters that time data is acquired and transformed into frequency data using the Fast Fourier Transform (FFT). Exactly the same method is used here, but twice, namely for the signal channel and for the reference channel. For both of them, the FFT block is taken, its spectrum calculated and finally the signal spectrum is divided by the reference spectrum to obtain the transfer function. Based on the transfer function then the impulse response can be calculated using an inverse FFT (iFFT).

Tech-Note:
The deconvolution concept introduced above is applied to each FFT block of input and output. But it is important to emphasize that for most real-world systems the output is to some extent delayed compared to the input and this has to be considered for the deconvolution.

Imagine that the audio streams for the signal (output of system under test) and for the reference (input to system under test) are continuously cut into FFT blocks of equal size. For both signal and reference, the start of each block is initially located at the same point of time. After that, the complex division is applied as described earlier. Under such circumstances, if the output is now slightly delayed compared to the input, it is clear that some samples are lost for the deconvolution. The FFT block of the output will be missing the very last samples of the processed input at the end, and it will contain some samples belonging to the previous block at the start instead. With increasing delay time, the correlation between the input and output blocks decreases. For a delay that is larger than the FFT size, the correlation is lost entirely (unless the signal fed is periodic).

It is obvious that for maximal measuring accuracy the actual delay must be found in order to adjust the two channels against each other time-wise. Only then, the FFT blocks for signal and reference will be optimally correlated. In SysTune this is a very simple process: First you usually measure with a long FFT size in order to determine the delay. This value can then be entered for any subsequent measurements, whether it be with longer or shorter FFT size. This ensures synchronicity of FFT blocks from input and output.

Real-Time Deconvolution (RTD™)

There are various measuring systems available and many of them can derive an impulse response and a transfer function for the system under test. Currently however, only EASERA SysTune is capable of calculating the full impulse response of a real-world system in real-time.

There are two aspects to that. On the one hand, SysTune has to calculate the spectra of the signal and reference channel via the FFT and then apply the iFFT to obtain the impulse response. On the other hand, this process not only has to be done once but continuously and at a high refresh rate.

On standard personal computers, this functionality is only possible by exploiting modern processor architecture and utilizing the fastest compiler available. Extensive use of multi-threading guarantees that the graphic user interface remains responsive while these enormous calculations are performed in the background.

Tech-Note:
Having been developed for high-performance applications, EASERA SysTune exploits much of the low-level functionality available in modern processor architecture. This specifically includes SSE instruction sets and vector optimization.

To keep the graphic user interface (GUI) responsive even when significant performance is needed to run the real-time deconvolution, EASERA SysTune relies heavily on multi-threaded programming. This means that processes for audio streaming, data transfer, calculation of results and user interaction work independently from each other and are distributed over several processors if available. The resulting gain in performance and comfort on multi-core processor systems is remarkable, although the development of multi-threaded applications is significantly more difficult.

Transfer Function in EASERA SysTune

In SysTune the transfer function and the impulse response are presented in several ways:

- In the time domain, you can view the impulse response directly. This is an important function for aligning delays of sound sources but also to investigate how the system responds to an input signal time-wise.
- Also in the time domain, you can view the log-squared impulse response as a curve displayed in dB. Called an energy time curve (ETC) in SysTune, this presentation relates better to our hearing system and lets you identify reflections for example.
- In the frequency domain, the magnitude of the transfer function is available as a graph. It is probably the most important display for tuning loudspeaker systems, because it gives you insight into how the system under test actually transfers a signal. A fairly flat frequency response is one of the criteria for a well-tuned system.
- Another aspect of the transfer function is the phase, it is available as a graph as well. The phase response is particularly useful to analyze frequency-dependent delay times which is not directly possible with the broadband impulse response.
- Based on the impulse response, some basic acoustic results can be derived as well, namely the reverberation time (RT) and the speech intelligibility (STI). For the sake of simplicity both measures are implemented only in a limited manner, if further data is needed more dedicated processing tools like the standard EASERA software should be employed.

It is important to note that in real-world venues, neither the transfer function nor the impulse response are simple functions. They will usually show a lot of detail and they will of course depend on the location of the microphone (or any other receiver at the output of the system under test) and the loudspeakers (or any other sources at the input of the system under test).

Example

Before we start to look at the measurement of transfer function and impulse response in SysTune, we shall briefly repeat the very fundamentals of what these data are good for. For that purpose we will use one of the example files, FinalMP2New.etm, that is shipping with the program.

Such a typical impulse response measurement is shown below.

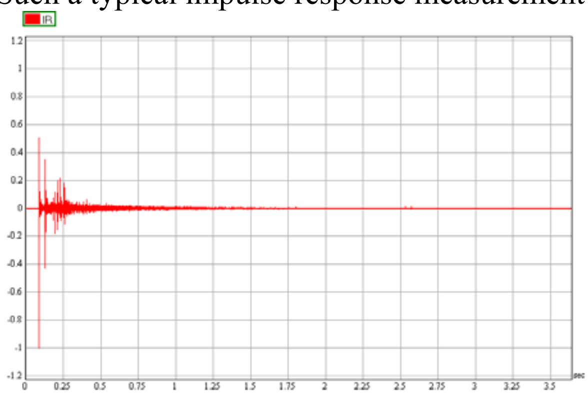

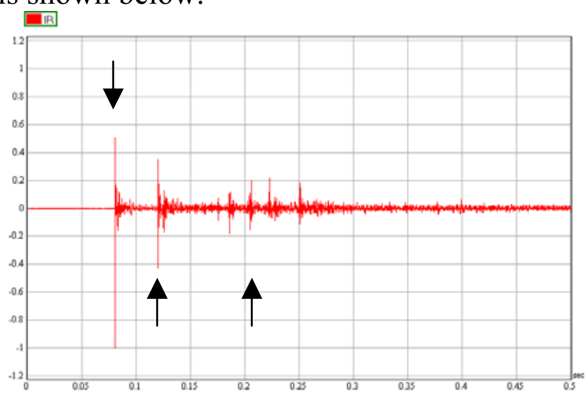

The very first peak in time is the direct sound, located at about 80 ms. Over the following 300 ms there are a number of peaks, identified as first-order reflections from the ceiling, the floor and some side walls. After that, the statistical decay takes place, which means the reflections become dense and cannot be separated anymore; in this regard one also speaks of the diffuse field. That part is very obvious in the ETC graph below.

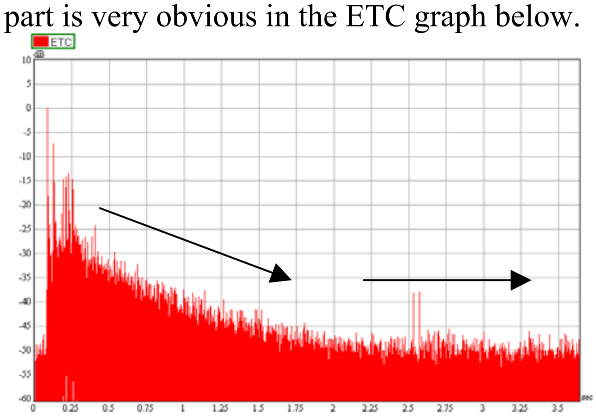

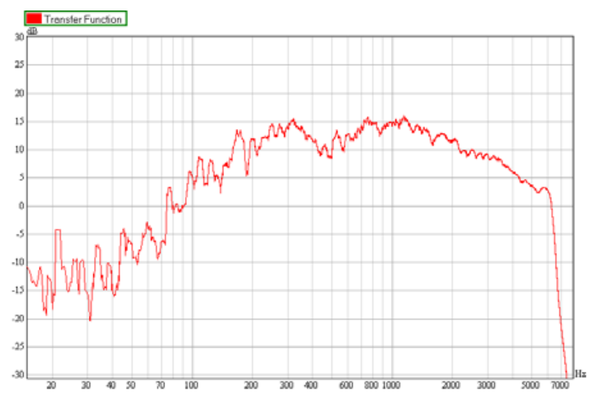

After about 2 seconds this ETC display also shows a constant noise floor which is typical for all room-acoustic measurements. This noise floor is related to the noise floor during the measurement in the actual venue. The higher the noise is, the less you will be able to precisely measure direct sound, reflections and the decay tail. In addition, the late part will clearly show any measuring errors (here two spikes at 2.5 seconds) and possibly distortion effects, because in the late section they are not covered by the main signal.

The picture to the right of the ETC shows the magnitude of the transfer function, smoothed to 1/12 octaves. Although this spectrum corresponds to the full impulse response, it is not necessarily connected to our listening experience. The reason for that is simple: for the impression of tonality and balance our ears mainly use the energies arriving early, but not the full length of the response. A better correlation to the subjective experience can often be achieved by only looking at the spectrum of the early part of the impulse response. The process of cutting out and analyzing this section separately is called windowing.

The magnitude function for only the direct sound (first peak) from the impulse response is shown below. You can immediately see a fairly wide gap with a level drop of about 10 dB located at around 500 Hz. It is quite probable that this dip can be recognized by the audience, so in practice one would think about counter-measures. Do not worry about the cut-off at the low end and at the high end right now. The drop below 100 Hz is caused by the window that we applied. Its size of ca. 30 ms naturally excludes frequencies with a wavelength longer than that. The roll-off at 7 kHz is caused by the measuring system that was used to gather the data originally, it was not SysTune in this case.

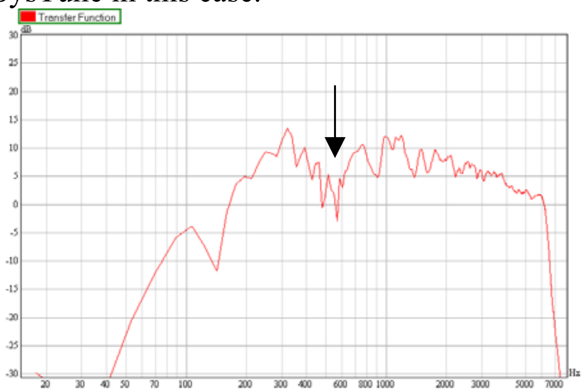

The phase function of a room-acoustic impulse response can be viewed in SysTune as well, but it normally has no meaning because there is no coherence in the overall sum of direct sound, reflections and diffuse energy. However, the phase response of the direct sound or the early arrivals is often useful to analyze the time alignment of loudspeaker components in the frequency domain. A typical graph for the windowed phase is shown above, to the right of the graph of the windowed magnitude. When a sound system consists of multiple sources covering different parts of the frequency spectrum, like a horn and a woofer, small differences in their arrival times at the measuring location will show up very clearly in the phase graph. In the plot given here one can roughly see a steady decay of phase up to 500 Hz as well as a rising curve starting at 2 kHz. This varying frequency behavior would be typical for a system with LF and HF components that are not exactly aligned.

Once a full impulse response is acquired, basic room-acoustic measures can be derived in SysTune, as well. The reverberation time (RT) of the room can be calculated for each octave band, for the given example the RT curve is shown in the left graph below. It is typical for mid- to large-size venues that the reverberation time decreases with higher frequencies due to air absorption. The speech intelligibility can be estimated as well, SysTune provides a graph for the speech transmission index (STI). The red curve is what we are interested in at the moment, it shows an STI of about 0.53 which is fairly good.

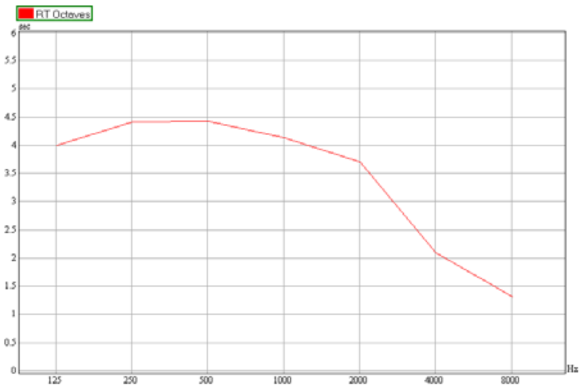

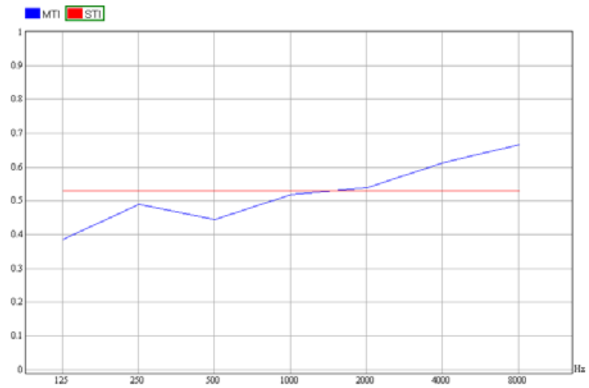

These two calculation options are implemented in SysTune, for more advanced results and detailed processing options please refer to another software, like EASERA or EASERA Pro.

Summary

This section gave you a quick overview of the theoretical background for transfer function measurements in SysTune. To determine the transfer function of a system under test, it is needed to know both the input signal into the system and the output signal as its response to the input. Therefore transfer function measurements generally require two channels, called the signal channel and the reference channel. The latter channel can be purely virtual if the excitation signal is known in advance. We emphasized that at present SysTune is the only software that can perform real-time measurements of full-length acoustic impulse responses. Finally, we discussed some examples for the use of acquired data in system tuning and room analysis.

4.2. Impulse Response

To exercise the measurement of a transfer function and impulse response, we will start with an internal stimulus and use that as a reference signal as well. In the previous chapters we have already covered how the input channel and the excitation signal are selected. We want to stay with that choice for the moment, so please select input `1` as the `SIGNAL CHANNEL`, pink noise of 3 seconds length at – 6dB as the stimulus signal and choose output `1` as the `PLAY CHANNEL`. The `FFT SIZE` should be about 3 seconds and the number of `AVERAGES` just `1`. After that, the control panel on the left will look approximately like this:

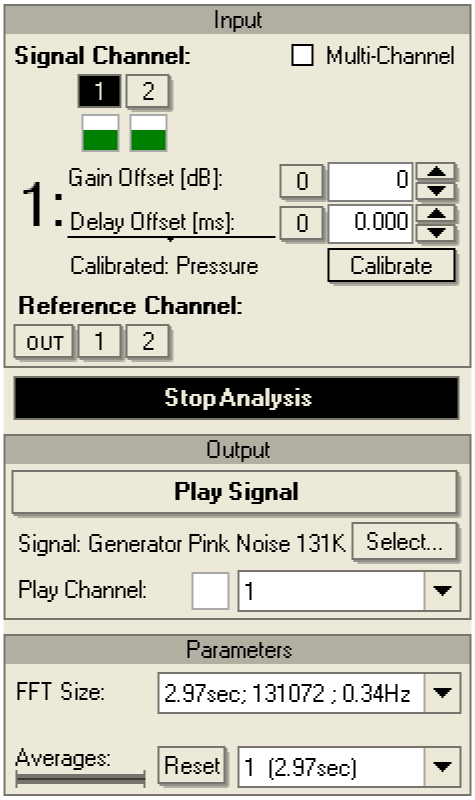

Disregard the fields for GAIN OFFSET and DELAY OFFSET for the moment, we will come back to them later. Below these items, the REFERENCE CHANNEL can be selected. Essentially, there are two choices available:

- The setting OUT tells the software to use the excitation signal as the reference, this is the virtual loopback we spoke about in the previous section.
- By selecting the input channel numbers 1 up to a maximum of 8 an external reference signal can be enabled. (The maximum number is determined by your sound hardware.)

Activate OUT to use the pink noise sent to the output as the reference signal as well. We do not need to connect another physical input for that.

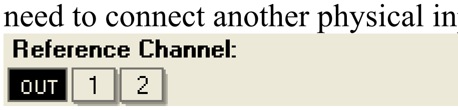

Immediately after setting the REFERENCE CHANNEL, the button rows at the top of each graph are unlocked completely. The buttons under TRANSFER FUNCTION and RESULTS are only available with a reference signal.

If you have not started the analysis yet, it is time to do so: press START ANALYSIS at the left. Also, we have to play a signal again, so ensure that your output gains are at a safe setting. Press PLAY SIGNAL to start the output. After that, switch the upper graph (we assume a split view here) to the impulse response: select IR under TRANSFER FUNCTION in the button row at the top. You should see a display like the ones below:

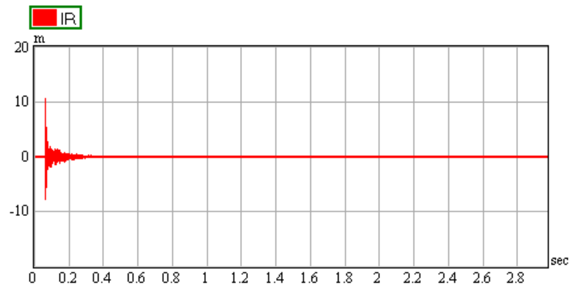

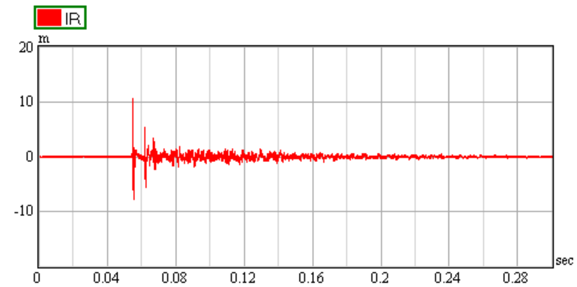

There should be a stable graph with one major peak that corresponds to the direct sound path from the loudspeaker to the microphone. If you are using multiple loudspeakers you may see several peaks. Also, depending on the room, you will see some reflections and the decay behavior we introduced in the previous section. You may want to zoom in to investigate the impulse response more closely like we have shown in the display above on the right.

Trouble-Shooting

If you do not see that or if you experience strange effects, here is a list of solutions for some frequently encountered problems:

- Make sure that your inputs and outputs are not clipping. Any clipping will lead to noise in the IR, severe clipping will destroy the IR completely. Check the mini-meters for input and output to be always in the yellow or preferably the green range.
- If you see a moving peak or a peak that jumps to a totally different point of time once in a while, make sure you are using a reliable soundcard. This effect is mostly caused by input and output not being clock-synchronized.
- If you do not see any peaks, make sure that the total delay of the output of the system is not longer than the `FFT Size` (here 3 seconds). In air, 3 seconds correspond to 1 km of path length, so this is normally not the problem. It is more probable that your measuring chain includes a delay setting somewhere.
- If you see a high noise floor, make sure that your system is sufficiently linear and time-invariant (see first section of this chapter).
- If the main peak appears at the end of the IR, the driver of the soundcard is possibly reporting a wrong internal latency to SysTune. Try to restart the analysis, to restart the software or reboot your computer to reset the driver. Make sure you have the latest driver update installed.
- Generally, a peak shifted toward the end or located at a time that does not correspond to a physical delay is often caused by the time alignment of the input compared to the output. Verify that your soundcard drivers, console or DSP settings do not use extensive delays.
- If you see several identical peaks at equal spacing, make sure that you have selected an `FFT Size` that is equal to or shorter than the time period of the signal played. Otherwise you will see such "phantom echoes".
- If you are not using pink noise but a periodic signal, make sure that the repeat time of your stimulus signal is not larger than the FFT size. Otherwise the FFT block will only contain a part of the signal and thus possibly only a part of its bandwidth.

If none of these points could resolve the problem, please refer to the trouble-shooting section in the next chapter for more help.

Impulse Response Measurements

Now that we have established the IR graph, we would like to examine the various options associated with it. By default, you should see on the right a DISPLAY panel configured as shown below.

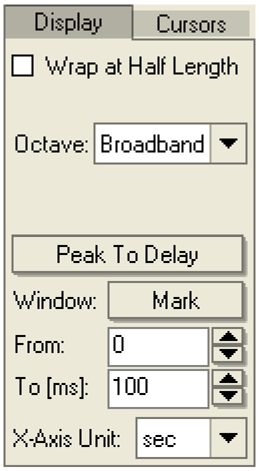

The first display option, WRAP AT HALF LENGTH, allows you to move the zero point of time to the horizontal center of the graph. We will see later that this is very useful when working with IR windows. But for room analysis, it is more common to have the starting point located at the very left of the graph. For right now, activate this check box and the IR graph will change to a form as shown below:

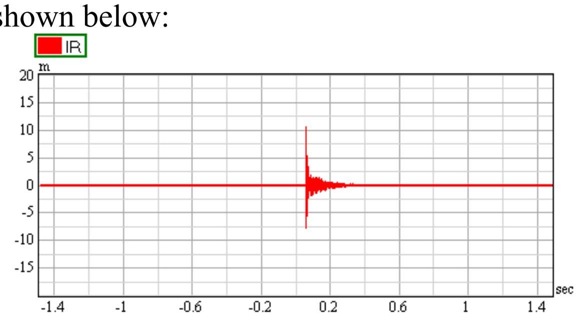

A bit earlier we emphasized the significance of considering the delay between input and output. In the IR graph of this measurement you can see that there is a remarkable delay indeed. Normally, it consists of the time of flight from the loudspeaker to the microphone and of the electronic latency in the signal chain.

In order to synchronize the signal channel with the reference channel we need to determine the delay and enter it into the software. In the control panel on the left you find a text field labeled DELAY OFFSET.

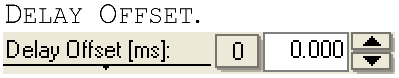

By default, it will show only zeros. Here the delay of the input channel should be entered in ms relative to the reference channel. To determine the exact delay of a channel up to the accuracy of a single sample is sometimes difficult, especially when the system consists of multiple signal paths with different frequency behavior. However, there are some simple functions in SysTune that will take you close enough.

In many cases, the delay of the system under test corresponds to the maximum value of the IR. In many other cases, the time of the IR's maximum value will be at least close to the actual delay. So, in any case this would be a good value to start with. Press the button PEAK TO DELAY on the right hand to automatically detect the absolute maximum in the IR and to transfer the associated time value to the DELAY OFFSET text field. This will move the peak to zero time within a few seconds.

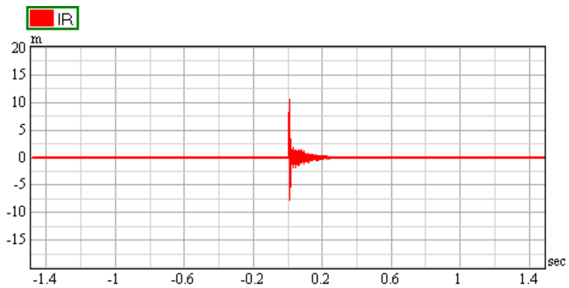

To try out different delay settings, you may enter a new value manually or use the spin buttons to adjust the delay time in sample increments. The button 0 will always reset the delay. We will come back to this topic when we look at the PHASE graph later on. For the moment, we leave the delay set to the time of the peak.

So far we have only looked at broadband impulse responses. In SysTune, you can also apply standard octave filters to an IR. To do that, select an octave center frequency from the OCTAVE drop down list. By default it is set to BROADBAND, but octave filters from 125 HZ to 8000 HZ are also available. At this point, it is noteworthy that these filters are only applied for display purposes. Like the setting WRAP AT HALF LENGTH they do not affect the underlying data or any other graph.

Hint: The real-time impulse response display can be used very easily to time-align full-bandwidth loudspeaker systems. The octave filtered IR graph can be utilized to align components of a system that work in the same frequency range if they are hidden or covered by other sources in the broadband IR.

To give a preview, SysTune supports an advanced filtering function as well. It is located on the TOOLS tab on the left and allows you to define custom band limits for the IR calculation. We will explain that later when we talk about reference signals from an external source.

The next items on the right side are related to adjusting the IR WINDOW. The button MARK activates a special mouse mode to place the window with the left and right mouse buttons. The fields FROM and TO can be used to enter the window coordinates directly and adjust them precisely by means of the associated spin buttons. We will be concerned with that shortly, when we talk about the graphs for windowed magnitude and phase.

The last element in the DISPLAY tab is a drop down list labeled X-AXIS UNIT. In addition to using seconds (SEC) for the horizontal axis of the IR, the length units M and FT can be employed as well. To calculate the scale for the horizontal axis based on distance, the speed of sound is assumed to be fixed at 340 m/s. Note that depending on the environmental conditions during the actual measurement, there may be slight deviations between displayed and real distance values.

Of course, as for all measurements before, the panel on the left controls the measurement parameters. Noteworthy here is that the signal-to-noise ratio for the impulse response can be increased significantly by choosing longer measuring periods. This can happen in two ways, either increase the FFT SIZE or increase the number of AVERAGES. For every doubling of the measuring time you will gain an increase of 3 dB in the signal-to-noise ratio, if the noise floor has a random nature. The obvious disadvantage of longer measuring durations is that the graphs will respond to changes in the measurement setup more slowly. For the purpose of time alignment, as explained in the section further below, you will have to find a sufficient compromise for these parameters, although mostly high S/N is only needed for room analysis and the calculation of RT and STI.

Tech-Note: Using Short FFT Sizes

In the introduction we emphasized how important it is that input and output channels are synchronized with regard to their delay. Just a little bit above we also applied the DELAY OFFSET practically. This is especially necessary, if you want to look at shorter FFT sizes. The following pictures should clarify this principle.

This is a measurement of a full-length IR with an FFT Size of 3 seconds. The arrival time is located at about 175 ms.

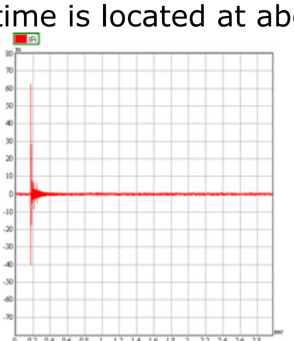

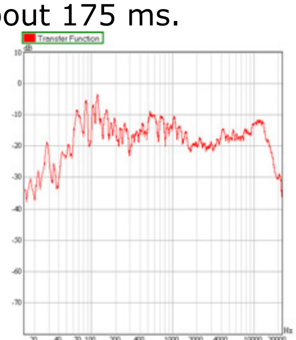

After entering this delay in SysTune, we can look at a smaller FFT size of perhaps 90 milli-seconds and obtain the accurate impulse response and transfer function for this shorter time section:

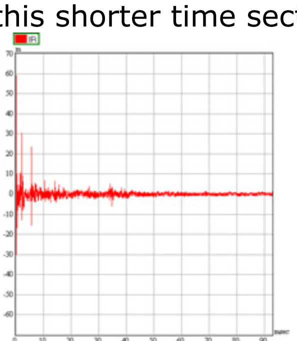

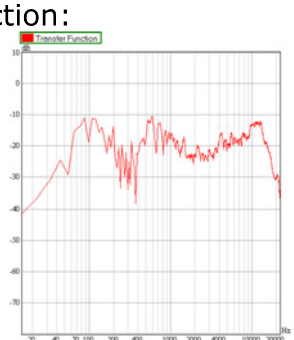

However, if we do not adjust the delay first we will receive the erroneous results shown below:

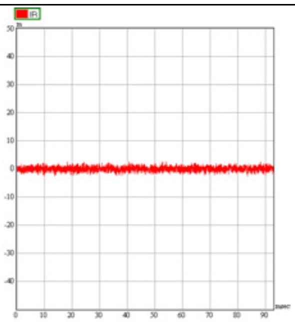

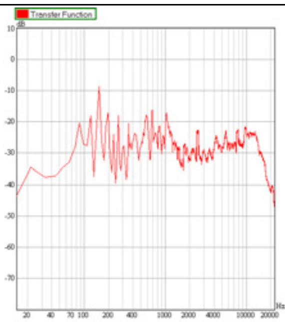

Now, the general level is about 10 dB lower and the are mainly showing noise. Obviously, without the ability to look at the actual IR it is difficult to estimate the validity of the transfer function obtained.

Time Alignment of Loudspeakers

Similar to the INPUT graphs, overlays can be captured for the TRANSFER FUNCTION graphs as well. In fact, a captured overlay always consists of all data available to the program at the time when it was stored. In the following short excursion we will use SysTune's ability to capture impulse responses to show how easy it is to time-align several loudspeakers to each other. The measuring setup that you are using to work through this document may not include multiple loudspeakers or delay control functions like we will need them. If so, only read this paragraph and try out the sequence of steps later.

Let us assume there are two loudspeaker systems which we want to align in time. First of all, they have to be connected to the output of our soundcard in order to make measurements. The impulse response with both loudspeakers active might look like this:

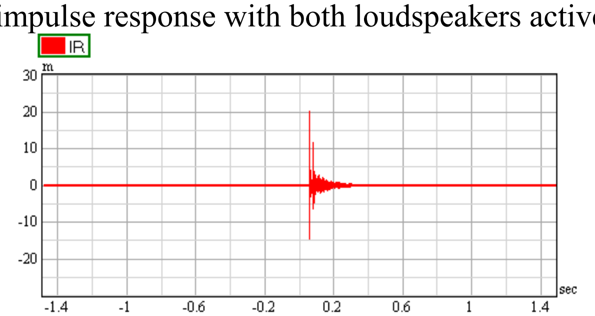

Assuming that we can only apply additional delay to the loudspeakers, we will have to delay the loudspeaker arriving first (A) compared to the one arriving later (B). Since we will be using the last arrival as our reference time, we want to measure the late loudspeaker B first, then the early loudspeaker A. So first, disconnect loudspeaker A and start the measurement. Then the impulse response should only show a single peak. Press PEAK TO DELAY to move the time arrival to zero and then press CAPTURE MEASUREMENT to capture the data and store it as an overlay.

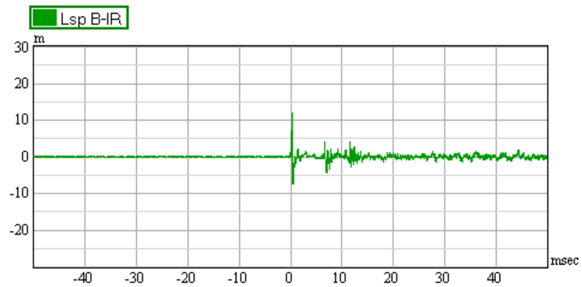

After that, disconnect loudspeaker B and connect loudspeaker A. Measure again and you will see a graph like the one below.

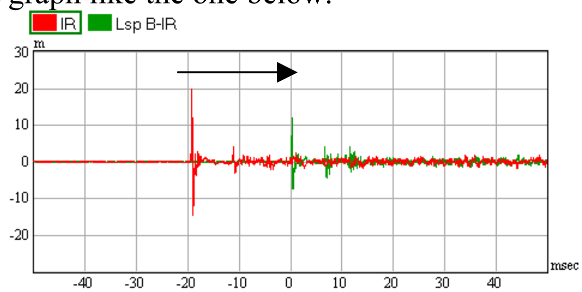

Because we adjusted the DELAY OFFSET for the last loudspeaker B, all other peaks in the sequence of arrivals must now occur at negative times. In the picture above the reference loudspeaker B is shown in green color, loudspeaker A in red.

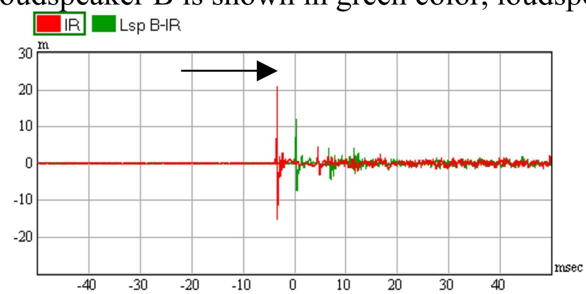

Once the reference time is set, it is very simple to align all other sources to it. You may read the time differential from the IR graph and enter it into the controller for the loudspeaker A. Or you simply adjust the loudspeaker delay slowly and use SysTune to watch the peak of loudspeaker A coming closer to the reference in real-time. A good time alignment is reached when the direct sound peaks of the loudspeakers A and B match sufficiently.

Hint: You can use the same method for aligning subwoofers. For that purpose, you have to enable an octave filter, like 125 Hz, to exclude any other sources from the IR that are radiating sound in the mid- and high-frequency range.

> Tech-Note:
> In reality, the time alignment of loudspeakers can be a time-consuming task. When the direct sound peaks are not similar it is often unclear how to arrange them for the best match. In such cases, a closer look at the phase graph may help. We will talk about that in just a few pages.

Also, it should be clear that in this example we only tuned the delays for a single receiving location. In practice, you will have to adjust loudspeaker delays in a way that the alignment is acceptable for the whole audience area covered by both loudspeakers. This is always a compromise, because exact time alignment can only be achieved for a single spot but not for a set of different locations.

Saving Impulse Responses

Once acquired, you may want to save the impulse response as an audio file. To do that go to the FILE menu and select SAVE TO AUDIO FILE | IMPULSE RESPONSE. This will bring up a dialog that allows you to select the CURRENT MEASUREMENT and any of the overlays for saving to a file.

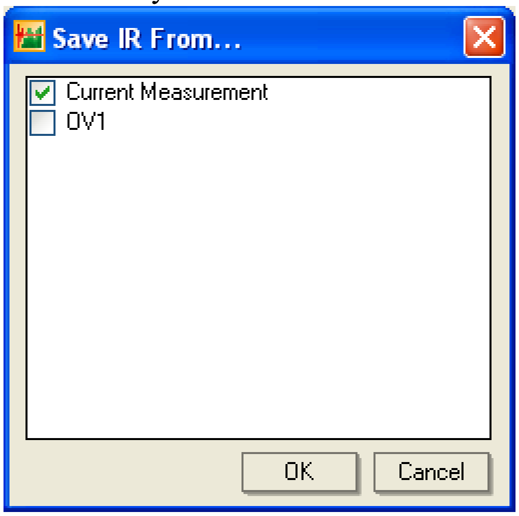

You may then use these files in a different software, like EASE or EASERA, or load them in SysTune at a later point of time.

Level Meters

In the last chapters we have worked quite a bit with OVERLAYS already. You may have wondered what the second tab is about that hides behind the OVERLAYS tab and is called LEVELS. This tab shows two large meters to allow better control of the levels at the inputs for signal and reference. With respect to their functionality, the two meters are similar to the mini-meters, but not identical with them. A picture is shown below; you may have to increase your overall window size a little bit to see it better. It displays our current setup with the reference being the pink noise signal at a gain of –6 dB.

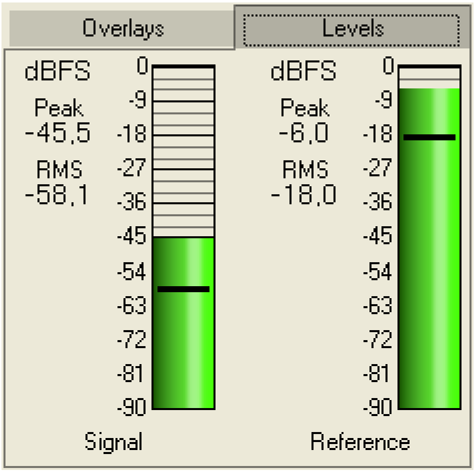

Remember that we selected a SIGNAL CHANNEL and a REFERENCE CHANNEL a short while ago. The level at these two inputs is shown here. Since the main purpose of these level meters is to check for clipping and to obtain an impression of the signal dynamics, they always display in digital units, that is full-scale (FS) numbers. In fact, two values are shown for each channel:

- The PEAK value as a number and as a vertical bar. The upper part of the bar will turn from green to yellow when a level of –6 dBFS is exceeded. Should the level reach –1 dBFS and more, the highest part of the bar will be filled with red color. By default, the hold time is 3 seconds. This setting can be adjusted in the OPTIONS window.
- The RMS value as a number and as a black horizontal line on the vertical bar. This is the RMS value over the last FFT block.

The next picture shows the level meters at clip level.

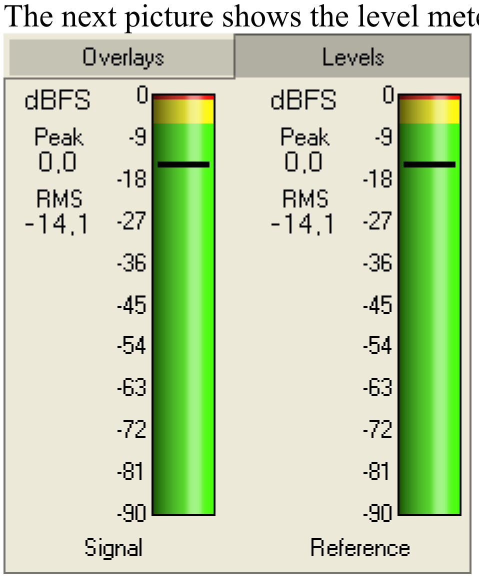

Hint: Make sure that you are always in the green range with both signals. During your first measurements it will occur more often than you may think that you find the reference channel clipping. This happens commonly because you cannot hear the reference signal when using an internal loopback or electronic input.

In addition to the Levels tab, there is another way to monitor the level of the reference channel. It is especially comfortable in a dual-channel setup. By enabling the Reference Channel, both Time Signal and Spectrum graph will show another check box in the Display panel. It is labeled Show Reference and it adds the reference channel as a second curve to the graph when activated.

Summary
To summarize, the acquisition and analysis of impulse response data was the main point of this section. We have described a simple measuring configuration using the stimulus signal directly as the reference. Advice was given to overcome potential problems in gathering the impulse response. After that, we demonstrated briefly how real-time impulse response measurements can be used to time-align loudspeakers. We also saved IR measurement data to a file. At the end we looked at the level meters in EASERA SysTune.

4.3. ETC

In SysTune, the energy-time-curve, or ETC, is defined as the logarithm of the squared impulse response, displayed in dB. Compared to the IR, its biggest advantage is the display of levels on a logarithmic scale. This corresponds better than the linear scale of the impulse response to what we hear and how we perceive sound over time. From the ETC graph engineers experienced in room-acoustics can immediately estimate the overall structure of the room, its reverberation time, early reflections and its suitability for music and speech performance. We have already discussed typical characteristics in the introductory section of this chapter.

Because the ETC graph is closely related to the IR graph, you will find the ETC button under Transfer Function, as well. A typical display like you may be seeing now is shown here:

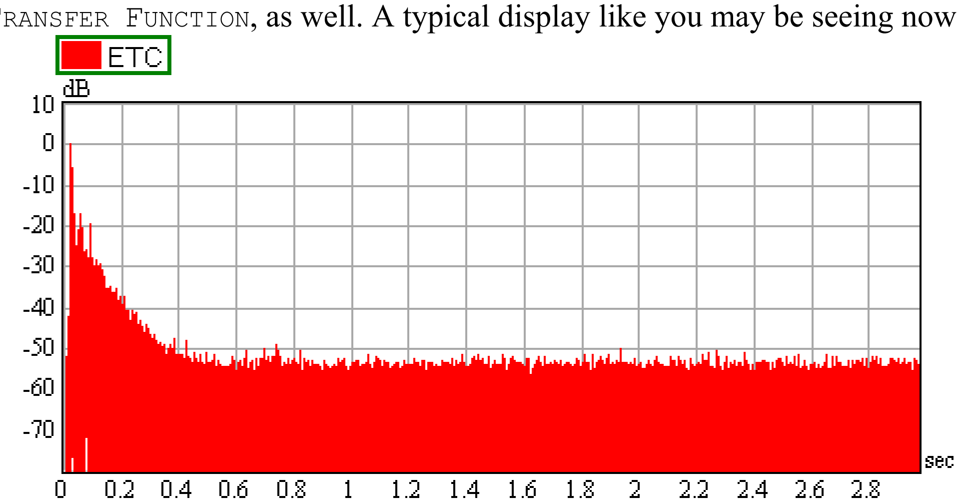

The Display panel to the right is very similar to the Display panel of the IR graph. The only difference is that the option Show Schroeder can be enabled here. When active, the Schroeder backward integral is added to the plot. This is particularly useful to analyze the reverberation time of a room more deeply and the decay of the energy in the volume in general. We will explain RT measurements and the associated parameters a little bit later.

The ETC graph is also a good indicator for the signal-to-noise ratio of the measurement. You only need to estimate the average level of the noise floor from the end of the curve (here ca. -50 dB). Since the peak of the ETC is always normalized to 0 dB, you have the S/N ratio immediately. Good acoustic measurements tend to have about 50 dB S/N. Significantly more than 60 dB is usually not possible with standard test equipment and reasonable measuring times.

Analyzing Reflections

As stated before the ETC graph is particularly useful to analyze reflections in detail. The CURSORS function is of great help in that. We already introduced it when we spoke about the INPUT SPECTRUM.

For a better overview please switch to the maximized graph now. For that, remember to use the small buttons in the menu bar at the top and very right. Then, zoom in to the first 300-400 ms which is normally the period of time where strong discrete reflections are visible.

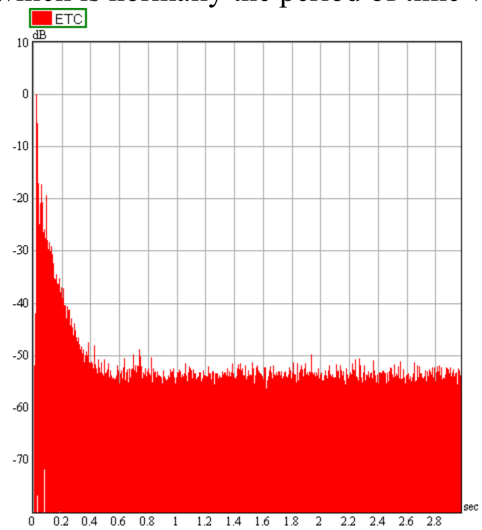

Now change the mouse mode to PEEK (left-most of the three tool buttons in the mouse bar) and use the mouse to insert 3 or 4 fixed cursors: move the mouse to place the tracking cross at the direct sound or a prominent reflection and make a right click to add a cursor at the location of the cross. After that, the window should look similar to the picture below:

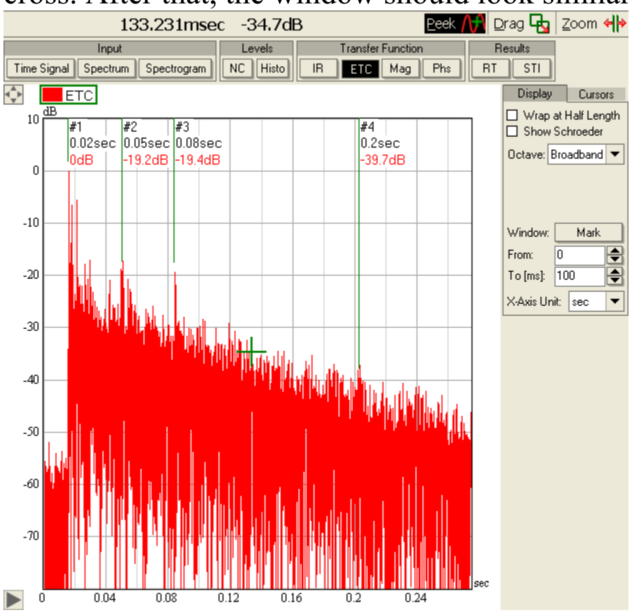

We are interested in the level and time difference of the reflections compared to the direct sound. Because it is very useful in this context, let us also briefly repeat how to display differential values using the inserted cursors. To select a reference cursor open the CURSORS tab to the right (hidden by the DISPLAY tab) and left-click on the R button in the first row.

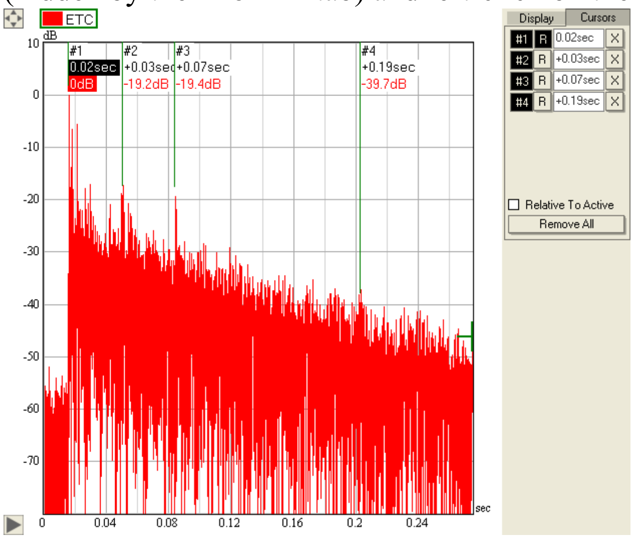

After that, the cursors in the ETC graph are shown with differential values, for both time and level. Check at which times the selected reflections arrive and which level they still have compared to the direct sound. The later they come and the louder they are the more probable it is to perceive them as an echo.

Summary

In SysTune, the log-squared impulse response is called energy time curve or ETC. The display options available for the ETC are similar to the IR. We have shown that the fixed cursors can be easily used to analyze the ETC with respect to its structure in time and level.

4.4. Magnitude

So far, we have discussed the time domain results of a dual-channel measurement, now we would like to look at the frequency domain, namely, the `TRANSFER FUNCTION`. To the right of the buttons for `IR` and `ETC` you will find the buttons for `MAG` (magnitude) and `PHS` (phase) as well.

We have already explained that the transfer function of an electronic or acoustic system is usually a complex function, which consists of a real part and an imaginary part. Equivalently, and practically much more useful, the complex transfer function can be separated into magnitude and phase. We should emphasize that these two entities - although often related due to the underlying physics - are independent and cannot be derived from each other. Magnitude and phase data individually represent only a half of the whole transfer function.

Relationship between Transfer Function and Impulse Response

The complex transfer function as defined in the introduction of this chapter can be derived from the impulse response by means of a Fourier transform. Vice versa, the impulse response (IR) can be directly calculated from the complex transfer function (TF). It must be underlined that these two functions are equivalent representatives of the same thing, the response of the system under test to an ideal impulse. The conversion between IR and TF is a lossless, purely technical process. Therefore, measuring one of them always gives you the other one too.

This relationship is sometimes confused in commercial documentation, where due to historical reasons of available computing power or memory the transform between time and frequency domain is understood as a lossy process. But loss of information happens only when a cut, a window or a filter is applied to the data and the actual remainder is transformed to the other domain. In EASERA SysTune, time domain and frequency domain views are principally equivalent views. They only show the response of the system in different ways. In SysTune, like in the real world, any change of the system in the frequency domain will be visible in the time domain as well and vice versa.

In the same manner, the resolution and size of IR and TF are controlled by the same measuring parameters, namely sample rate and FFT size. The length of impulse response is defined by the length of the FFT block, and the resolution of the IR by the sample rate. The size of the transfer function, that is the upper frequency limit, is determined by the sample rate and the resolution of the TF by the FFT size. So, changing either parameter will affect the resulting data in both domains.

Magnitude of the Transfer Function

The magnitude display of the transfer function of the system under test is probably one of the most important views of any acoustic measuring platform. If available in real-time, as in SysTune, changes of the system during the measurement can be monitored continuously. In practice and especially with respect to loudspeaker systems this has two important consequences. On the one hand, the system can be watched passively and its performance verified before and during an event. On the other hand, the system can be configured and tuned actively, all changes to the setup become visible immediately while the system is operating.

The magnitude display of the transfer function (or frequency response) fulfills a variety of purposes in acoustic measurements:

- It is useful when equalizing a single loudspeaker or a whole sound system. This often includes multiple measurements for different spots, because like time-alignment, frequency equalization in a room can also only be perfect for a single spot.
- It is useful to detect comb filtering effects which distort the listening impression.
- It is useful in the low frequency range to identify and measure room modes.
- It is useful to determine the overall gain of the system under test.

In most cases, the full frequency response of a room will look similar to the curve shown in the introduction.

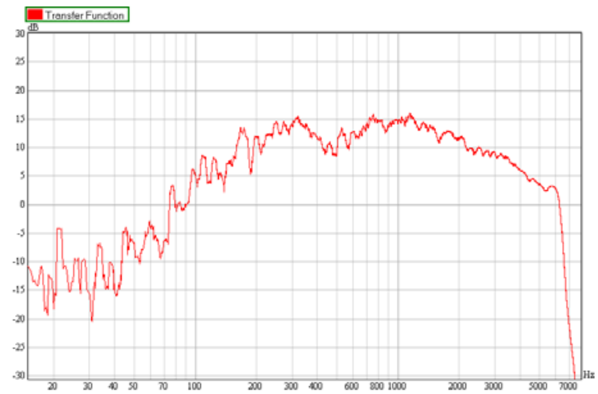

Toward the very low frequencies the level drops, because a normal room with loudspeakers is not a very efficient transmitter of sound energy around 20 Hz and below. Toward the high frequencies the level decreases as well, on the one side because the power output capabilities of the sound reinforcement system are usually lower at the very high frequencies, and on the other side because in medium- and large-size rooms the energy absorption of the air attenuates the propagating sound more at higher frequencies. How flat or volatile the frequency response is in the intermediate range depends on the sound system and on the room.

Remark:
Often a flat frequency response with minimal variation in the audible frequency range is considered optimal. However, it should be said that as outlined in the introduction the full frequency response of the room does not necessarily coincide with our listening impression. The frequency characteristics of only the early part of the impulse response, including the direct sound and the first major reflections seem to be equally important. Already, you may imagine that these two criteria are difficult to satisfy at the same time and for all listening locations. In addition, time arrivals and directions of the first reflections are important. Room modes will influence the impression at the low frequencies. The sensitivity of the human hearing system as another point of significance was discussed earlier. Various other factors enter into the equation as well.

It should be clearly stated that for a good listening impression many more things play a role than just the flatness of the frequency response. And beyond that, not only are the properties of the room and audio system important, but also the contents of the signal to be used. Speech and music performances have different requirements and different optimal conditions. Different kinds of music, like rock, jazz or classical music, do as well. Last but not least, the subjective *impression must not be underestimated. Objective measurements of a venue may yield close-to-perfect results, but still the audience may not like the sound. Also, a room may sound awful to one person, but someone else may like it.*

As a conclusion, we cannot give a general guideline here for how to create a perfect-sounding loudspeaker system or room. We can only describe the tools to analyze a sound system and to understand it and we can explain how to solve basic problems. But the question of the optimal configuration is to a large extent a subjective issue.

Now, switch to the magnitude display using the Mag button under Transfer Function. For the next steps it is most useful to have the upper graph display the impulse response and the lower graph the magnitude function. The picture below shows the approximate configuration, note that we work with the same settings as for the IR measurements before. For the IR we have used Wrap at Half Length and Peak To Delay to center the peak in the graph.

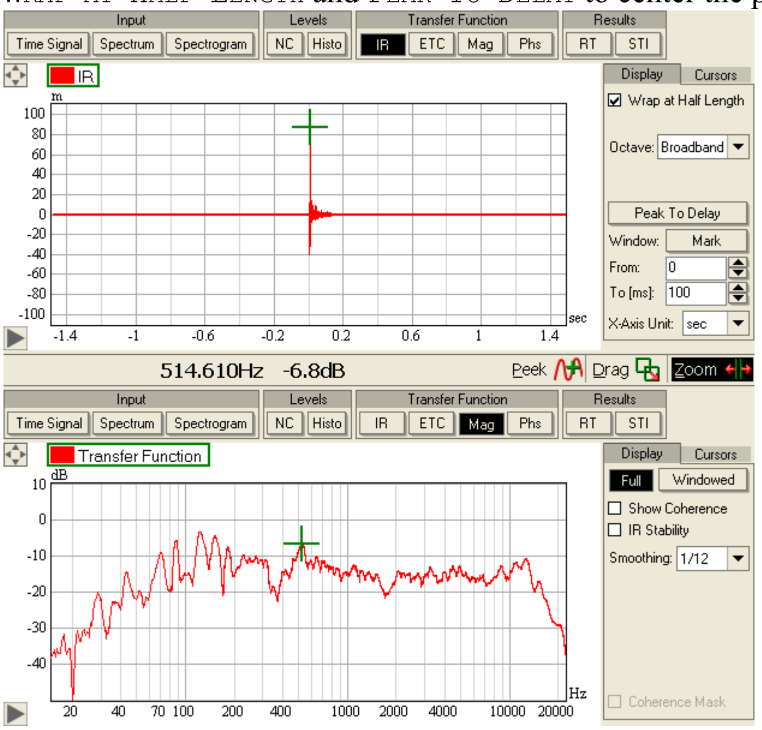

Let us look at the Display panel to the right of the Mag graph. The first row shows a check box to toggle between the magnitude display for the Full transfer function and for the Windowed transfer function. We will return to windowed measurements in just a few paragraphs.

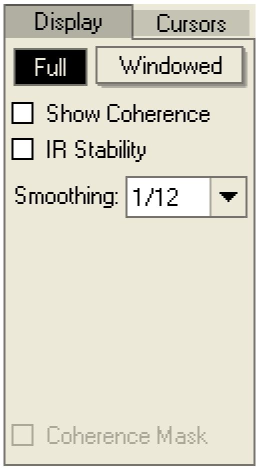

The next display options are labeled SHOW COHERENCE and IR STABILITY. These two check boxes are related, each of them adds a curve to the graph that describes the validity of the measured data. Before we explain these functions in detail, we want to mention the last control here, labeled SMOOTHING. This drop down list provides a number of bandwidths in fractional octave steps, from `1/1` down to `1/96`. In SysTune, smoothing just means that each data point is calculated as an average value over the chosen bandwidth. This helps to remove undesired details and noise from a plot, but sometimes it can also hide problems, when working with a resolution that is too coarse.

Coherence and IR Stability

Back to the options SHOW COHERENCE and IR STABILITY. Activating either check box draws an additional curve in the graph to indicate for which frequencies the measured data is valid and for which not. To understand this precisely, we need to come back to the introduction at the beginning of this chapter. There we defined that the transfer function in SysTune can only be measured for a linear, time-invariant system. We also stated that all real-world systems are nonlinear and time-variant, to a small degree at least. So the question arises, how do we determine if our transfer function measurement is correct and what kind of error is related to it?

The functions called coherence and IR stability provide a limited answer. They allow us to estimate how valid individual data points of the transfer function are. Both of them use the data from the signal and reference inputs to calculate the degree of linearity and time-invariance of the system. Ideally, for each frequency point the ratio between system input and system output is constant for all times and for any magnitude of the input signal. The basic idea of the coherence and IR stability functions is to quantify how much the actual measurements deviate from this ideal situation.

The coherency function is a well-defined entity in system theory, for the mathematical expression we refer you to standard text books. Essentially, it measures the correlation between signal and reference, output and input, respectively, of the system under test. This results in a frequency-dependent function with values between 100% and 0%, which represents the level of linearity and constancy of the measured system. There are some fundamental properties worth mentioning with respect to the software implementation.

- First, its evaluation relies on the set of available FFT blocks. Coherency is always 100% if there is only a single FFT block included with the measurement, which is 1 average. In

SysTune for this situation, you may only see values different from 100% if the reference signal is not defined for a particular frequency. For these frequencies the coherency will be 0%.

- Second, coherency accumulates variation and occurrences of nonlinearity over time. That means, the more averages you include in the measurement the lower will be the coherency, as every additional second of measuring time will always pick up a little bit of noise.
- In practice, a coherency value of 100% or close to that will be difficult to achieve. For a small number of averages, normally a coherency of 50% is a good value, for a large number of averages about 25% will usually be sufficient.

The IR stability function has been developed by SDA to complement the coherency function. It also measures the correlation between input and output of the system under test. But in contrast to coherency it does not relate signal and reference directly to each other, rather it is a differential quantity that describes the similarity of subsequently measured impulse responses or transfer functions. Here the value is also given as a percentage between 0% and 100%. IR stability is available for any number of averages and because it continuously compares only the last two measurements, it indicates temporal variations in the data quickly.

To summarize, coherency is a standard measure to quantify data validity. It is especially useful to indicate band limitations in the excitation signal or transmission chain, effects of noise and changes in the system under test. IR stability was developed by the SysTune software team to indicate fast changes in time. For both functions, a value close to 0% represents invalid data points; a value close to 100% confirms data validity. However, these tools are not the Holy Grail. They can not show every problem in the measuring chain and sometimes they may indicate errors which are practically irrelevant.

The following picture shows our current measuring setup with the stimulus muted. Note that we have temporarily switched to a number of 2 AVERAGES in order to show meaningful coherency values.

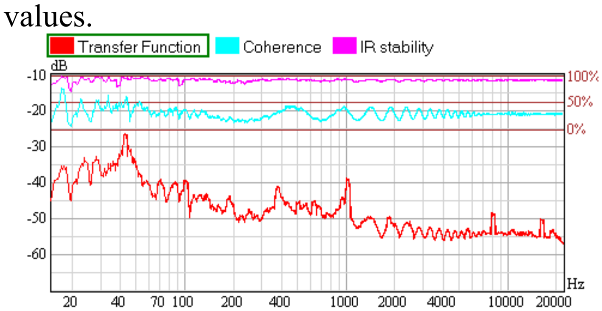

By default, both coherence and IR stability are displayed in the top quarter of the magnitude graph. If you would like to work at a higher or lower vertical resolution, you may adjust the area used under OPTIONS | COHERENCY GRAPH COVERAGE.

When SHOW COHERENCE is enabled, the switch COHERENCE MASK at the bottom of the DISPLAY panel can be activated as well. It allows you to partially or fully hide those parts of the displayed magnitude curve that belong to frequencies with a coherency value that is too low. By default, magnitude data points with a coherency below the threshold have a transparency of 74%. The default coherency threshold is 50%. Both parameters can be changed under Options.

Hint: This function is very useful when looking at transfer function measurements with an external stimulus. Music and speech signals especially are not broadband and they often change their spectral composition over time. By hiding invalid frequency points from the magnitude curve you can focus on valid data only.

Adjusting the Gain Offset

A program feature that has great importance when capturing and comparing magnitude curves is the Gain Offset parameter in the control panel to the left.

The gain setting can be defined for each input channel and each overlay individually. It controls the relative level of the measurement data for all graphs. In practice its functionality is two-fold:

- The gain parameter can be used to compensate for changes in the external gain structure after the system was calibrated. For example, if you have calibrated the microphone already, but later you decide to decrease the overall sensitivity a bit by reducing the preamp gain. The reason might be that the microphone is now closer to a loudspeaker and your input is clipping. In such a case you can just enter the differential gain into SysTune and re-establish the calibrated state of the software.
- Assume you make several measurements in a venue and you would like to compare the magnitude curve of the transfer functions in a joint plot of overlays. But because you measured at locations from the front to the back of the room, the curves are spread over a level range of 10 dB or more. By entering a gain value only for display purposes you can move the curves to a common range of levels.

For each input channel the Gain Offset can be entered directly as a numerical value into the text field. You can use the spin buttons to adjust it precisely in steps of 0.1 dB if needed. To reset the gain press the 0 button.

The equivalent control for a measurement that has already been stored is the Gain text field in the Overlay Properties window. Press the ... button in the Overlays list at the bottom of the control panel to open it. Feel free to go back to the third chapter to repeat the details of using this window.

Before we proceed with a detailed application of this function, let us practice it quickly to become familiar with it. You should still have the IR graph displayed at the top and the Mag graph at the bottom of the window. Make sure you are playing pink noise at –6 dB gain with an FFT Size of about 3 seconds and just 1 Average. Verify that your Reference Channel is set to Out. The Mag graph should be showing the Full transfer function with Show Coherence and IR Stability disabled. The Smoothing bandwidth is set to 1/12 octave. Check also if the Gain Offset is still 0 dB.

Now, we would like to capture a measurement and call it “Gain Offset Zero”. After that, enter a Gain Offset of 10 dB into the text field for the current input channel. You should see a picture like the following.

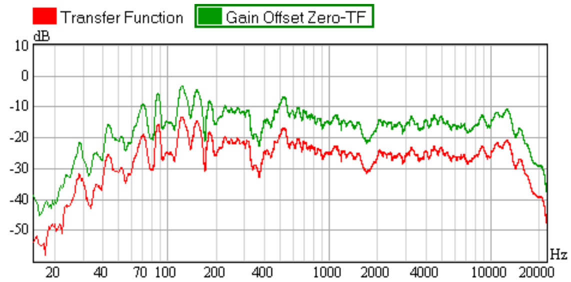

The program attenuates the current input channel by 10 dB. Because we did not add any gain externally, the graph will show exactly that reduction in the measurement. Just to practice, adjust the captured overlay "Gain Offset Zero" to match the current measurement again: Press the ... button in the OVERLAYS list and enter -10 dB into the GAIN text field. Press OK to confirm and close the window. As a result, the captured curve (green) and the current measurement curve (red) will be at the same level again. If you like, play a little bit with the gain functions. But before moving on with the next section clear the list of overlays using the REMOVE button.

Averaging Transfer Functions

For the majority of applications where you are tuning sound systems, you will not be measuring at a single spot only. You may take the microphone to several locations and record a transfer function for each of them or use multiple microphones to measure a set of transfer functions simultaneously. A representative frequency response for the whole audience area or for a part of it is often useful to make some judgments about the need for equalization. Like earlier for the INPUT SPECTRUM we will use the AVERAGE function to calculate that transfer function but we will also use the GAIN setting to adjust the individual measurements before actually taking the average of them.

Make three measurements for different microphone locations and label them "Meas 1", "Meas 2" and "Meas 3". Then stop the analysis and hide the CURRENT MEASUREMENT. The resulting graph will show 3 curves, each with a different frequency dependence. An example is shown below, it uses a SMOOTHING of 1/3 octave.

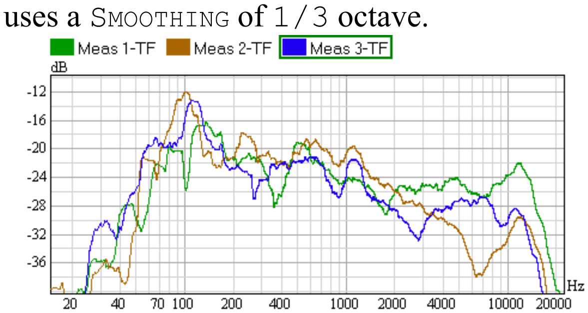

We know that curves with higher level will receive more weight in the average than curves with lower level. In the end, this is what averaging is about. But still, we would like to use the averaging process to extract the characteristics that are common to all of the curves. For that to work we have to make sure that they enter equivalently into the average.

In the example picture given, we see some correlation between 400 Hz and 3 kHz. Based on that, our intention is to remove specific variations and end up with a curve which represents all included measurements in their fundamental properties. In the given example, we would first change the SMOOTHING to 1/1 octave in order to be able to estimate the gain needed for each curve.

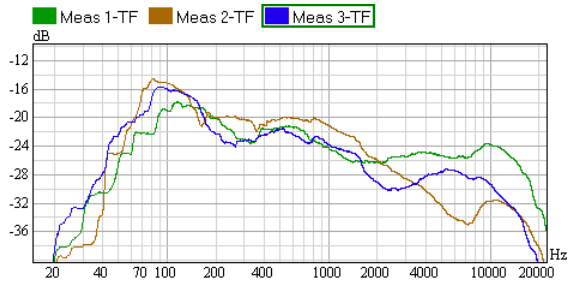

In the frequency range of interest, we find that we have to use a gain of about 1 dB for the 3rd measurement and a gain of about -1.5 dB for the 2nd measurement in order to align them approximately to the 1st measurement.

Hint: Such correction values you can just read from the graph, but you can also use the Peek mouse mode, the tracking cross or the fixed cursors as well. Remember that all of these functions allow you to determine absolute or relative magnitude values for the given graph.

Having determined the offsets, we enter these dB values into the GAIN field of the corresponding OVERLAY PROPERTIES window and switch back to a SMOOTHING of 1/3.

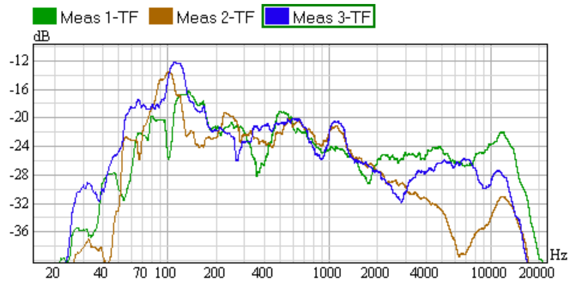

The overlays will be adjusted according to their gain setting and we can finally calculate the average of them. Press AVERAGE to open a window to select the overlays to include.

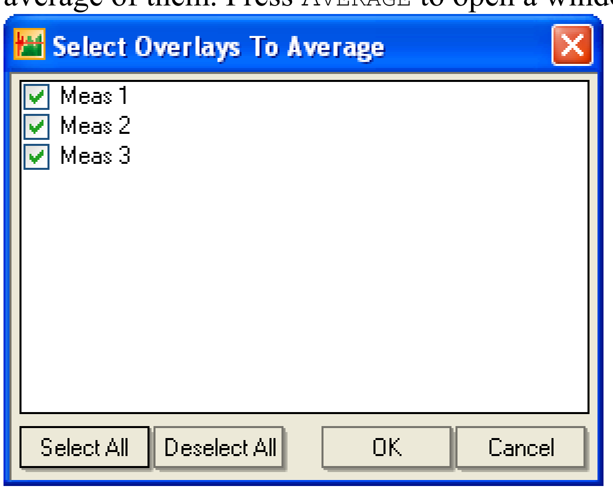

Activate all of the overlays for averaging and confirm with OK. A new data set will be added to the list of overlays. Label it “Avg 1-3”, pick dark red for the color and set all other curves invisible for the moment (switch visibility with the second button V). The resulting curve for our example looks like this:

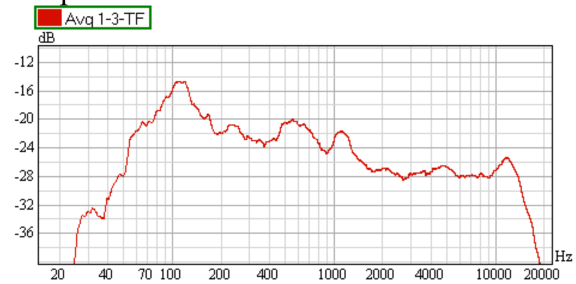

It is obviously smoother than any of the individual curves, which is a consequence of the averaging process, of course. But also, because we adjusted the gains properly, it is a data set representative for all of the included overlays; it is not similar only to the ones with the highest level. Therefore this result can be considered a starting point for equalization.

Tech-Note:
As explained earlier, when we introduced the input spectrum, in SysTune an averaged curve is the power average of the included measurements. It does not take into account phase or time alignments.

Because of that, we do not calculate an average impulse response or an average phase. It is worth mentioning, that the corresponding access buttons are disabled for averaged curves. In overlay plots with other data sets these curves will not be displayed.

The choice of gains for the individual curves is not trivial. It depends on the bandwidth that you want to look at. You may even use gains to emphasize certain measurements and lessen others in their weight for the average curve. A location centered in the audience area may have higher importance than seats close to the side wall.

Exporting Graph Data

To make use of the graph data in an external application we can export it into a text file. To do that go to the menu FILE|EXPORT DATA AS TEXT|FROM LOWER GRAPH when you are still in the split view with two graphs. A file dialog will open, that allows you to save the contents of the graph to a text file. Note that this works for any of the graphs and includes all of the overlaid data currently displayed.

However, using this command the data will be saved in the native resolution of that graph. For a transfer function, as we have derived it, it is often more practical to export it as fractional octave data and possibly invert it for later application in an equalizer. The menu command FILE|EXPORT DATA AS TEXT|FREQUENCY RESPONSE does just that. It opens the following window to select further export options:

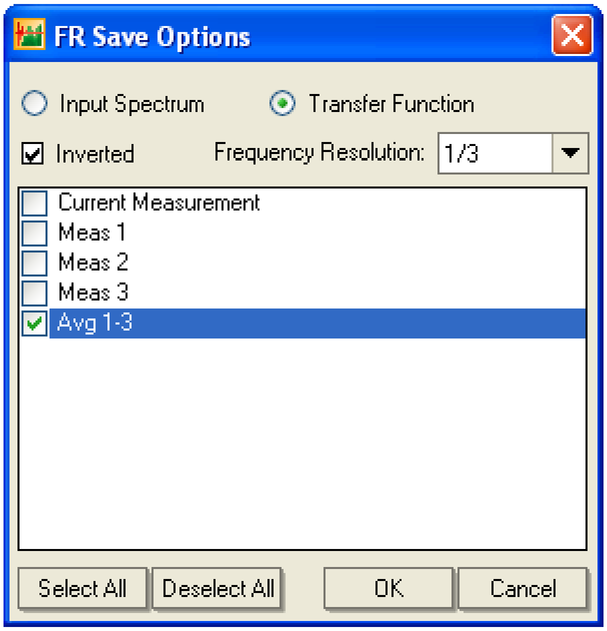

Here in our example, we would like to export the data of the TRANSFER FUNCTION, not of the INPUT SPECTRUM. We would also like to activate the switch INVERTED, it is on by default. The FREQUENCY RESOLUTION can be one of various fractional octave selections, at this time we want to export 1/3 octave data. Because we need only the data for the average curve "Avg 1-3", we activate that entry in the list below and deactivate all others. Press OK to confirm these options and then select a save location for the text file.

Introduction to Windowing

An important function that we have not discussed in detail yet is windowing. So far it has been mentioned only briefly, but in this and the following section we will explain the background for windows as well as their application in SysTune.

In the specific sense of audio measurement platforms we understand windowing as a process in the time domain that is used to isolate part of the impulse response for further analysis and at the same time exclude the rest of the data. A window is basically a weighting function in the time domain with values between 0 and 1. It consists of an area where it is 0; this is the part that is removed from the original data, and an area where it is greater than 0, which remains in some form. Often windows include a range where they are exactly 1; there the original data is left unchanged. Along with the definition of how the window is exactly curved its starting point relative to the data to be windowed must be defined also. Applying a window to a data set means to simply multiply each value of the original data with the value of the window function. Of course, this must happen with consideration of the starting point of the window. The result of this process is the windowed data set.

A commonly used window is the rectangular window. This is a function which has a value of 1 for a defined period of time and a value of 0 otherwise. Only the data of the section where the window is 1 is let through, the rest of the data, namely where the window is 0, is discarded. However, because the rise and the decay of the rectangular window is very steep - better said, there is an abrupt change from 0 to 1 or 1 to 0 - it may cause some undesired artifacts when processing the resulting data later on. As a consequence, a variety of windows has been developed to overcome such flaws and to instead provide a smooth transition. A common choice is the Tukey or flat-top window. It consists essentially of a part where the weighting function is 1 and a part where it is zero with a gentle roll-off in between which is based on a cosine function.

In the next picture the blue curve between the two vertical lines shows a typical Tukey window. In this example, the window is constant over 50% of the window length.

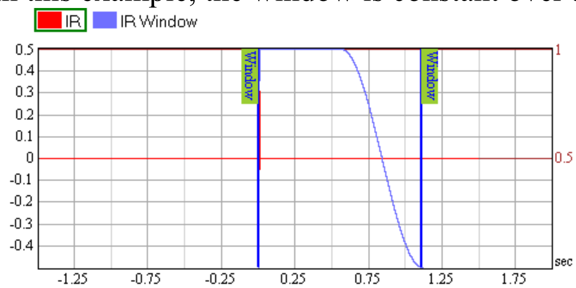

For a complete overview see for example: Fredric J. Harris, On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform, Proceedings of the IEEE, Vol. 66, No. 1, January 1978.

In general, windowing is a complex topic. There are many aspects to it and a variety of textbooks are concerned with that. We would like you to refer to such literature if you would like to learn all about the details of windowing. To summarize, the choice of the window mostly depends on the specific application and the data to be windowed. SysTune offers a small set of windows too, but generally we recommend using the Tukey windows. There is only one exception; in typical audio applications the new TFC window can be more effective. This window function was developed by SDA to resolve some of the shortcomings of general-purpose windows which are specifically obvious in acoustic measurements. We will explain this window type after we have become familiar with the windowing process in SysTune.

Windowing in SysTune

To show how to apply windows in SysTune, let us come back to the measurement setup that we last used. We had the IR shown in the top graph and the magnitude shown in the bottom graph.

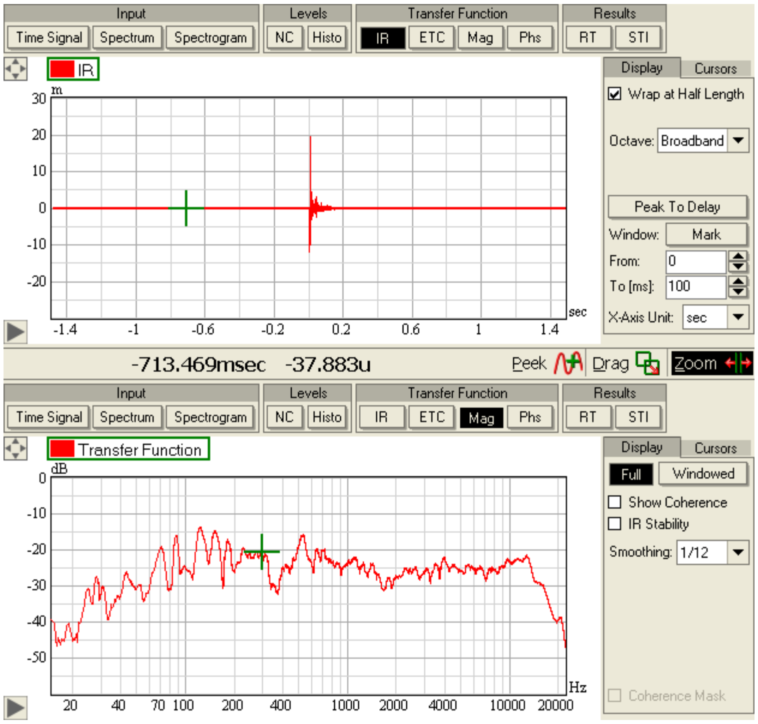

Assume, we have a clearly distinguished time arrival of the direct sound; this is normally the first and highest peak in the IR. Now we would like to look at the transfer function of the system only with regard to the direct sound. So far, the magnitude graph showed the Full transfer function, but we want to see it without any reflections. To accomplish that, we have to do two things:

- Define the window location in the time domain, that is in the IR,
- Switch the magnitude graph to the Windowed display.

For the first step, look at the Display panel of the IR graph. It shows several controls related to windowing. The button Mark activates a special mouse mode that lets you place the window start and stop coordinates with the left and right mouse button respectively. Click on Mark and then left-click in the IR graph on a location to the left of the peak, right-click on a location to the right of the peak. You should see a display like below.

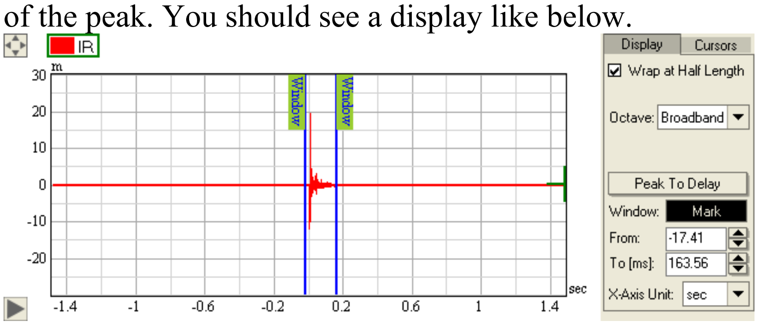

The start and the stop of the window are indicated by two blue vertical lines, we will call them window markers. The left marker always represents the start and the right marker the end of the window. You will also see the exact coordinates of both markers in the From and To text fields directly below the Mark button.

To define the window coordinates more precisely, let us zoom in a little bit. To use the mouse zoom you would have to deactivate MARK first, but we can do it more easily using the axis settings. Click on the small triangle in the left bottom of the graph to open the control area and enter "–20m" for the start of the X view and "200m" for the end. (These values assume that your IR peak was already centered at a time of 0 using the PEAK TO DELAY button.)

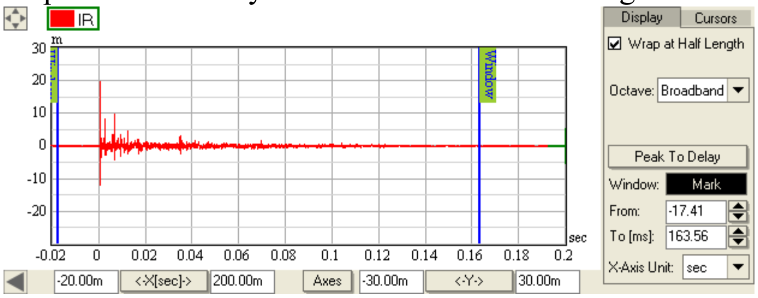

Again, use the left and right mouse button, but this time left-click on a location just a little bit left to the main peak and right-click at a time of about 150 ms. Alternatively you can enter "–2" into the FROM text field and "150" into the TO field.

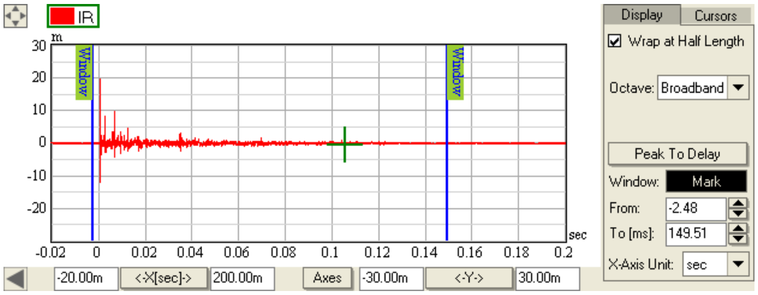

Note that the window markers are always defined in ms, in contrast to the time axis where we have entered the range in seconds or added the suffix "m" for milliseconds to the number.

In most real-world applications, you will already see a set of reflections within the first 150 ms and not just the direct sound. But for now, let us use this window length and adjust it later on exactly.

Having selected the window location, we switch the magnitude graph to WINDOWED. Actually, we could have changed that before defining the window, since the order of actions does not matter in SysTune. Both graphs are always synchronized. After switching, the magnitude graph will show you a smoother response than before. Because the windowed data includes less of the environment, the direct sound will dominate the resulting transfer function.

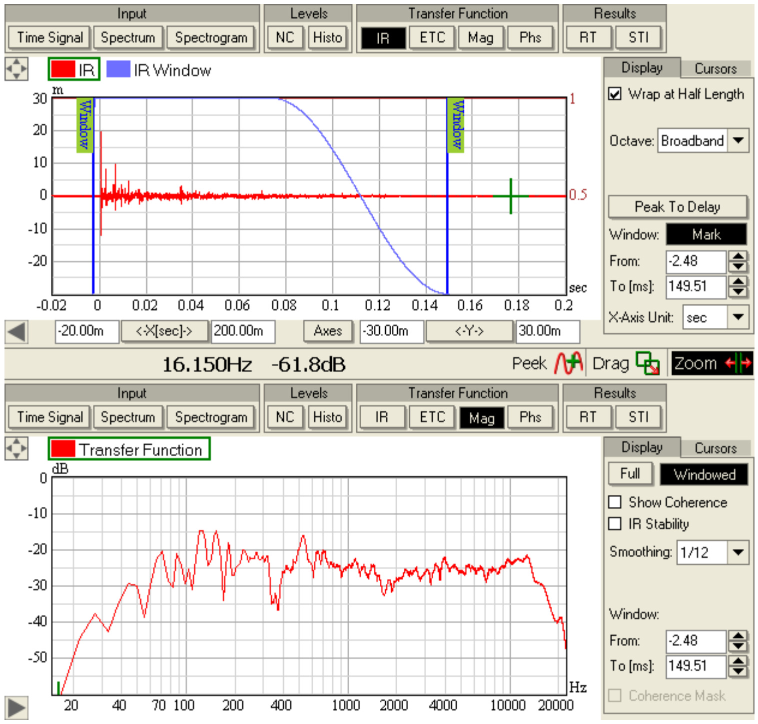

By activating the WINDOWED view in the MAG graph, immediately a few things have changed with it:

- When windowed data is displayed, the IR graph will show the window as a blue curve in the overlay. You can see the typical shape for a Tukey window, there is a flat area over the first 50% of the time and then a smooth decay toward the end of the window. This ensures that any data in the first half of the window area enters at full magnitude into the computation; later data points will be attenuated. Any data points before the window starts or after the window ends will be discarded.
- In the DISPLAY panel of the magnitude graph the same FROM and TO controls as in the IR graph are now accessible. This will help you when you do not have the IR visible as the second graph or when working in a maximized view with just a single magnitude graph.

Let us play a bit with the window controls. Using the spin buttons of the FROM field, move the starting point of the window toward 0 and even a little bit beyond, like 1 ms. You will notice that the windowed magnitude does not start changing immediately at the time of zero but a little bit later. We would expect the direct sound to be excluded from the magnitude response as soon as the start of window passes beyond the arrival time of the peak. Why does that not happen? The answer is that SysTune uses a second, small rising window before the actual left window marker. To see it, zoom in just a little bit into the area around zero time, maybe from –5 ms to + 5 ms.

Directly before the window marker, there is another, smaller window. By default, it includes just a few milliseconds of data before the actual marker. This function has multiple purposes. On the one hand it allows for a bit of tolerance when placing the left window marker. You may not exactly catch the peak when setting the starting point of the window, but you still get a sensible result. On the other hand, in contrast to a discontinuous step-like start the smooth onset ensures data validity, especially for the lower frequencies. The direct sound signal might be spread slightly, depending on the dispersion of the signal and the time-dependence of the sources. A gently rising window start will include all components.

Tech-Note:
Before we look at further windowing options, we should briefly discuss why windows are applied and what their limitations are.

As stated earlier, the human hearing system derives a large part of the subjective listening impression from the direct sound and the early reflections. This includes tonal balance, coloration, perceived directions and precedence. Therefore, it is often useful not only to look at the full transfer function of the room, but also specifically at its early part. Windowing is used for that.
It may also be useful to study the frequency response of individual reflections. This can give some insight into where particularly high levels in the low or high frequency range come from.
Equally important, windows are utilized to remove noise disturbances and distortion effects from the impulse response. A window allows you to obtain a much cleaner view of the response behavior of the system under test. All of the noise that we have seen at the end of the ETC plots earlier can and should be removed for an accurate investigation of the transfer function.

There are some limitations for windowing as well. We have already spoken about the negative side-effects that they can have on the data. Rectangular windows especially may create undesired artifacts in the frequency domain, including phantom side lobes or an artificial noise floor. In addition, similar to the discussion of the relationship between the full IR and full TF, the length of the windowed impulse response determines the resolution of the windowed transfer function. It is important to note that the length of the window directly determines the lower frequency limit. For example, a window of only 5 ms length will have a lower frequency limit of 200 Hz. Here the same relationship applies that we introduced in the very beginning when we talked about the frequency resolution as a result of the FFT block size. The lower frequency limit equals the frequency resolution given by $\Delta f = 1 / \Delta t$, where Δt is the time length of the window.

Windowing Options and Options Window
In this section we would like to look more closely at the available window types. We have not explained the `OPTIONS` window of SysTune yet, so this is a good time to do it. You can open it by selecting `OPTIONS` from the `FILE` menu or by pressing the key `F9`.

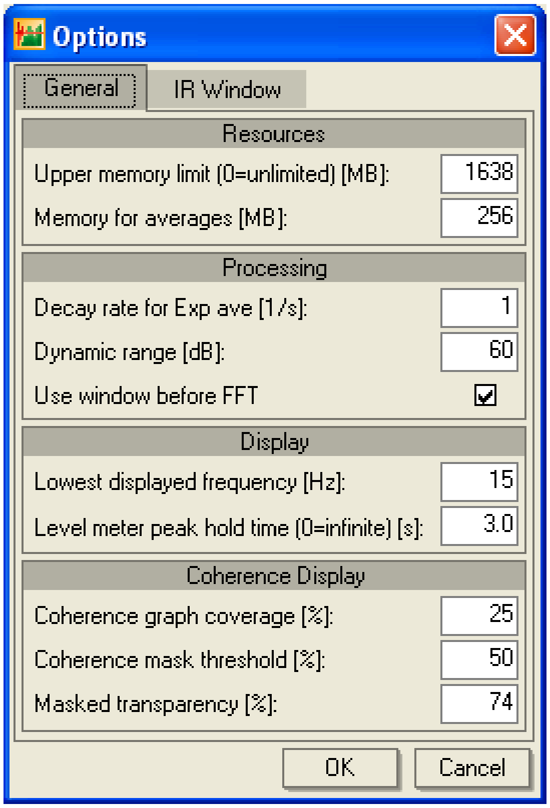

Before we return to the window types, we have to make a small excursion to introduce some general program parameters. These can be configured on the first tab of the Options window, the General tab. It contains options applicable to various areas of the software:

- Resources|Upper memory limit defines the maximum memory allowed for overlay data. This includes captured curves, loaded files and averages. By default, the program sets this memory limit to 80% of the physical RAM available.
    - If you experience problems with other applications when running SysTune try reducing this memory limit.
    - If you have a lot of RAM available and you can still use some more for SysTune, increase this limit.
    - A value of 0 will let SysTune use all memory available, even if this can lead to software instabilities.
- Resources|Memory for averages sets how much memory is used for the time averaging. The higher this memory allowance, the higher the refresh rate will be and the more time averages will be available in SysTune.
    - Use a value of 64 or 124 MB if you are limited in the PC memory available and not much interested in high refresh rates.
    - Select a value of 256 MB or higher if you would like to use high refresh rates and long averaging times.
    - The default value is 50% of the RAM available, but not less than 32 MB and not more than 256 MB.
- Processing|Decay rate for Exp ave sets the time constant for the exponential time average Exp. For more details please see the second chapter where we introduced this parameter.

- PROCESSING|DYNAMIC RANGE defines the threshold value for the deconvolution. When we discussed the deconvolution concept in the beginning of this chapter, we already talked about the Wiener method and the use of thresholds to avoid division by zero. This value represents such a threshold and is defined relative to the maximum of the reference signal. Any frequency points of the reference data below this level are discarded in the deconvolution.
    - If you are looking at an electronic circuit or a very stable environment (anechoic chamber) you may want to increase the dynamic range beyond 60 dB, if the stimulus signal provides that.
    - If you are receiving your reference signal from an electronic or even acoustic input, you may want to decrease the dynamic range to the approximate S/N of the reference signal. Otherwise noise in the reference signal will be included in the results of the deconvolution.
    - In most cases, a value of 60 dB is just about right. Normally, you should only have to change that for complex measurements with special requirements.
- PROCESSING|USE WINDOW BEFORE FFT ensures that a window is applied to the input signal before the FFT converts it into the frequency domain. The window used here is a Tukey (flat-top) window of 90%. It will be seldom necessary to switch the window off. One reason to do that is if your input signal has exactly the periodicity of the FFT size. In such a case, an FFT window is not needed because the window would only remove valid signal contents. The sole purpose of the FFT window is to suppress artifacts due to the non-periodicity of the input relative to the FFT block size. Its effect is most easily seen when you compare the spectrum of a sine signal at the input with and without the FFT window.
- DISPLAY|LOWEST DISPLAYED FREQUENCY controls the default for the lower view limit for graphs in the frequency domain.
- DISPLAY|LEVEL METER PEAK HOLD TIME defines how long the peak level is shown after the value was measured. For more details please refer to the section of this chapter where we introduced the level meters. Note, that for a value of 0, the peak is held infinitely long (it can still be reset manually by a right mouse click on the meter).
- COHERENCE DISPLAY – please also see the explanations about coherence for the background of these options.
- COHERENCE DISPLAY|COHERENCE GRAPH COVERAGE sets the area of the graph (magnitude, phase etc.) that is used for the coherence and IR stability curves. The default is 25%, which is the top quarter of the graph. Sometimes it may be useful to enlarge this area, especially when working with small screens. If you use coherence masking, you may also want to decrease this value to make the covered area smaller, because you do not need the exact numbers.
- COHERENCE DISPLAY|COHERENCE MASK THRESHOLD defines the value below which data points of the graph (magnitude, phase, etc.) are masked, that is made transparent or invisible depending on the next setting.
- COHERENCE DISPLAY|MASKED TRANSPARENCY is used to select the level of transparency of data points whose coherence value dropped below the threshold for masking. A value of 100% makes these invalid points fully invisible. With a value of 0% invalid data points can no longer be distinguished from valid data points.

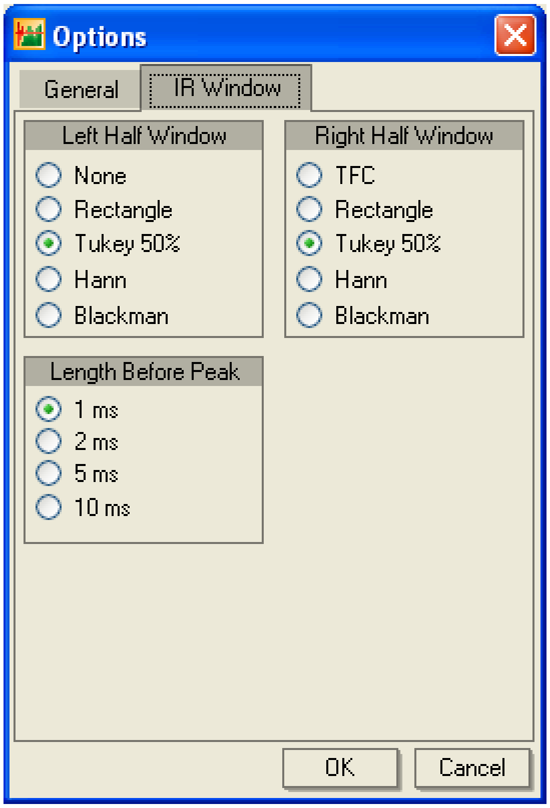

The second tab is the IR WINDOW tab. It controls the composition of the IR window. The LEFT HALF WINDOW is the short window to the left of the left window marker. Its length can be selected just below in the LENGTH BEFORE PEAK selection area. The RIGHT HALF WINDOW is the main window whose length you can adjust with the window markers.

For the LEFT HALF WINDOW you have the choice of:

- NONE means no window at all. This would include all data before the actual window start in the windowed IR.
- RECTANGLE leaves the IR at full magnitude for the length of the short window but discards everything before that.
- TUKEY 50% is a flat-top window with the weight of the later half being 1 and a rising cosine for the first half. This is the default.
- HANN is a rising cosine window that has non-zero and non-unity gain over its full length.
- BLACKMAN is an advanced cosine window and is a common choice in electronic and audio applications. Please refer to standard textbooks for more information.

The RIGHT HALF WINDOW can be one of the following windows:

- TFC is the new patent-pending windowing function (Time-Frequency-Constant Window), developed by SDA. It is actually not a simple window, but a complex weighting function that corresponds to a window with different lengths for different frequencies. See the next paragraph below.
- RECTANGLE is a rectangular window from start to stop. It leaves data unchanged between the starting point and the end point.

- TUKEY 50% is a flat-top window with the weight of the first half being 1 and a decaying cosine for the second half. This is the default.
- HANN is a decaying cosine window that has non-zero and non-unity gain over its full length.
- BLACKMAN is an advanced cosine window and is a common choice in electronic and audio applications. Please refer to standard textbooks for more information.

As an example, the following picture shows a LEFT HALF WINDOW that is a RECTANGLE of 10 ms length and a RIGHT HALF WINDOW that is a HANN window of 170 ms length.

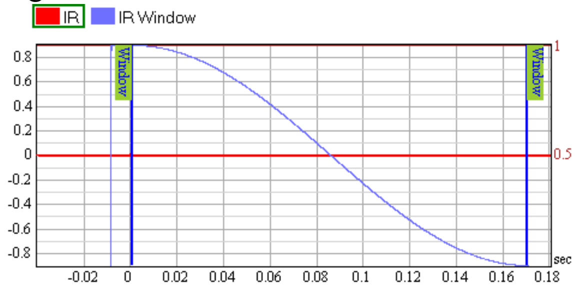

TFC Window™

The Time-Frequency-Constant (TFC) window is not really a window but only similar to ordinary fixed-size windows. In fact, the TFC window consists of a weighting function that depends on both time and frequency. For each frequency data point there is a different window length, namely for each doubling of the frequency the window length is cut in half. For example, the window length at 1 kHz is twice as long as the length at 2 kHz but only half as long as at 500 Hz. The actual window for each frequency is a Tukey 50% window.

Tech-Note:
In acoustic measurements there is often the need to look at the frequency response of the direct sound or the early arrivals only. This is difficult, when reflections that should be excluded are very close in time. On the one hand, you can not make the window very short, because then you lose all of the information for the low frequencies. In this respect we mentioned the limitations of windows in the previous section. On the other hand, you cannot select a very long window because then you would include all of the reflections as well, especially in the mid and high frequencies.

A solution to this is a window that has a longer length for low frequencies and a shorter length for high frequencies. The TFC window provides this functionality. It allows you to exclude reflections in the mid-to-high frequency range, because it can be set as short as needed. But at the same time, it is significantly longer for the low frequency range. This allows you to extract data for the lower frequencies, which may include some reflected energy, but that is better than no data at all.

The TFC window also images a characteristic property of our hearing system: Low-frequency signals have a longer integration time than high-frequency signals. Although in-depth investigations and listening tests for TFC window applications have not yet been performed, the usage so far indicates a strong correlation with our subjective impression of the signal's frequency content.

Finally, another remarkable property of the TFC window should be mentioned. Unlike other frequency-depending window implementations available until now, the TFC window is not based on octave or 1/3rd octave bands. In contrast, it uses a continuous weighting function that is able to provide a consistent magnitude and phase response without discontinuities or gaps.

The picture below shows a three-dimensional graph to help visualize the TFC Window.

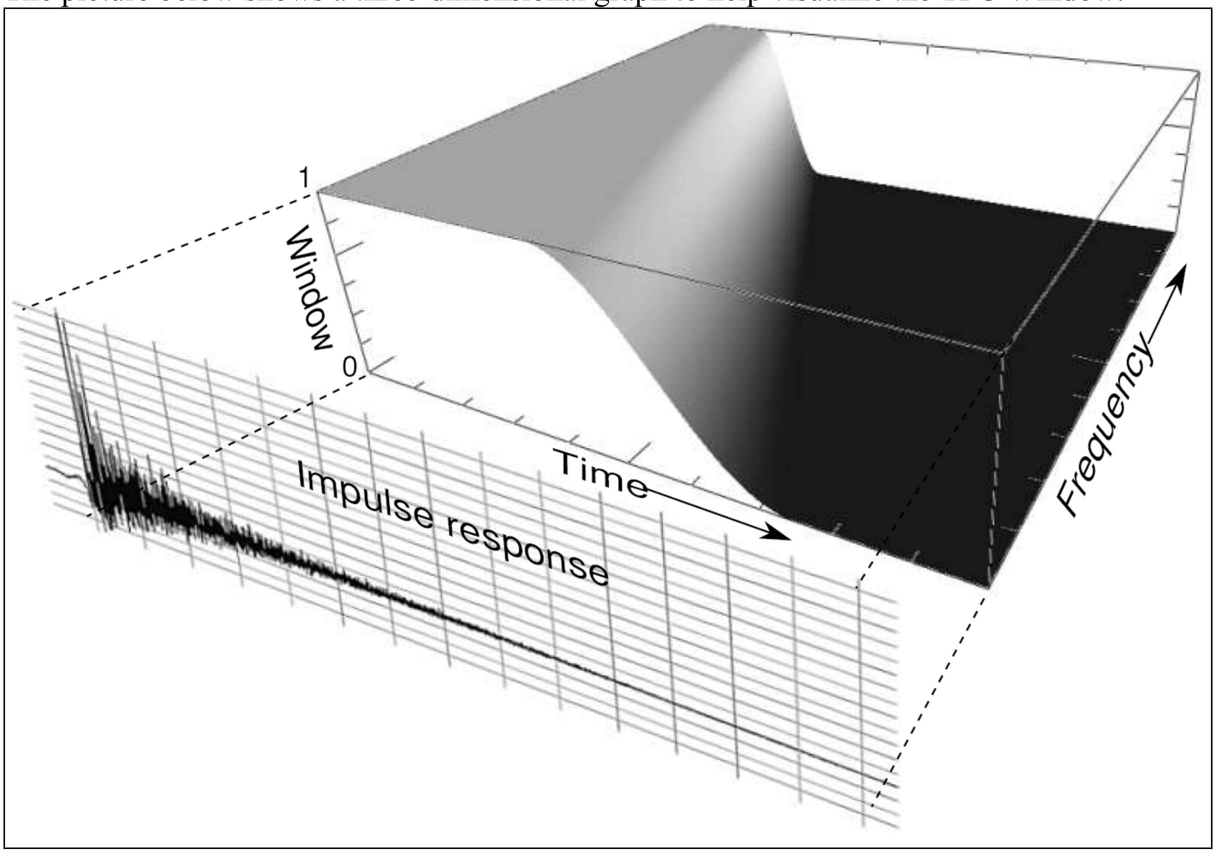

Let us come back to the Options window. The time differential between the left and right marker determines the time length of the TFC window at 1 kHz. To try this window out, select TFC for the Right Half Window and press OK. After that, the window curve in the IR graph will change to a different format. In fact, three curves will be shown, of which you should see two if you are still in the same zoom range as we selected before. Then you will see a picture like the one below on the left.

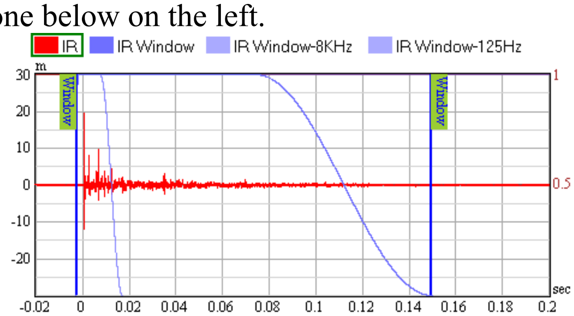

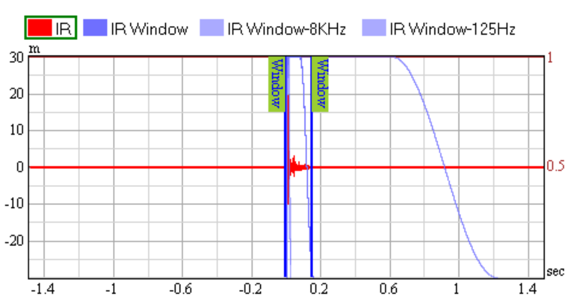

In addition to the window curve for 1 kHz it shows the curve for 8 kHz as well. It uses only 1/8th of the selected length. Zoom out to see the window length at 125 Hz; it is 8 times longer than the selected length for 1 kHz. This view is shown above on the right.

Depending on the impulse response you measured you will see some changes in the magnitude graph as well. If the room is rather dry you may get a result very similar to the original Tukey window, because there is not much to remove from the IR by means of any window. But if there is a fair amount of reflections and reverberant energy, you will see some change. For a dense distribution of reflections there will be an increase in level for the low frequency range, because of the greater length the TFC window provides, and a slight decrease in level for the high frequency range, due to the shortness of the TFC windowing function at higher frequencies.

Summary
This section discussed the measurement and display of the magnitude of the transfer function in great detail. We have learned how closely transfer function and impulse response are related, that they are two different views of the same data. We have talked about the reasons to investigate the frequency response of the system under test in audio applications. A variety of display options are available for the magnitude graph. In that regard our concern was mainly about the coherence and IR stability functions, which help to estimate the validity of results in a dual-FFT setup. Also, we introduced the gain setting for input channels and overlays and how it can be used to compensate for external gain changes or for averaging data in a meaningful way. After that we exported the results to a text file.
In the second part we explained what windowing is in general, where its limitations are and how you can apply windows in EASERA SysTune. We used different windows to look at the direct sound part of the impulse response. In this respect we also visited the Options window and talked about all of the program settings available here. Finally, we introduced the TFC window as an important new development for better windowing results in acoustic measurements.

4.5. Phase

The phase response is the second part of the complex transfer function. In SysTune it can be viewed using the `PHASE` button under `TRANSFER FUNCTION`. In contrast to the magnitude graph, the phase plot is often useless in acoustic measurements unless the measurement is windowed or it happens to be under special circumstances (electronic measurements, measurements of a very dry room, loudspeaker measurements). The reason is that the reflections and the reverberant energy of a room are usually incoherent and thus the phase becomes a random entity when derived from the whole impulse response or transfer function.

Nevertheless, phase is not meaningless. It is still being discussed widely if we can actually hear phase or not. This issue has not yet come to an accepted result among professionals. But there is still much to gain from looking at the phase response. To understand better why, we will talk a little bit about phase in general in the next section and afterwards we will discuss its practical usage in SysTune.

Meaning of Phase Data
We already mentioned that it is not quite clear if we can actually hear phase itself. But we can measure it and use it to analyze and adjust loudspeaker systems. That is possible because at least in audio applications, phase data consists typically of two components:

- On the one hand, there is the inherent phase of the system, you can think of it as a very small, frequency-depending delay that changes smoothly and slowly but often unpredictably over the frequency range of the system under test. For some systems the inherent phase can be approximately computed from the magnitude data, for so-called minimum phase systems. For other systems, the phase response is a fundamental property of their behavior, like filters in analog networks. Nevertheless, more complex systems like a loudspeaker in a cabinet will show an inherent phase response - as a function of frequency and of the measurement angle relative to the cabinet -, which cannot be described or calculated in a simple manner.
- On the other hand, there is always a delay included in the phase data, unless it was removed explicitly already. This delay is frequency-independent and may be composed of the latency in the electronic part of the system under test and of the time that sound needs to travel from the loudspeaker to the microphone. Often the electronic latency can be measured directly and the time of flight can be estimated based on distance between source and receiver and the speed of sound. Alternatively, impulse response measurements can be used to determine the delay as well.

This relationship can be quantified in a very simple equation for the phase of the transfer function:

Overall Phase = Inherent Phase + Delay Phase

Note, that all three values are still frequency-depending. Fortunately the phase change due to delay is well-defined, so that we can extend this equation

Overall Phase = Inherent Phase - 2 x PI x Frequency x Delay

This result is a very fundamental relation. We can use it to measure two things: Once the delay is known exactly, we can remove it from the overall phase and look at the inherent phase of the system. But we can also draw some conclusions about the delay. Roughly spoken, even if we do not know the inherent phase response, we know that its values are significantly smaller than the phase of the delay, at least for typical delays and frequency ranges. So the delay part will be dominant compared to the inherent phase.

These properties allow for some very practical applications in the tuning of sound systems. We can use the delay in the phase to align loudspeakers in time because in the phase plot it appears as a curve with a continuously rising or falling slope. This approach can be more accurate than the time alignment directly in time domain. Also, we can use this property to time-align components of loudspeaker systems that cover different frequency ranges. Simply put, the phase response for a two-way system that is not time-aligned will show a different slope in the high frequency range than in the low frequency range. If time-aligned, the phase curve will show a consistent slope over the full frequency range.

Tech-Note:
Because phase is so sensitive to small changes of the included delay time, it has long been considered as a side product only understood by a few enlightened people. But to be clear, phase itself is no mystery but just another simple number we can measure.

Like the magnitude, it is calculated from the real and imaginary part of the complex transfer function. Normally, phase data points are understood as points on a circle. They can only take values between 0 and 360 degrees, or −180 and +180 degrees if you like. Once computed, we can use phase to look at the transfer function in yet another way. Display options like unwrapping make the phase plots even more usable in practice.

Using the Phase Graph For Time Alignment

Some pages before we used the IR graph to time-align loudspeakers in SysTune. Now, we will come back to that and use the phase graph in addition. Again, let us start with our standard measuring setup, still using pink noise and the IR in the top graph. But this time switch the bottom graph to PHASE. Assuming you have already applied the PEAK TO DELAY function, you will see a picture like the following.

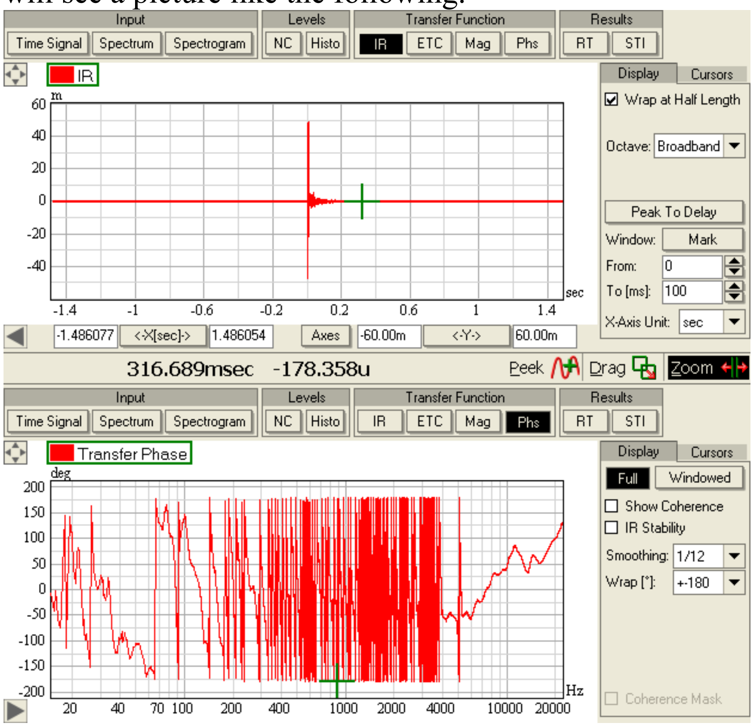

Depending on your measuring environment the FULL phase response may look quite chaotic. We have not switched to the WINDOWED view yet, but we will do that in a minute. Before that, have a look at the DISPLAY panel. Most options look familiar. You can enable SHOW COHERENCE and IR STABILITY like the magnitude graph. You can also apply SMOOTHING to the phase curve; it is set to 1/12 by default. We will talk about the subtleties of smoothing phase a little bit later. The second drop-down list labeled WRAP is new and only available for phase data. It allows you to select how the phase data is displayed. The default here is to show phase wrapped between –180 and +180 degrees.

In the previous section we discussed a lot about the delay included in the phase response - here we can demonstrate this principle easily. Just to exercise, press the 0 button for the DELAY OFFSET. This will delete the previous delay value for the current input channel and replace it by

0. After that, switch the WRAP drop down list to NONE to look at unwrapped phase and you will see a graph like this:

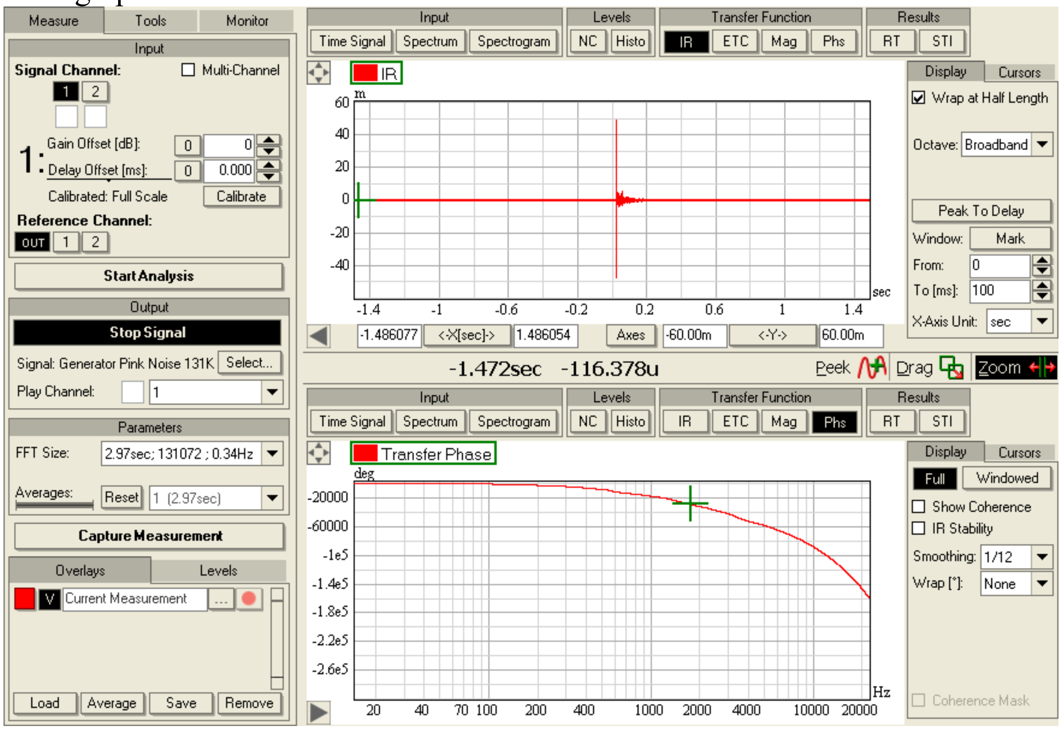

Note, that when we set the DELAY OFFSET to 0, the peak of the IR moved to the right, toward higher times. At the same time the phase response changed too, because it now includes the full delay of the impulse relative to zero. The phase graph shows just that. Look back to the equation above, it is obvious that this relationship is reflected here and one can see how strongly the delay dominates the phase response. The slope was not so steep before, if there was any at all. You may press PEAK TO DELAY and after that 0 once more to compare how the delay in the IR affects the phase response. With the delay uncompensated the phase response is much steeper.

Obviously this effect can be used to find the exact time arrival (peak) of one loudspeaker and align other loudspeakers to it to achieve a coherent response behavior. But in order to avoid measuring errors and include a lot of data that degrades the accuracy, we need to place a window first. We have already practiced that in the previous section, and actually you may still have that window in place. Press PEAK TO DELAY to move the peak of the IR to the center of the top graph. Zoom into an area of –50 ms to 200 ms. Switch the PHS graph to the WINDOWED view and enter 0 and 30 as the window limits. Make sure you are working with the Tukey 50% window, otherwise switch to it under OPTIONS. You should get a view like this one:

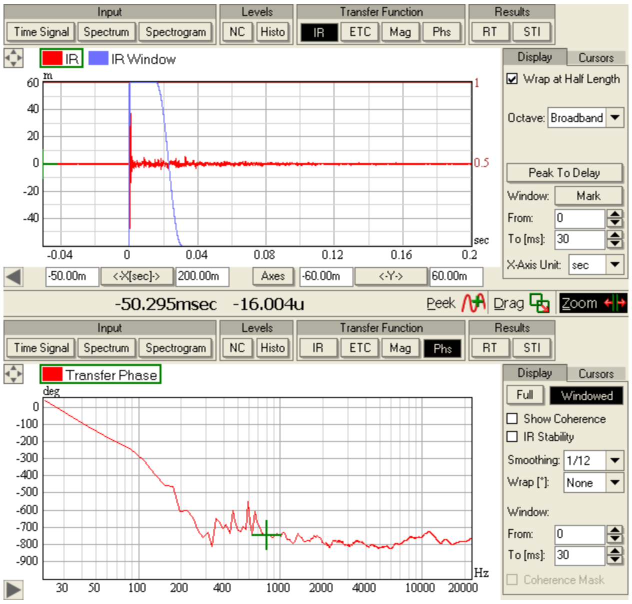

Of course, the phase graph depends on a lot of variables; do not expect yours to look exactly like the one in the picture. You will also have to rescale the phase plot to view limits adequate for your data. However, it is important to note, that if the IR has the major peak located at a time of 0 and any other significant peaks are excluded by the window, the phase response will be approximately flat in the frequency range where data is defined (check magnitude graph and/or coherence). To determine the delay more accurately than by the maximum value of the IR – that is what PEAK TO DELAY does – we can manually try to make the phase response even flatter.

We can accomplish this in two ways, either by adjusting the delay of the input channel under DELAY OFFSET or by manipulating the starting point of the WINDOW. For the moment, use the spin buttons for the DELAY OFFSET to shift the IR peak slightly to lower or higher times. You will immediately see an effect in the phase plot as well. The curve will start to change its slope with every sample that you add. You can keep updating the DELAY OFFSET setting until the phase response is maximally flat. When that is the case, the same text field will tell you the exact overall delay time for that channel. You can keep this measurement and delay setting as a reference and then align loudspeakers to it in the same manner.

Another way to change the delay time included in the phase response is to manipulate the starting point of the window. Using the spin buttons close to the WINDOW FROM text field in either IR or PHASE graph you can adjust the start of the window, which is the reference time for the windowed phase response as well. However, now your final overall delay will consist of the delay under DELAY OFFSET plus the delay that was added by shifting the window onset. Generally, this second option is more useful to investigate the phase response of parts of the IR that are not located at time zero. The advantage is that you can do that without having to change the delay compensation for the whole input channel.

Tech-Note:
We would like to emphasize very briefly that this section talks about delay and phase in a very simplified fashion. There are more subtleties to this topic than it may seem at first glance, here are some examples:

We already mentioned that you can only time-align loudspeakers to a single spot or to a particular direction if you are far enough away. Another point worth discussion is where and how to place the IR window as it will naturally affect both magnitude and phase. In addition, it is often hard to examine phase if there are several peaks of equal level, which are located very close to each other, whether that is direct sound arrivals or early reflections. In such cases there may be no flat phase response at all or vice versa different delay settings provide a similarly flat phase response. Environmental conditions such as temperature and wind affect the propagation of sound, as well.
All of this should not scare you, but should make you aware of the complexity of this issue. This software is only a tool and not a silver bullet; only knowledge and experience will help you to deal with such topics.

The second powerful application of the phase response is the time-alignment of multi-way or frequency-shaded loudspeaker systems. As shown above, here the delay in the phase response also plays the main role. Imagine that we have a raw two-way loudspeaker instead of a fine-tuned full-range system. Its woofer and horn may not yet be aligned. With the crossover filters in place, you can imagine that the phase response will show the delay of the woofer in the low frequency range and the delay of the horn in the high frequency range. To align them precisely in time is rather difficult using impulse responses, because the peaks have very different shapes. But often they can be aligned relatively easy using the phase response.

The next picture shows the phase of the transfer function of such a simple two-way loudspeaker. It is already roughly aligned, but as you can see from the plot, its phase response still contains two different delay times. There is an approximate plateau up to 1.5 kHz and then, beyond the crossover range, there is a slope of fairly constant decay between 1.5 kHz and 10 kHz. This roughly corresponds to the woofer and the horn with the window start being adjusted to the delay of the woofer.

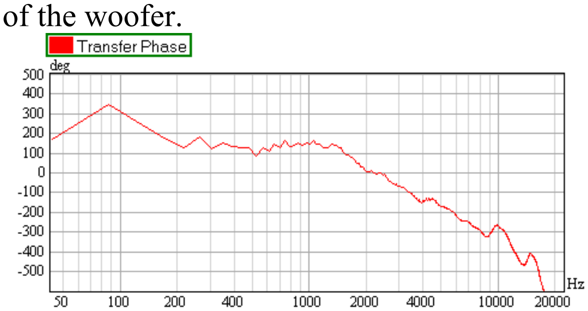

An improved configuration may look like the following measurement.

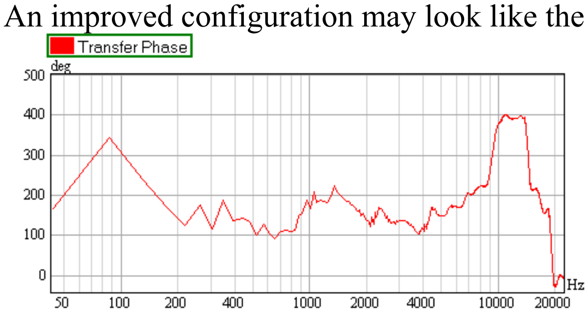

Here, both components have an approximately flat phase curve, which corresponds to a delay of zero. This may be even more obvious in the wrapped display shown below.

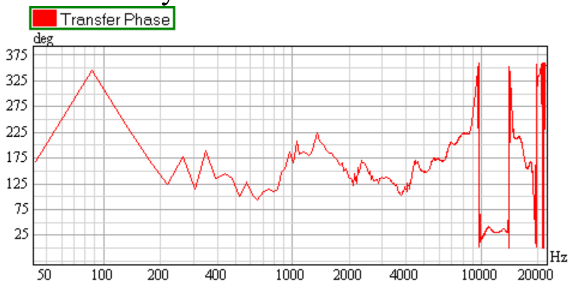

Smoothing and Wrapping Phase

Previously we have used two options called SMOOTHING and WRAP without much explanation. But their background is certainly noteworthy and we would like to make up for that here.

Wrapping phase means compressing phase values to a range of 360 degrees. Mathematically this is an absolutely accurate procedure without loss of information. In fact, when transforming the complex transfer function from real and imaginary part into magnitude and phase, there are only phase values in the range of 360 degrees. A process called unwrapping unfolds phase values to an unlimited range for easier viewing and analysis. The reason for this is simple. By nature, (wrapped) phase is a circular value which means that on a scale between 0 and 360 degrees the data points of 0 and 360 are identical, 1 and 361 and so on. In fact, there are no values greater than or equal to 360 or values less than 0 degrees. However, a smooth function on a circle shows discontinuities when mapped to a linear, noncircular scale. A small step between 359 degrees and 0 degrees turns into a large step there, because these two points are on opposite ends of the phase axis.

Unwrapping processes the phase values sequentially from the first point to the last. Whenever a gap of more than 180 degrees is encountered between adjacent data points, the value of all following points is shifted by 360 degrees, so that the difference between the two neighbor points is again smaller than 180 degrees in the processed data. The resulting curve will show more clearly the general trends in the phase response. They are often hidden by the many discontinuities when wrapped data is displayed on a linear scale.

The available options for WRAP include:

- `+/-180` shows the phase wrapped to a range between –180 degrees and +180 degrees.
- `0-360` shows the phase wrapped to a range between 0 degrees and 360 degrees. Depending on where the discontinuities in the wrapped phase data are, this plot may be more useful than the previous. Remember, that for well-aligned systems you will have a fairly flat phase response and may not need to unwrap phase significantly.
- NONE shows the phase data unwrapped as explained before.
- `+/-360`, `+/-540` and `+/-720` show the unwrapped phase, wrapped back to a larger range of values. These ranges are useful when the phase response is rough but not very steep.

SMOOTHING phase works like the smoothing function for magnitude data. Every data point represents the average value over the selected bandwidth. For phase, the data is unwrapped first, then smoothed and then wrapped back to the desired display range. However, especially with respect to phase one must be careful to use smoothing for large bandwidths.

The main purpose of smoothing is to remove details and noise from the response. But before a phase curve can be smoothed it must be unwrapped first. Only a bit of noise can already make the unwrapped curve look quite different, if the variation happens in a range where the difference between adjacent data points is close to 180 degrees. Here smoothing the resulting curve can lead to almost arbitrary values. Therefore it is strongly recommended to check the continuity of the phase data before activating wide-band smoothing. Adjacent phase data points should be close to each other, generally within 90 degrees on the circular axis. For example, degree values of 90 and 180 or 300 and 30 are allowed for neighbor points, but 30 and 210 are not. If this is not taken into account, you may see smoothed phase curves that are less coherent than the original phase curves.

5. Further Measurements

This chapter is concerned with investigations that are based on the further processing of measurements we have introduced in the previous chapters. While our walk through SysTune from the input signal to the transfer function was rather straightforward so far, we will now look at additional applications of the software. They will not require any extra hardware but just the setup that we already have.

5.1 Level Measurements

Noise Criteria

In many aspects of daily life we encounter noise as a disturbing effect, whether this is random noise like street traffic and air conditioning systems or more structured noise from machines and vibrations. Several standards exist that define how to measure and quantify such noise levels, and one of them is NC which stands for noise criteria (ANSI standard S12.2-1995 (2003)).

NC can be measured in SysTune very easily. First, make sure that you are not playing any signal and that you have calibrated your `Signal Channel`. Then switch to `Levels|NC` from the button row at the top of the graph. After that, the display will change to an overlay of the `Current Measurement` and the major NC curves from `NC20` to `NC70`. According to the standard, these are all octave-based curves.

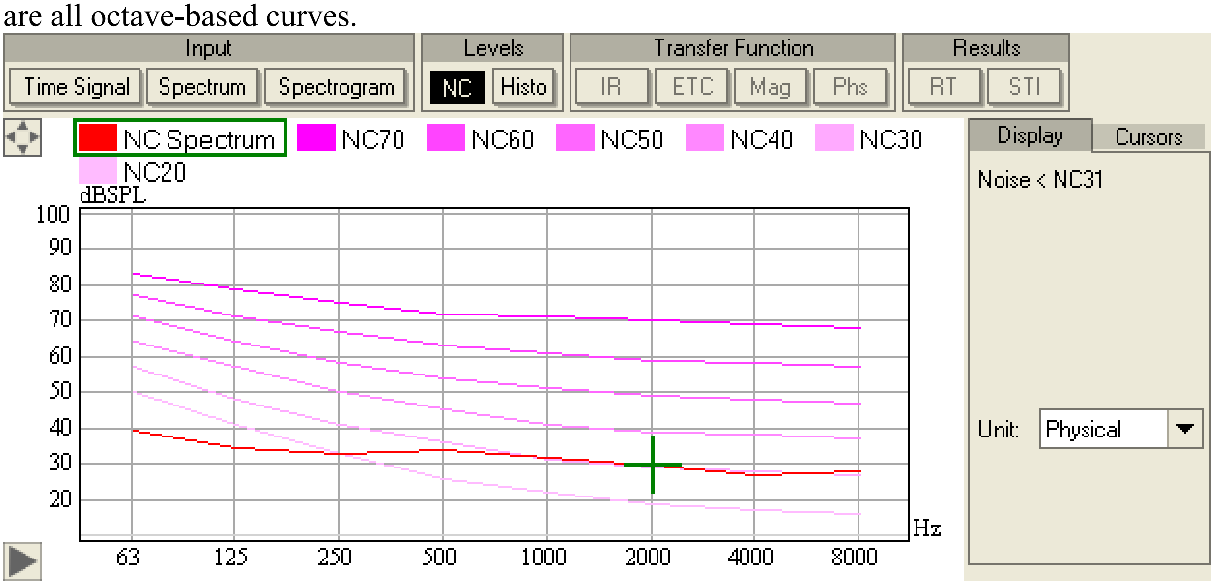

This graph already gives you an indication of how much environmental noise there is. The broadband NC number derived from the octave data is shown in the `Display` panel to the right. Typical NC measurements are:

- NC10 Broadcast and Recording Studios (distant microphone pickup used)
- NC10-15 Concert halls, Opera houses, and Recital halls
- NC15-20 Large auditoriums, Drama theatres, and Churches
- NC15-25 TV and Broadcast Studios (close microphone pickup only)
- NC25-30 Conference rooms (large), Lecture halls and Classrooms
- NC25-30 Small auditoriums

- NC25-35 Meeting/banquet rooms
- NC27-37 Movie theatres
- NC30-35 Conference rooms (small) and private offices
- NC30-35 Churches (small)
- NC33-37 Courtrooms
- NC38-43 Restaurants

Be aware that if the SIGNAL CHANNEL is not calibrated or you switch the UNIT to DIGITAL FS you will still see a curve but you cannot measure NC. Note also, that at present you cannot capture NC data in an overlay. But as for other graphs, the NC curve can be saved to a text file using the FILE|EXPORT DATA AS TEXT menu.

SPL and LEQ Measurements
The question of what is considered noise and what is not, depends on the point of view. The audience of a rock concert will usually not think of the music as noise. However, people living nearby who do not share the enthusiasm for fast guitars and loud drums will think differently.

In practice, such contradictions can only be solved by an agreement on acceptable sound levels at various spots. The organizer of an event has to make sure that locations in the audience area are covered well enough while spots outside of the venue should more or less not be exposed to sound pressure levels that are too high. In addition, to avoid hearing damage it is often necessary to monitor the average levels that listeners in the venue are exposed to.

The exact definition for the measurement and averaging of such levels is given by a set of standards. In EASERA SysTune, some of these standardized methods are implemented as SPL and LEQ measurements (IEC-61672-1:2002 and ANSI_S1.4-1983_(R2006)).

Essentially, you can view SPL and LEQ measurements in SysTune in two ways:
- as an instantaneous value, that is displayed as a number, or
- as a histogram, that is a plot of the measured levels over time.

SPL and LEQ Monitor
To look at the current sound pressure level at the input of the Signal Channel, switch the control panel on the left to the MONITOR tab.

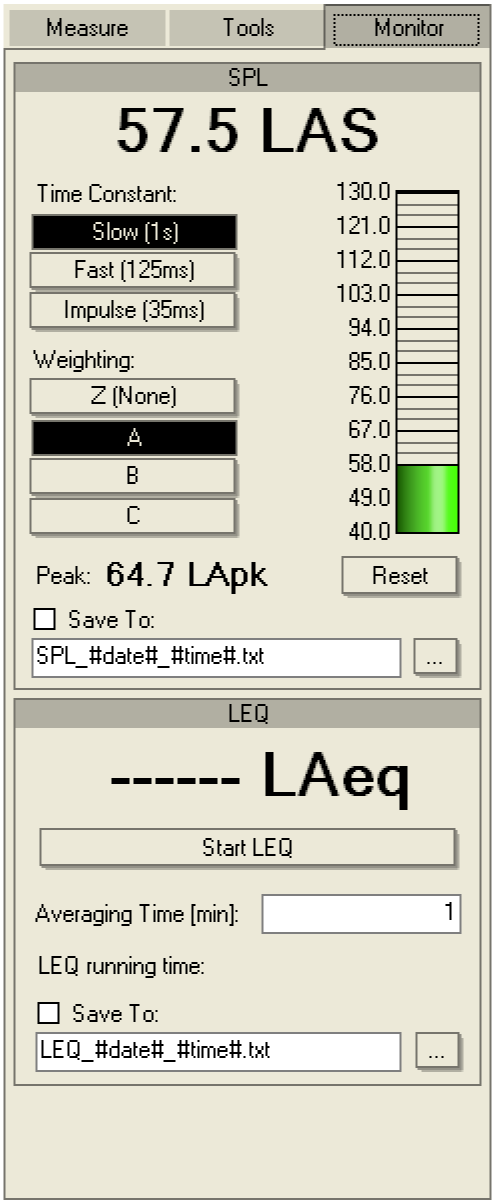

Assuming that your SIGNAL CHANNEL is calibrated to pressure, the upper frame labeled SPL shows the actual, broadband sound pressure level at the input. Below, the frame labeled LEQ shows the controls for the measurement of equivalent long-term averaged sound pressure levels. Most of the items on the MONITOR tab are self-explanatory, but for the sake of clarity let us go through them step by step.

The first number shown in the SPL frame is the broadband sound pressure level at the input. It is derived directly from the time signal and uses the TIME CONSTANT and WEIGHTING as selected below. The TIME CONSTANT is a decay constant that defines how an instantaneous event loses energy in the level average as time passes by. SLOW denotes a rather long decay time, the event is remembered in the average level for a while. FAST is a short decay time, thus the average level changes more quickly. IMPULSE is a time constant with a fast rise and a slow decay. All three constants have been implemented in SysTune according to the existing standards and are comparable to corresponding selections of a handheld sound level meter or RTA. The WEIGHTING of the data can be either Z (NONE), A, B or C. These choices have already been introduced in detail in the second chapter. A visual indication of the SPL value is given as a meter to the right. By default, the SPL display uses SLOW for the TIME CONSTANT and A for the WEIGHTING.

Below the WEIGHTING controls another numerical value is shown, the PEAK value. This is basically the maximum of all SPL values since the monitoring process was started or reset last. The RESET button to the right of the display allows you to restart the PEAK metering.

At the very bottom, the check box SAVE TO allows you to save the SPL values continually into a text file. Its location can be selected via a file dialog that appears when you press the small … button to the right of the text field. You may also enter a file name manually. Note, that when the placeholders #DATE# and #TIME# are used, the starting date and time for the logging process is included with the file name.

The LEQ frame has the same controls for file saving as the SPL frame. However, it stores the A-weighted sound pressure level based on a long term average rather than values for each short period of time. That AVERAGING TIME can be entered in the same frame. It is located right above the indicator LEQ RUNNING TIME which shows the overall time duration since the LEQ was last started. The default value for the AVERAGING TIME is one minute. To actually start the calculation process you have to press the START LEQ button. In contrast to SPL, this command is needed for LEQ, because the monitoring requires additional performance and memory that should only be used when needed.

Histo Graph
In addition to just viewing the instantaneous levels and writing them into a file, you may also view the level history graphically. This histogram graph is available through the HISTO button just to the right of NC in the LEVELS group of buttons. The first time you press that button, the graph for the SPL HISTORY will be shown:

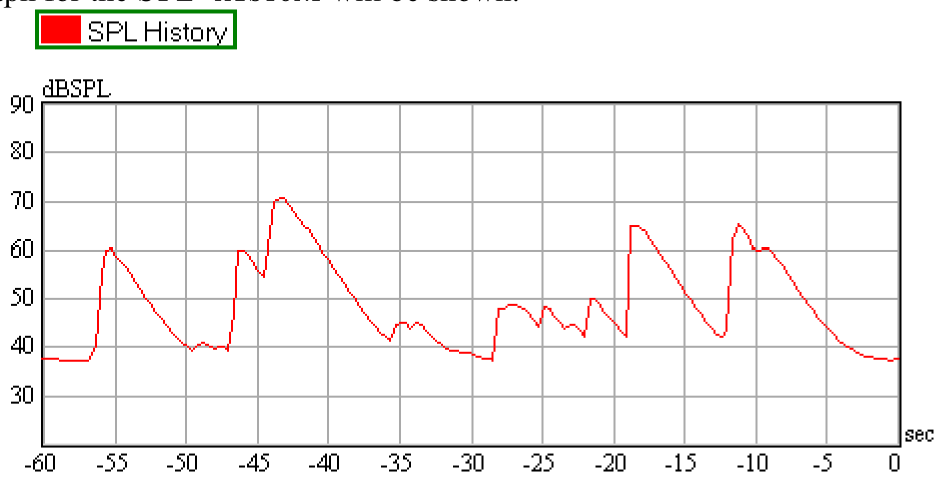

The picture here shows that graph after a monitoring time of over 60 seconds. To the right of the graph you will find two buttons, SPL, which is active by default, and LEQ, which is off by default.

The graph will switch to LEQ HISTORY when you activate the LEQ button.

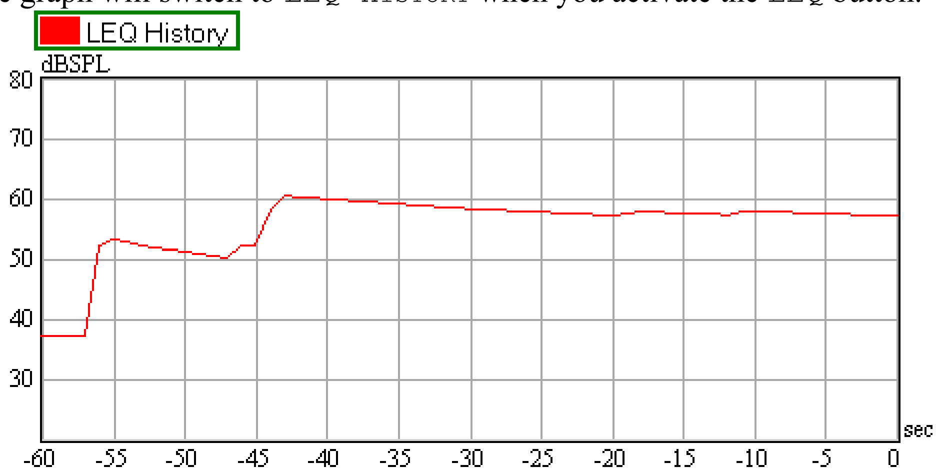

If you are still in a split view with two graphs, adjust the top graph to show SPL HISTORY and the bottom graph to show LEQ HISTORY. You may notice, that if LEQ was not started yet using the START LEQ button mentioned above, the program will do that for you when you select LEQ from the DISPLAY panel. It will also automatically switch to the MONITOR tab on the left. We remark that in contrast to all other graphs, LEQ HISTORY and SPL HISTORY each are only available for one graph at a time.

By default, the two graphs will only show the first 60 seconds of time. Use the axis settings (small triangle at lower left of the graph) to change the view to longer time scales like one hour. Because we are looking into the past, you will have to enter –3600 seconds for the start value of the X-axis. After a time period of one hour the two graphs will look similar to this:

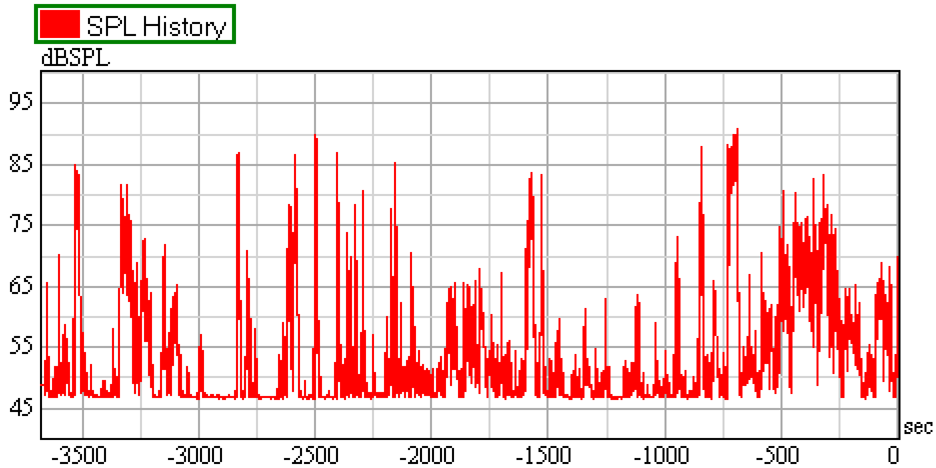

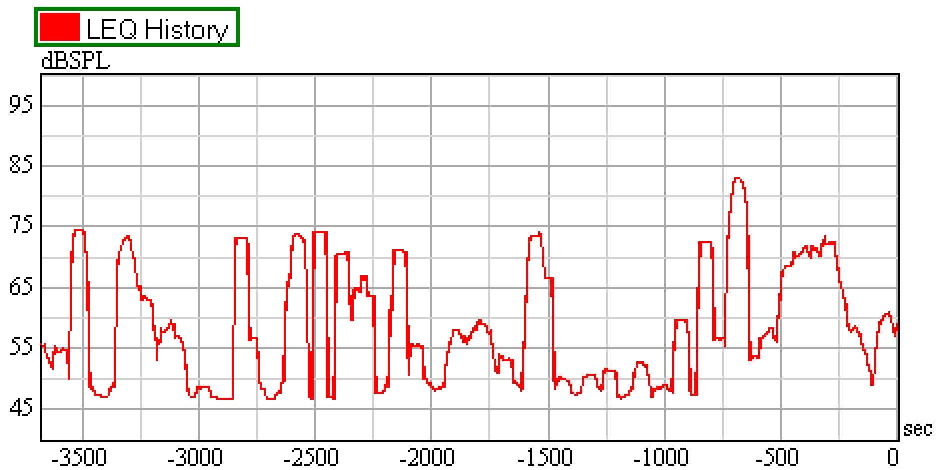

Note, that STOP ANALYSIS will also stop the monitoring process. When restarted, the old history graph will be cleared. Any other changes, such as parameter changes or switching input channels will also restart the histogram plot. With regard to the log file, such changes will be annotated automatically.

Summary
In this section we have shown how to make simple NC, SPL and LEQ measurements. We have also discussed how SPL and LEQ measurements can be logged to a file for later evaluation. The Histo graph can be used to view the level history while the program is running.

5.2 Measurements using Multiple Signal Channels

The previous chapters of this manual were concerned with measurements using a single channel or two channels with one channel being the reference. But EASERA SysTune supports up to 8 channels for simultaneous monitoring, measuring and display of averaged data.

Imagine a typical stadium with a number of different zones. Using a simple dual-channel setup you would have to carry the microphone around to make measurements sequentially. Or you would need a microphone switcher and still you can only use one microphone at a time.

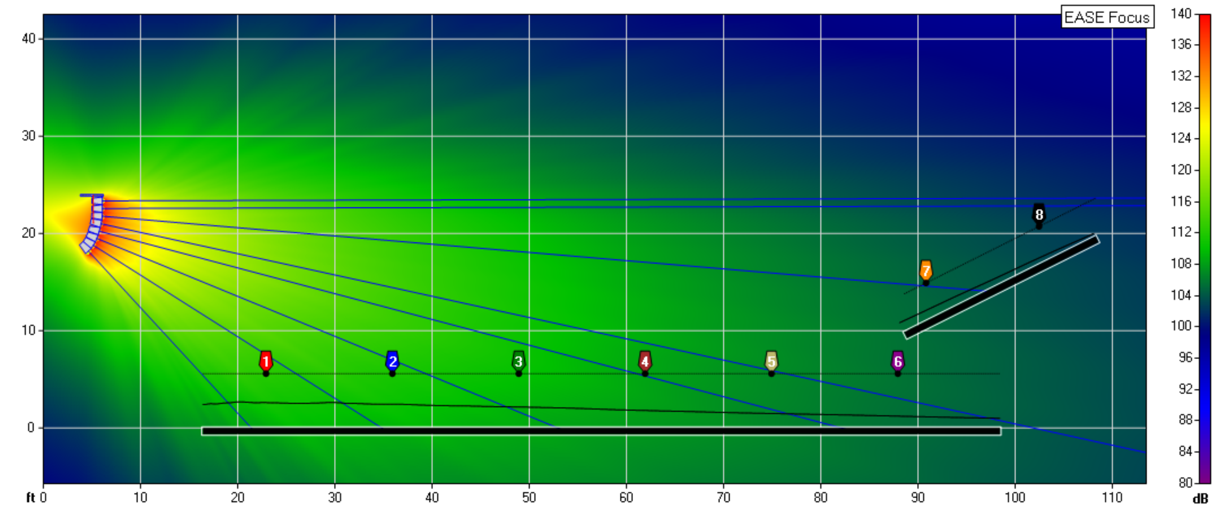

In contrast, with SysTune and a multi-channel soundcard you can handle up to 8 microphones at the same time. The software enables you to monitor all of the channels and to make detailed measurements with each one of them. You can also look at the averaged spectrum and transfer function of up to 8 channels in real-time. In this part we will go through these innovative new features step by step.

Changing Soundcard and Driver

By default, SysTune starts with the Windows default sound driver. This is normally a driver using the Direct Sound protocol which supports only 2 channels. Up to now, only professional sound drivers on the basis of Steinberg's ASIO provide access to more than 2 channels. Therefore we need to switch SysTune to use a different configuration.

To do that, press Select Device ... in the Configure menu. This will open the Select Device dialog window. Note, that immediately when the command is processed the real-time analysis is stopped.

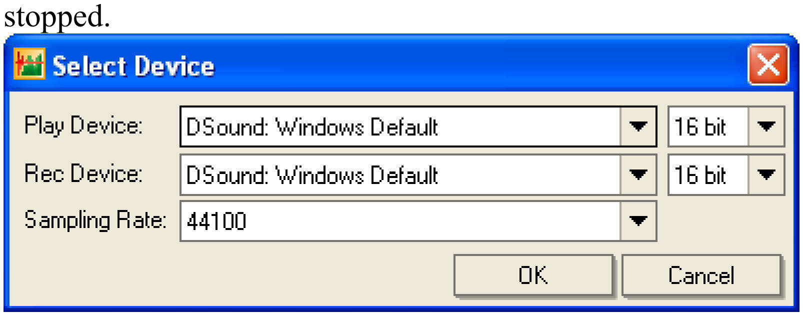

Through this window all of the control parameters for the configuration of the sound driver can be accessed. The drop down list to the right of the label Play Device allows you to select any of the available DSound (Direct Sound), Wave or ASIO drivers available for the output from the software. Similarly, the drop down list labeled Rec Device shows the active driver for the input to the software. Below that, the Sampling Rate can be chosen.

On the right hand side, the bit depth for input and the output can be selected as well. This setting controls the dynamic range used for the analog/digital and digital/analog conversion. For acoustic measurements 16 bit is usually enough as this corresponds to 96 dB between the maximum and the minimum value of a sample. For electronic measurements a higher bit resolution is recommended, such as 24 or 32 bit. However, be aware that for most AD/DA converters these are just nominal rates. Standard soundcards will not be able to provide more than an effective number of 20 or 21 bits of resolution. This is equivalent to a dynamic range of about 120 dB.

Hint: For ASIO drivers, the Select Device window will show an additional button that gives you direct access to the ASIO control panel. For some devices, you can use that to adjust the gains and the sample rate of the soundcard.

If you have a soundcard connected that supports more than 2 channels, select the corresponding ASIO driver for the input and for the output. If you do not have one at hand, you may as well stay with the default driver DSound: Windows Default and use its 2 channels for the following exercise. Here we will use the ASIO driver of the EASERA Gateway AD/DA.

Press OK to confirm and close the dialog.

Status Bar

After the Select Device window was closed, the status bar at the bottom of the window will immediately reflect the changes.

Depending on your current configuration it will show the following text:

- With a Signal Channel only:
  Live Signal <Name of Rec Device> <Sample Rate> / No Reference

- With a Signal Channel and an Output Signal selected:
  Live Signal <Name of Rec Device> <Sample Rate> / No Reference / Play Signal : < Name of File or Name of Internal Signal> on <Name of Play Device>

- With a Signal Channel, an Output Signal and Out selected for the Reference Channel:
  Live Signal <Name of Rec Device> <Sample Rate> / Play + Reference Signal : < Name of File or Name of Internal Signal> on <Name of Play Device>

- With a Signal Channel, an Output Signal and 1 to 8 selected for the Reference Channel:
  Live Signal / Live Reference <Name of Rec Device> <Sample Rate> / Play Signal : < Name of File or Name of Internal Signal> on <Name of Play Device>

- With a Signal Channel and 1 to 8 selected for the Reference Channel:
  Live Signal / Live Reference <Name of Rec Device> <Sample Rate>

- In File Mode:
  Opened File: <Name of File>

- Otherwise:
  Setup not configured.

The status bar will also be updated by the software when you change the Reference Channel or the Signal.

Hint: The status bar is especially useful to quickly check the current setup with regard to the currently selected soundcard and sample rate. This means it is not necessary to open the Select Device window to look it up.

Multi-Channel Measurements
Making measurements with more than a single channel is simple in SysTune. We have already seen in the second chapter, that the mini-meters right below the Signal Channel buttons allow you to monitor up to 8 channels. Depending on the current soundcard and driver, this row of buttons and meters will show up to 8 elements according to the number of input channels. To the right of the label Signal Channel you will find the Multi-Channel check box. Select this check box and then click on another Signal Channel button, like 2.

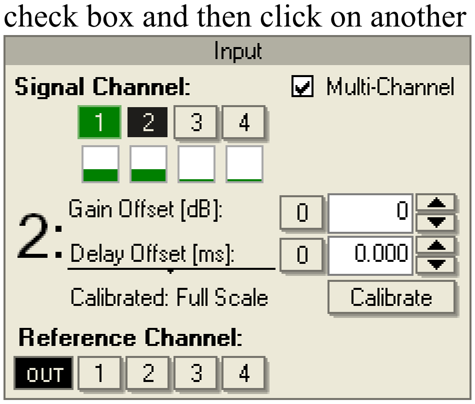

Now your display will have changed in two ways. At first, when activating Signal Channel 2 the button labeled 1 was not deactivated but it changed its color to green. This means that this channel is used in the multi-channel averaging. Second, the top and bottom graphs will be replaced if they are not showing Spectrum, Spectrogram or Mag graphs already. That is because the averaging process happens as a power average in the frequency domain. Essentially, it is the same calculation that we used on a sequential basis before when we discussed averaging overlays for the input spectrum or for the magnitude of the transfer function. Therefore, only the buttons for SPECTRUM, SPECTROGRAM and MAG are available when MULTI-CHANNEL is active.

With respect to all other functions almost all of them are also available in multi-channel mode. For example, you can apply time-averaging as well and you can capture overlays. Just as it was done for averaged overlays, you can apply gains to the individual channels to add weight to them in the averaging calculation. Because of performance reasons, IR windows are not available in multi-channel mode.

Hint: You can employ the settings for FFT Size and Delay Offset to achieve the same effect like the IR window. First you will have to set the delays for the channels individually, then choose your desired window length as the FFT size for the multi-channel measurement. However, note that the FFT block will use either a Tukey 90% window or none, depending on the setting for the FFT window in Options.

To return to the single-channel mode just deactivate the MULTI-CHANNEL check box. In multi-channel mode you can add input channels by just clicking on the corresponding button in the SIGNAL CHANNEL row. You can activate a channel that is already part of the multi-channel measurement (indicated by green color) by clicking on the related button. To remove a channel from the averaging process, activate it first and then click again on the same button. After that, another channel will be set active automatically.

Tech-Note:
It has already been mentioned but should be emphasized once more here that averaging measurements is a method that can provide some interesting answers and improved insight. But no system tuning process should rely solely on that.

For example, a deep and broad gap in the frequency response of one of the measuring locations will cause a bad listening impression at that spot although it may hardly be recognized in an average of 4 channels or more. Vice versa, a large peak in level may cause the whole average to show that peak, although it does not exist for any other location than just one. Averaging only makes sense if the included data sets are reasonably comparable to each other.

To be clear: Although innovative and new, the multi-channel averaging function in SysTune can give you only one more powerful tool to look at the whole system from yet another perspective. It is not the Silver Bullet.

5.3 Reverberation Time and Speech Intelligibility

In chapter 4 we have already introduced typical properties of room-acoustic impulse responses and how they must be understood. The reverberation time (RT) is a fundamental measure in room-acoustics and can be derived directly from the impulse response, we will discuss this in the following section. After that we will look at one of the most accepted measures for speech intelligibility, namely the speech transmission index (STI).

Reverberation Time

The reverberation time is used to characterize rooms of small to large size, whether it is a recording studio or a large cathedral. For this, a principle assumption is made: Exciting the room with a stimulus signal loud enough to fill the whole volume and then switching off the signal provides a decay curve for the sound energy in the room. When there is no more signal energy supplied to the room, the energy in the room will continuously decrease due to absorption by walls, ceiling, floor, by the air and by the audience as well.

This decay rate is measured in seconds. It states how long it takes for the pressure level of the sound field to drop by 60 dB. This time corresponds roughly with our subjective impression of how long it takes after we clapped our hands once and cannot hear the room echoes anymore.

Tech-Note

In most software applications the reverberation time is derived by calculating the Schroeder backward integral - this corresponds to the sound energy in the room – from the squared impulse response. This integral represents the energy decay over time from which the reverberation time can be derived. In practice, due to signal-to-noise issues it is usually not possible to obtain the full time for a drop of 60 dB. Therefore, shorter times are measured, such as for 20 dB and then extrapolated to 60 dB.

In SysTune the reverberation time displayed is T20, as defined by ISO standard 3382. It corresponds to the decay time from –5 dB of the Schroeder integral down to –25 dB and is then extrapolated to include 60 dB of decay. Generally, calculations are made based on the octave-filtered impulse responses, so the resulting RT curve is octave-based also.

For most rooms the frequency-dependence of the RT cannot be ignored. It typically shows a decrease toward higher frequencies which is mainly caused by the attenuation of sound energy by the air. When people speak about the RT of a room they usually mean the average reverberation time in the mid-frequency bands, usually from 500 Hz to 2 kHz.

Note also, that the reverberation time is only an approximate measure for the whole room. Depending on the measuring location, on the sources and the geometry of the room, it may vary significantly between measurements. Make sure that you evaluate several RTs before drawing any conclusions. It is also good practice for reports including RT numbers to add the variation throughout the venue.

To view the reverberation time as a function of frequency, press the button RT in the RESULTS group of buttons, directly to the right to the PHS button. Since you may not have a measuring setup available to derive a reasonable reverberation time, let us use a loaded file this time.

Go to FILE|OPEN AUDIO FILE and load the example file FinalMP2New from the IR subdirectory of your SysTune data folder. This will switch the software to the File Mode, which is a static program mode to view files. You cannot perform any measurements in this mode, but you may switch back to the real-time analyzer at any time using the button Switch to Live Mode in the control panel on the left.

After the file is loaded, switch the top graph to ETC and the bottom graph to RT. You will see a window like the one shown below:

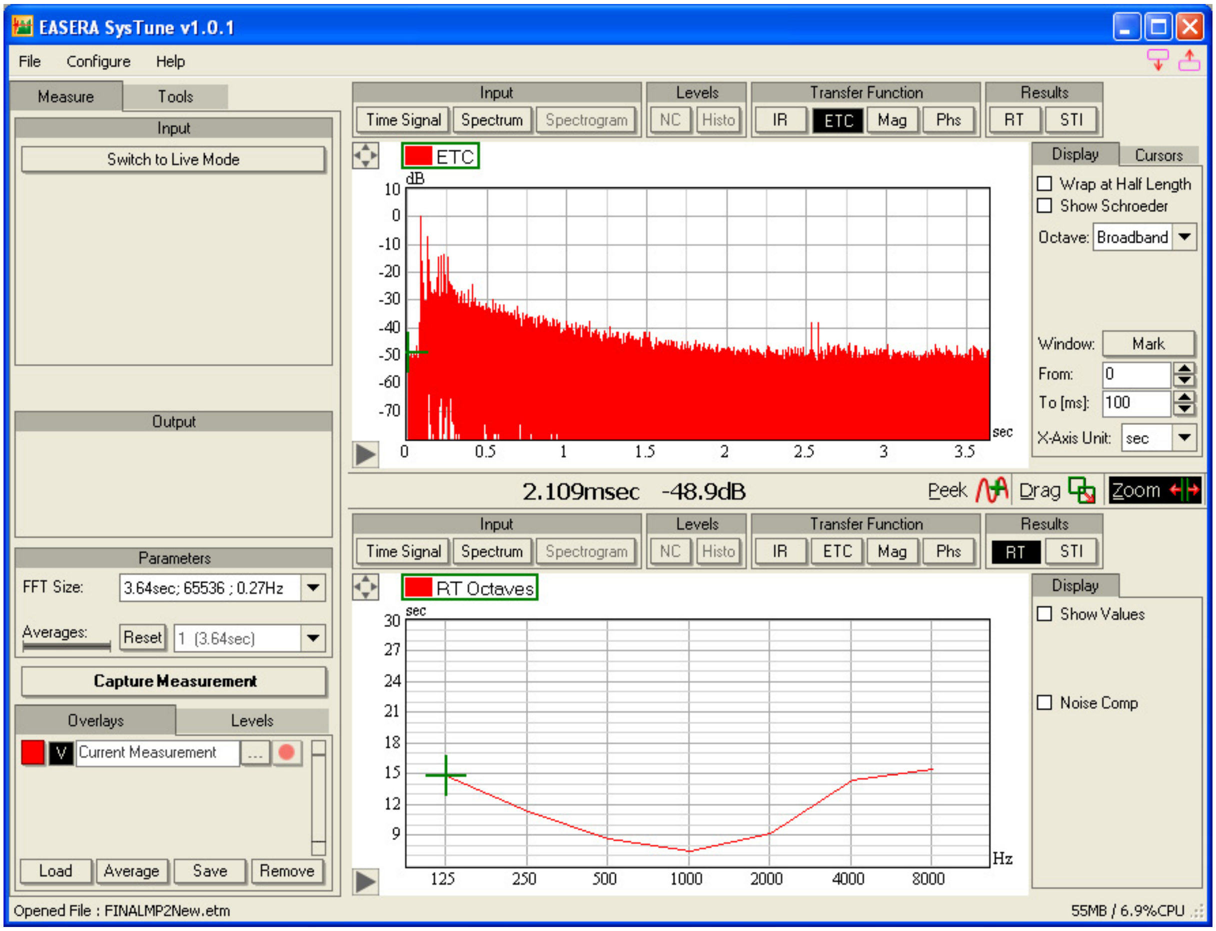

At first glance the RT curve does not seem right. It shows very high RT values (up to 15 seconds!) and in addition it is not decreasing for higher frequencies but increasing. The reason is fairly simple, if you look at the ETC. In the graph you can see that a large part of the log-squared impulse response consists of noise. We have already spoken about that a bit earlier. All of this noise is erroneously included in the RT calculation and therefore influences the results.

To understand better what is happening, enable Show Schroeder in the ETC view. This will add the curve for the Schroeder backward integral to the graph, which is nothing but the decay curve for the overall sound energy in the room as calculated from the impulse response.

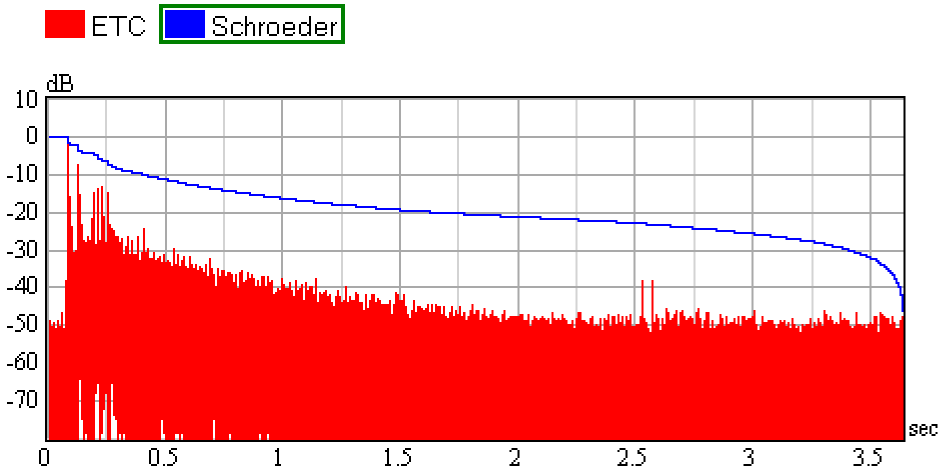

Here it becomes obvious that the integral contains all of the noise as well. To reduce the effect of noise in the RT calculation, there is a function called NOISE COMP in SysTune. Enable the noise compensation with the NOISE COMP switch in the DISPLAY panel of the ETC graph. After that the Schroeder curve will look much better.

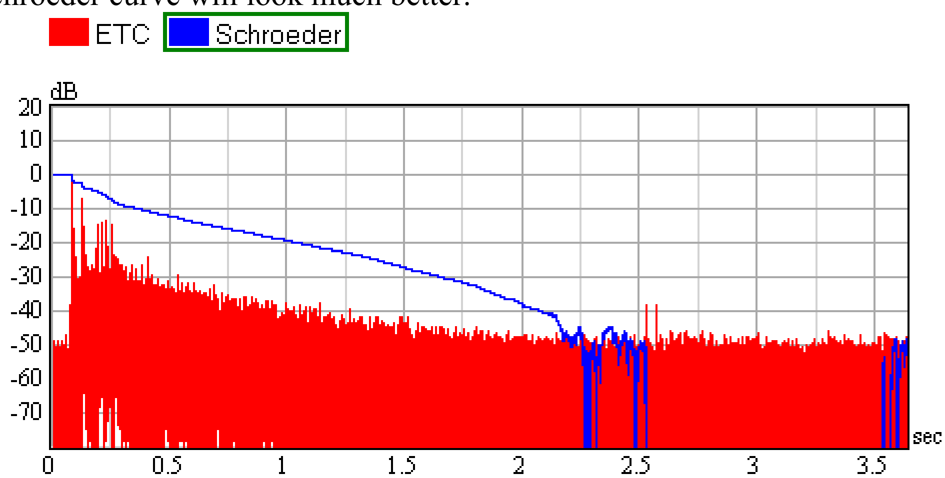

In the late part of it, for higher times, it should show gaps as shown above where the noise is subtracted correctly from the original integral. But if the curve is still too flat (as first shown above, without noise compensation) or too much energy is removed (the gaps occur very early in time), then you will need an advanced software tool to calculate the RT manually. To verify the noise prediction, you can also compare the noise value computed by the software with the ETC graph. It is shown directly below the switch NOISE COMP.

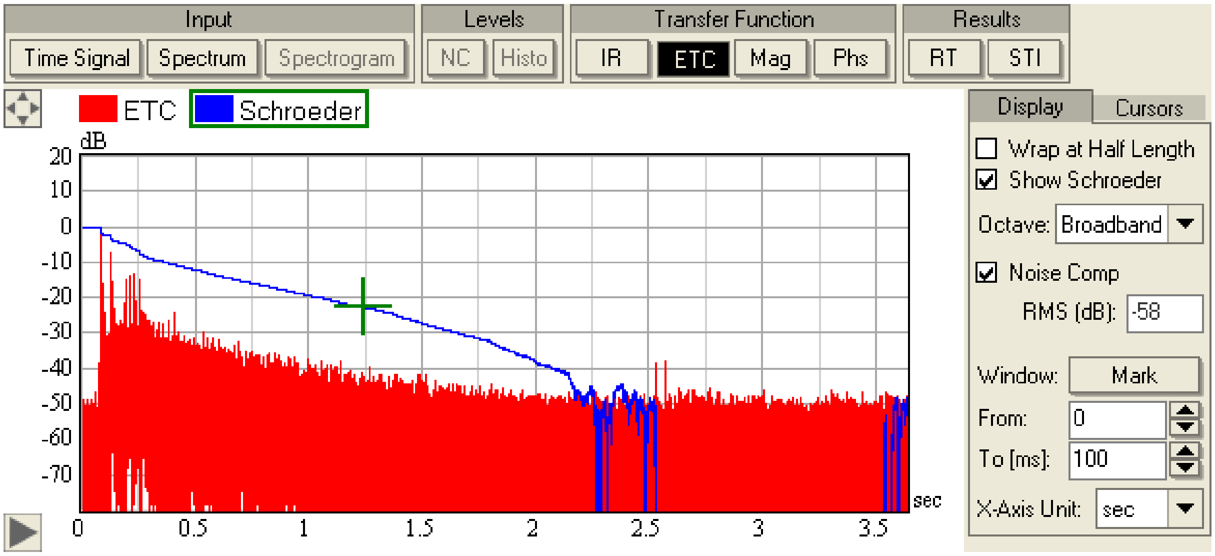

The value indicated here is –58 dB RMS which seems to coincide with our visual impression (remember that we are looking at the RMS, not the peak). At the same time, when you activate the noise compensation for the `ETC`, it is activated for the `RT` graph as well. However, it is not switched on by default, because the noise estimation routine requires quite some performance which is not always desirable especially in the real-time mode.

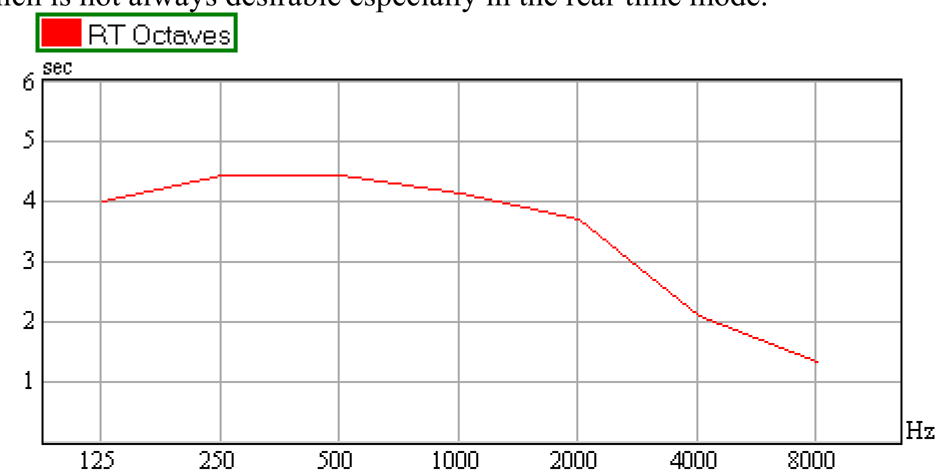

The resulting RT curve looks correct, finally. We can add the actual RT values to the graph by activating the switch SHOW VALUES in the DISPLAY panel on the right.

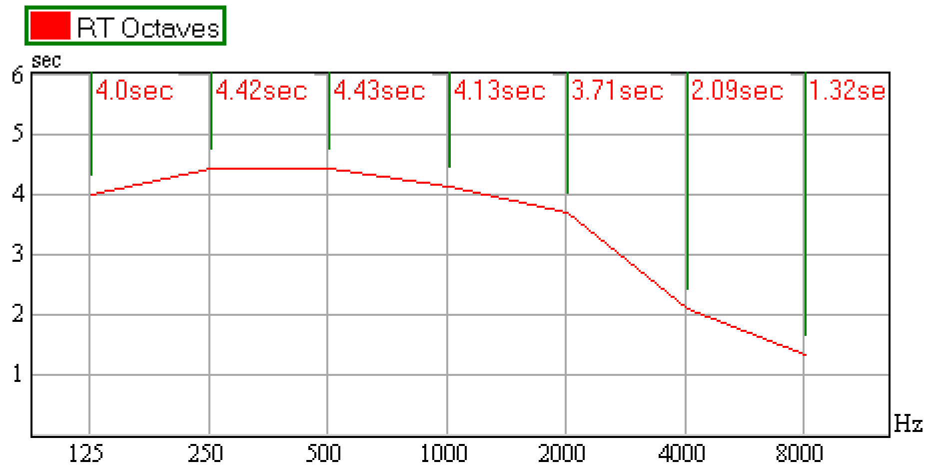

Hint: If you see very large values in the RT curve which do not seem right, use the ETC plot to investigate the reason in more detail. You can view the Schroeder curve for each octave band individually if you select the corresponding octave filter. Make sure that the noise compensation in the Schroeder plot yields the right values. Otherwise try to make a measurement with better signal-to-noise ratio.

If you would like to export the data to a text file you can do that like for any other graph. Use FILE|EXPORT DATA AS TEXT and FROM LOWER GRAPH if you would like to export data for the bottom graph or use FILE|EXPORT DATA AS TEXT and FROM UPPER GRAPH for the top graph.

Reverberation times depend largely on the volume of the room and the absorption materials on the walls, floor and ceiling. Typical reverberation times in the mid-frequency range are:

- Recording studios, small rooms ~ 0.3-1 second
- Multi-purpose halls ~ 0.8-1.7 seconds
- Theatres, concert halls ~ 1.2-2.2 seconds
- Large stadiums with roof ~ 3-5 seconds
- Cathedrals ~ 4-12 seconds

In this respect we would also like to emphasize that the reverberation time also depends on the occupancy of the venue. Often, a full stadium has a much reduced RT compared to the same but empty stadium. This is where SysTune becomes an invaluable tool, because you can measure the reverberation time with a dual-FFT setup, but instead of playing annoying stimulus signals, typical program material including trailers, advertisements and music can be used.

Until the introduction of SysTune, it was rather difficult to achieve this. Either one would have to make use of expensive, dedicated hardware or employ an acoustic simulation software, like EASE, to extrapolate the performance of the venue from the empty state, which was measured, to the occupied state, which is unknown.

Speech Intelligibility

Speech intelligibility measures describe how the system under test transmits human speech. Since the perception and recognition of words and whole sentences is a complex task, that involves psycho-acoustic issues as well as subjective effects, one cannot simply derive a formula that provides accurate results for all applications. However, one particular approach has turned out to be able to deliver results that correspond reasonably well with a broad range of listening tests. According to the standard IEC 60268-16 the speech transmission index (STI) defines intelligibility as a number on a scale from 0 to 1, where higher values are associated with better intelligibility.

Tech-Note
The principle concept of the STI makes the assumption that speech can be considered as a signal with very low modulation frequencies. In consequence, the STI is derived from the impulse response by evaluating how well the system under test can transmit modulation frequencies in the range of about 1 to 12 Hz. Accordingly, a set of 14 discrete modulation transfer function (MTF) values is computed for each of 7 octave bands. The result is a set of 7 modulation transfer indices (MTI), where each index quantifies the modulation properties of the system under test for a particular octave band. The final, single-valued STI number is then derived from these 7 coefficients as a weighted average.

As described above, the implementation of STI in SysTune follows IEC 60268-16 directly. However, it does not include the correction factors for signal masking and noise levels. Therefore it can only be considered as a first approximation. A detailed STI calculation for critical situations will have to be performed with a dedicated processing tool, like EASERA.

To switch to the STI graph, press the button STI from the RESULTS group of buttons. Like the RT before, we have used the example file FinalMP2New.etm here as well. The STI graph for this file will look like this:

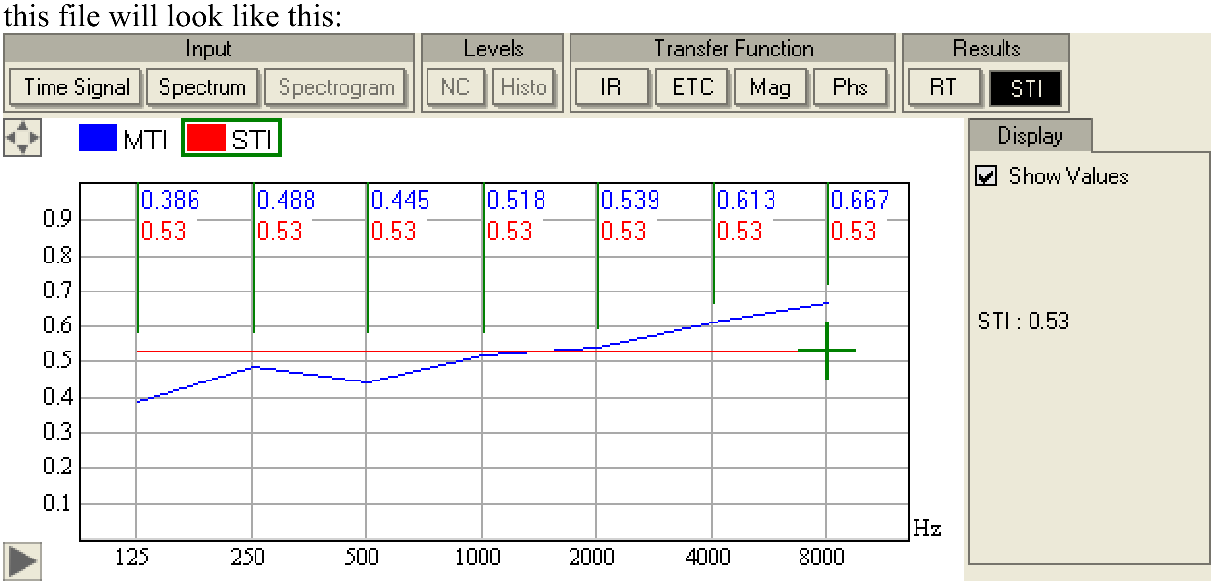

In this picture we have already activated the switch SHOW VALUES from the DISPLAY panel. This enables the numerical display of values in the graph. We can see that for this measuring location, the STI of 0.53 - shown in the DISPLAY panel - is acceptable. Looking at the blue MTI curve, the higher octave bands seem to contribute to a better intelligibility rather than the lower bands. In the above picture, the red horizontal line denotes the STI as the weighted average of the MTI values.

The rating for STI values according to the standard is:

- 0.00 – 0.30 unintelligible
- 0.30 – 0.45 poor
- 0.45 – 0.60 fair
- 0.60 – 0.75 good
- 0.75 – 1.00 excellent

Note that, this scale is given by the standard, but in the same document STI values greater than 0.5 are considered as sufficient, too. For more details about STI and its background please refer to the standard IEC 60268-16 as well as to the publications referenced by it.

Hint: Please also see above literature if you are interested in how to improve speech intelligibility in a venue. Like the tuning of a sound reinforcement system this is a complex task. Generally, well-aimed and time-aligned loudspeakers and enough absorption to reduce echoes and reverberation will increase the STI. There are acoustic simulation software packages available, like EASE, that allow you to predict the performance of a sound system in a venue.

Like any other data set, this graph can also be exported via the `FILE|EXPORT DATA AS TEXT` menu.

Summary

Reverberation time and speech intelligibility are primary measures to characterize rooms acoustically. We have learned how to calculate and display them in SysTune. It is important to remember that these measures are sensitive to noise and other measuring problems. Results in SysTune may only be an approximation and the analysis should be extended with more dedicated platforms as needed.

5.4 Measurements using Speech and Music Signals

In chapter 4 we talked in great detail about the measurement of impulse responses and transfer functions. In particular, in the beginning of that chapter we introduced two different setups, one setup was using the internal signal as a reference and the other one was using an external reference channel. All of the following exercises and functions were fully applicable to both configurations, because the quality of the measuring results does not depend on whether the reference signal is internal or external.

However, we should still discuss a little bit about the limitations for speech and music signals compared to dedicated excitation signals like a sweep or pink noise. We already touched that topic slightly when we introduced coherency and IR stability as measures to quantify the linearity and time-constancy of the system under test. In fact, how much the system deviates from these ideal conditions is not only a matter of the system itself but often also depends on the stimulus signal, its characteristics, its frequency distribution and its level. You can imagine that simple, random noise at a fixed level affects the measurement much more when a low level for the excitation signal is used than with a high level. A better signal-to-noise ratio will provide a more accurate impulse response and transfer function.

On this basis, we have to look at all uncommon measurement signals like speech and music. With respect to typical program material, time variance and harmonic distortion in the system under test are less of a concern. Rather, the signal must provide sufficient S/N over the full frequency range of interest. To obtain a broadband impulse response, a broadband excitation signal is needed. Signals that are band-limited, like speech, or prefer a number of tones, like baroque violin music, can only yield a part of the full transfer function. Other types of music, like rock music or classical music with a full orchestra, suffer less under band limitations.

As an example the next pictures show the time and frequency views of a typical piece of pop music. Note the dense signal distribution in both domains.

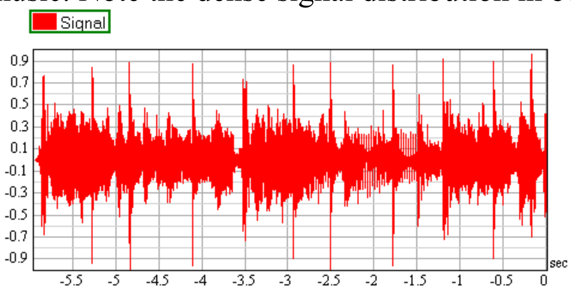

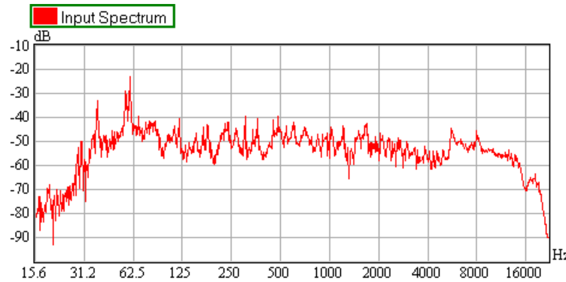

The next example shows the time and frequency data for a baroque dance. Note the systematic peaks and dips in the frequency domain due to the tonal character of the piece.

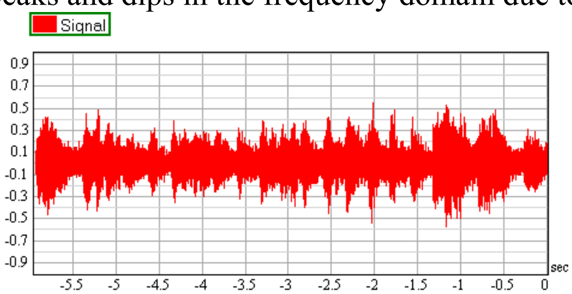

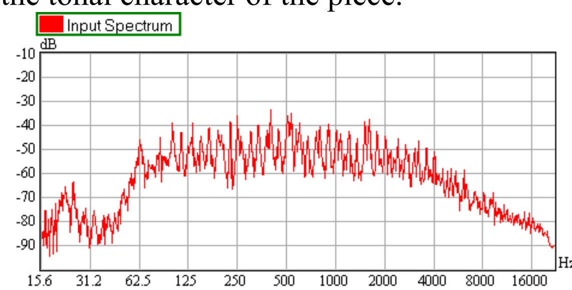

The last example is a speech sample. It shows gaps in both the time and frequency domain and will thus be most sensitive to any kind of noise (the additional peak at 8 kHz is an effect owed to the conversion from the original recording).

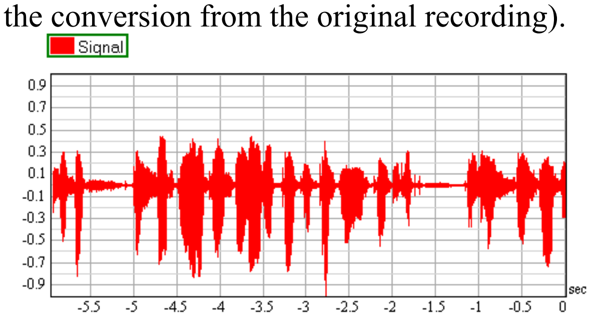

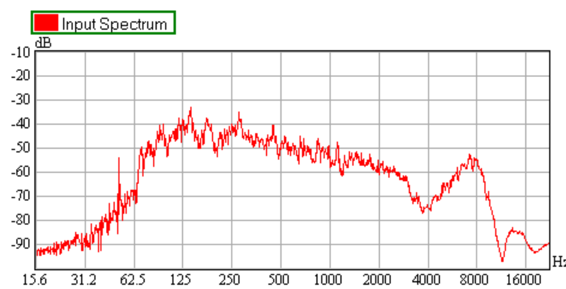

To increase the signal-to-noise ratio one can either make the signal louder or suppress the noise. The latter is difficult to achieve directly, but time-averaging helps to reduce the noise included with the measurement. Every doubling of the measurement duration reduces the noise floor by 3 dB, if the noise is randomly distributed. In consequence, longer measuring times can extend the band limits of a given signal and can increase the overall S/N as well.

Another point of concern when using arbitrary reference signals is the continuity of the signal contents. Speech and music material, like trailers, advertisements etc., often contain pauses of several seconds (see last of the examples above). During this time, there is only noise at the input and information about the system under test cannot be acquired. When running long time-averaged measurements such pauses should be excluded whenever possible.

SysTune offers two functions to remove noise from the gathered impulse response data. On the one hand, you can set a user-defined band-pass to exclude frequencies outside of the defined bandwidth of the excitation signal. On the other hand, a threshold function can be enabled to exclude data blocks below a certain level from the IR and TF calculations. We will show how to use these functions in SysTune in the following section.

Finally, it should be emphasized that the quality of the measurements also depends on the kind of source that is used as a reference. A signal that is supplied electronically and comes directly from the mixing console or control room is mostly free of additional disturbances. But a signal that is received from a microphone and supplies an acoustic reference will suffer more from environmental effects and will require more care. It is an especially difficult task to use a rehearsing choir or playing orchestra to measure the impulse response of a concert hall. So far, the experience and knowledge published in this field is very limited. The development team of SysTune has given some papers on this very topic; please refer to them for more information (see for example Ahnert, Feistel, Miron, Finder: Software-Based Live Sound Measurements Part 2, presented at the 123 AES convention, 2007 October 5-8, New York, NY, USA, preprint 7304).

Noise Suppression Tools

The place to enable noise suppression functions in SysTune is the TOOLS tab of the control panel on the left.

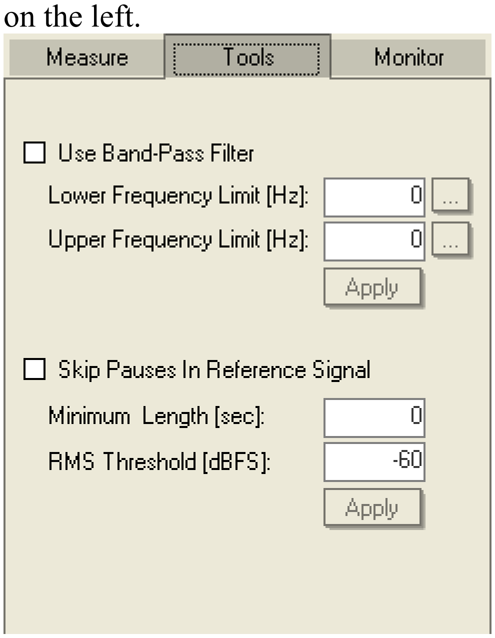

The first function, labeled USE BAND-PASS FILTER, is a user-defined band-pass filter that can be applied to the transfer function after deconvolution. Using the LOWER FREQUENCY LIMIT and UPPER FREQUENCY LIMIT text fields you can either enter the frequencies directly or press the respective ... button to open the frequency selection window. Any changes will only take effect after you have pressed APPLY. Note, that the filter will only be considered for impulse response and transfer function computations. The calculation of RT and STI will remain unaffected, since these use applied octave-band filters in any case.

The second tool called SKIP PAUSES IN REFERENCE SIGNAL activates a gating function. It will only allow those FFT frames to enter the deconvolution which exceed a certain threshold. If the reference signal is below the threshold level defined by RMS THRESHOLD for a time longer than the value specified for MINIMUM LENGTH, the current FFT block will be discarded. When calibrated, the threshold can also be given in a physical unit. Press APPLY to confirm your changes.

6. Additional Topics

This chapter should serve as an overview and as a reference of commonly used items. It also explains some elements of the graphic user interface that have not been covered in detail by the foregoing tutorial chapters.

6.1 Menu Structure

The FILE menu consists of the following entries:

- OPEN AUDIO FILE allows you to load another measurement in a variety of file formats and switches the program to the File Mode.
- SAVE TO AUDIO FILE saves the CURRENT MEASUREMENT or any of the OVERLAYS to a file. Use the submenu TIME SIGNAL or IMPULSE RESPONSE to save the corresponding measurement data to the file. Note, that IMPULSE RESPONSE data cannot be saved, when there is no REFERENCE CHANNEL.
- SEND [UPPER/LOWER] PICTURE TO saves the current graph to a graphics FILE or to the CLIPBOARD. In split view, when there is a top and bottom graph, this menu will be split into separate UPPER and LOWER menu items as well.
- EXPORT DATA AS TEXT|FROM [UPPER/LOWER] GRAPH saves the data of all overlaid curves of the current graph to a text file. In split view, when there is a top and bottom graph, this menu will be split into separate UPPER and LOWER menu items as well.
- EXPORT DATA AS TEXT|FREQUENCY RESPONSE saves the INPUT SPECTRUM or the TRANSFER FUNCTION of the CURRENT MEASUREMENT or of any of the OVERLAYS to a text file. We have used and explained this function in chapter 4.
- OPTIONS opens the OPTIONS window to adjust general display, performance and processing parameters of SysTune.
- OPEN SETUP FILE loads a SysTune configuration file. This file contains all general measurement and display settings in an XML format.
- SAVE SETUP FILE stores the current configuration of SysTune in a file that can be archived, transferred and loaded again. This file contains all general measurement and display settings in an XML format.
- EXIT quits the program.

The CONFIGURE menu provides the following functions:

- SELECT DEVICE ... opens the SELECT DEVICE window to select audio drivers for input and output as well as the sample rate. Please see chapter 5.2 for more details.
- COLOR SCHEME|WHITE ON BLACK switches SysTune to a display with bright colors on a dark background. This is a convenient setting for measurements in a dark environment.
- COLOR SCHEME|SYSTEM COLORS switches SysTune to a display with dark colors on a bright background. This is a convenient setting for measurements in daylight or an office environment.

The HELP menu contains the following options:

- HELP opens the EASERA SysTune help file, which is this document.

- CREATE STATUS REPORT – use this command if you have technical problems with the software. It will collect the system and license information that is needed to provide you with technical support. You will also be prompted to send an email right away, if you like.
- EASERA SYSTUNE WEBSITE takes you to the program's website, namely www.EASERASysTune.com. The latest software updates as well as other useful information are located here.
- ABOUT shows the ABOUT window of SysTune, it is useful if you need to know the exact version number of the program or the license that is currently installed.

6.2 Short Cuts

This list of short cuts implemented in SysTune should help you to access important functions quickly:

Ctrl+O	Opens an audio file, see menu command FILE|OPEN AUDIO FILE
Ctrl+A	Saves data to an audio file, see menu command FILE|SAVE TO AUDIO FILE
F1	Opens the help file, see menu command HELP|HELP
F4	Opens the PROPERTIES window for the selected overlay
F5	Starts the real-time analysis, equivalent to the button START ANALYSIS
F6	Starts playing the test signal, equivalent to the button PLAY SIGNAL
F7	Stops playing the test signal, equivalent to the button STOP SIGNAL
F10	Stops the real-time analysis, equivalent to the button STOP ANALYSIS
Space	Captures an overlay immediately, equivalent to the button CAPTURE MEASUREMENT
+ and -	Switches the active overlay of the current graph
Alt+Z	Switches the mouse mode permanently to ZOOM
Ctrl+Alt	Switches the mouse mode temporarily to ZOOM
Alt	Switches the mouse mode temporarily to EXACT ZOOM
Ctrl	Switches the mouse mode temporarily to MARK WINDOW
Alt+P	Switches the mouse mode permanently to PEEK
Shift	Switches the mouse mode temporarily to PEEK
Alt+D	Switches the mouse mode permanently to DRAG
Ctrl+Shift	Switches the mouse mode temporarily to DRAG
Right and Left Arrow Keys	Switches the REFERENCE CURSOR

6.3 View Limits

The view limits of any graph can be adjusted using the ZOOM and DRAG mouse modes. In ZOOM mouse mode, drag with the left mouse button pressed to zoom along the horizontal axis or drag the right mouse button to zoom along the vertical axis. In DRAG mouse mode, drag with the left mouse button to shift the graph area.

To return to the full view (auto-scale) double click in the graph or left click on the auto-scale button in the upper left corner of the graph.

Alternatively, a new range for the horizontal or vertical axis can be entered in the view limits section of the graph. To open it press the small triangle button in the lower left corner of the graph.

In this area you may enter the start and stop coordinate for the X-Axis as well as the start and stop coordinate for the Y-Axis. The view limits section looks like this for time domain graphs:

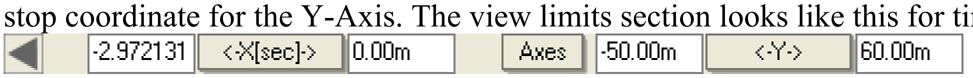

In the frequency domain, the view limits section includes two additional buttons, labeled ....

These buttons allow you to pick a value directly from a list of common frequencies.

If you would like to apply a partial auto-scale press the button <-X-> for the horizontal axis or <-Y-> for the vertical axis and the graph will be zoomed to include all data points along that axis.

Press the AXES button to open the AXES SETTINGS window for more options to scale the current graph. For example, in the frequency domain the window will be similar to the one in the picture below:

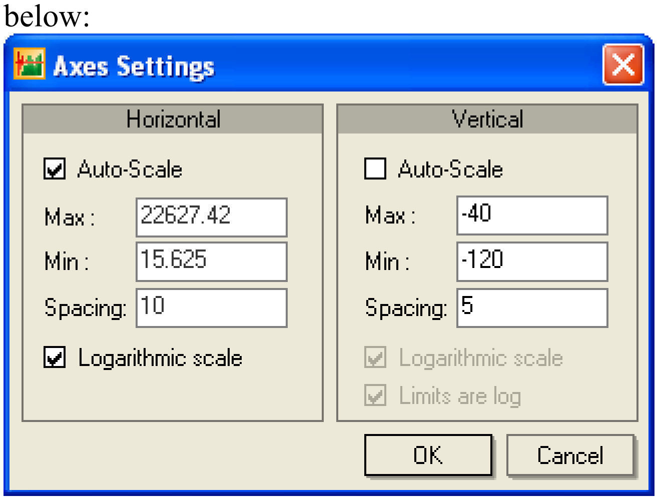

For both axes, HORIZONTAL and VERTICAL, you can disable AUTO-SCALE and enter user-defined view limits. MAX and MIN define the upper and lower limit for the respective axis. Using the entry for SPACING you can set the distance between adjacent lines of the grid. If LOGARITHMIC SCALE is switched off the corresponding axis will be displayed with linear spacing otherwise it will be logarithmic. Use the option LIMITS ARE LOG to define the view limits as logarithmic numbers even if a linear scale is used. Note, that not all options are available for all graphs.

In most cases, the software will pick the right axis settings for you automatically. If you enter new values in the view limits section as discussed above, AUTO-SCALE will be switched off automatically. It can be switched on for both axes simultaneously by double-clicking into the graph or for each axis individually using the <-X-> or <-Y-> button. Therefore, generally you will only need to use the AXES SETTINGS window to change the grid spacing or to toggle the axis scale between linear and logarithmic.

6.4 Trouble Shooting

Here are some additional hints for trouble shooting your software configuration if you encounter technical problems:

- If you cannot start EASERA SysTune, make sure you have correctly installed all of the files. Please refer to the installation instructions and verify that you have performed all of the installation steps properly.
- If SysTune worked before, but it does not start anymore or shows strange errors, do the following: From the Windows Start menu, start EASERA SysTune using the link PROGRAMS / AFMG / EASERA SYSTUNE (USE DEFAULT SETTINGS). This will discard the last application settings and start like a fresh installation.
- If SysTune worked before, but you tried out another soundcard or driver and now you receive some errors, do the following: From the Windows Start menu, start EASERA SysTune using the link PROGRAMS / AFMG / EASERA SYSTUNE (USE DEFAULT AUDIO). This will let SysTune use the Windows default audio drivers, which should always work.
- If you cannot capture any signal with the input or if you cannot play any signal to the output, make sure that the soundcard works otherwise. Use the Windows Media Player and Windows Sound Recorder to ensure the basic functionality. Make sure that you have the latest version of the soundcard driver installed.
- If you see strange measurements in any of the graphs, please also see the trouble shooting tips in the impulse response chapter 4.2. Make sure you have switched off all monitoring functions of your soundcard and other mixing abilities. To measure a stable impulse response you will need to have synchronized input and output clocks as well.
- Visit the website of your distributor and our public forum located at www.afmg-network.com. Also verify that you have installed the latest version of the software. You will find information about updates and other useful tools under www.EASERASysTune.com.

If you still cannot make EASERA SysTune run, please create a status report: From the Windows Start menu select the menu item PROGRAMS / AFMG / EASERA SYSTUNE / CREATE STATUS REPORT. After that, please contact your software distributor and send him this report along with a detailed description of the error you have encountered.

Graph Reference

Graphs [All]

Buttons

The `Graphs` have a set of buttons that can be used to set the limits for both axes. Select the button to display the buttons as shown in the next picture. Select to hide the buttons. Here are some typical examples.

Time:

Distance:

Frequency:

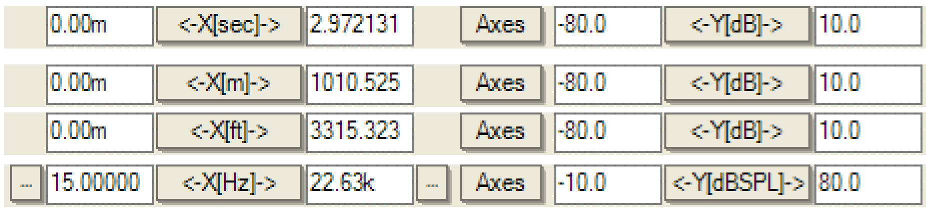

- `Axes:` Opens a dialog box to set the X and Y axis limits.
- `<-X[sec]->, <-X[m]->, <-X[ft]->, <-X[Hz]->:` Zooms the display to show the entire X axis. X axis values can be directly entered into the boxes to the left and right of this button. Click the `[…]` buttons to open a dialog box that allows you to select ISO standard frequencies for the values. Time Domain units are displayed as time (sec) or distance (m or ft) and Frequency Domain units are displayed as frequency (Hz).
- `<-Y[dB]->, <-Y[dBSPL]->:` Zooms the display to show the entire Y axis. Y axis values can be directly entered into the boxes to the left and right of this button. The units displayed are based on the specific graph that is currently displayed, see below.

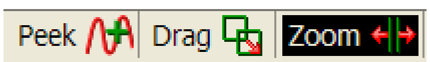

- `Peek:` Enables the mouse mode to track the values in the graph, see below.
- `Drag:` Enables the mouse mode to drag the graph, see below.
- `Zoom:` Enables the mouse mode to zoom into the graph, see below.

Mouse

- `Peek:` Click and Drag the Left button to show the difference in time or frequency and level between any two points on the current curve. Click the Right button to insert a cursor. Use the + and - keys to toggle the cursor from one curve to another. 'Left'/'Right' keys change the reference value for the X-value display. The text is shown in bold for both. Select the `Cursors` tab in the right hand panel to further manipulate the cursors or to remove them.
- `Drag:` Click and Drag with the Right or Left button to move the graph. This will also set the Axis limits to fixed values if they were previously set to Auto-Scale.
- `Zoom:` Click and Drag with the Left button to select the area to zoom into in the X axis. Click and Drag with the Right button to select the area to zoom into in the Y axis. By default, the mouse zoom snaps to the tick lines. Hold the Alt key pressed while zooming to disable the snap function.
- `Full View:` Double-click in the picture or select the button in the upper left corner of the graph to switch the view to Auto-Scale (like <-X[sec]-> and <-Y[dB]-> combined). This will automatically show all of the data automatically scaled to fit in the graph.

Graphs [Input]

Time Signal

This graph shows the signal that appears at the selected input.

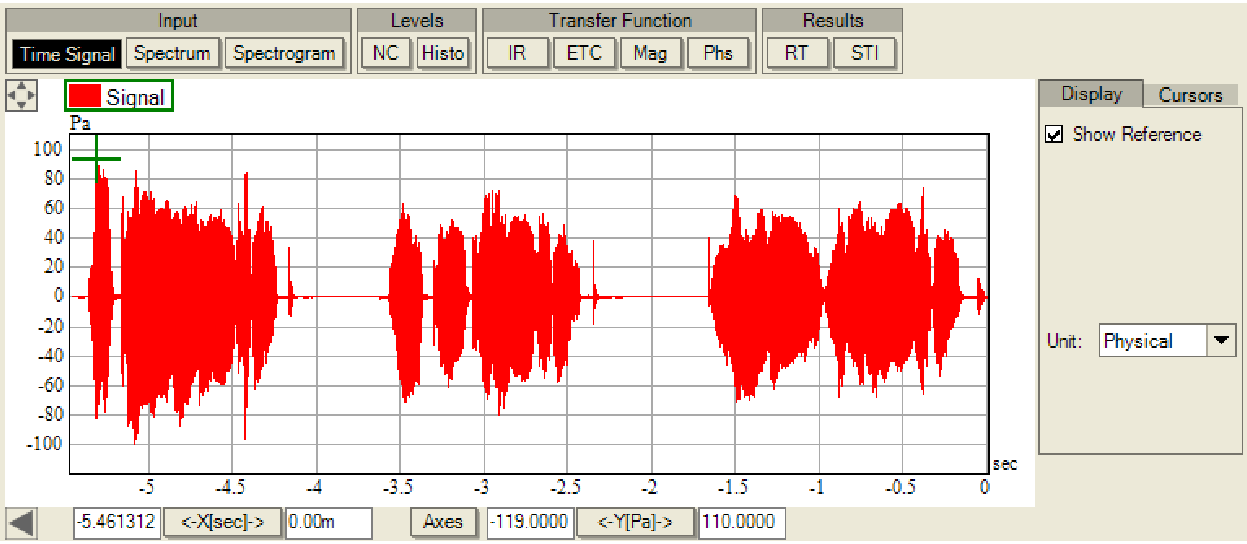

Parameters:

- `Show Reference`: Displays the Reference Channel in addition to the Input Channel.
- `Unit`: Default display is `Digital FS`. If the channel has been calibrated, select `Physical` to display the corresponding Pressure or Voltage.

Spectrum

This graph shows the spectrum of the signal that appears at the selected input.

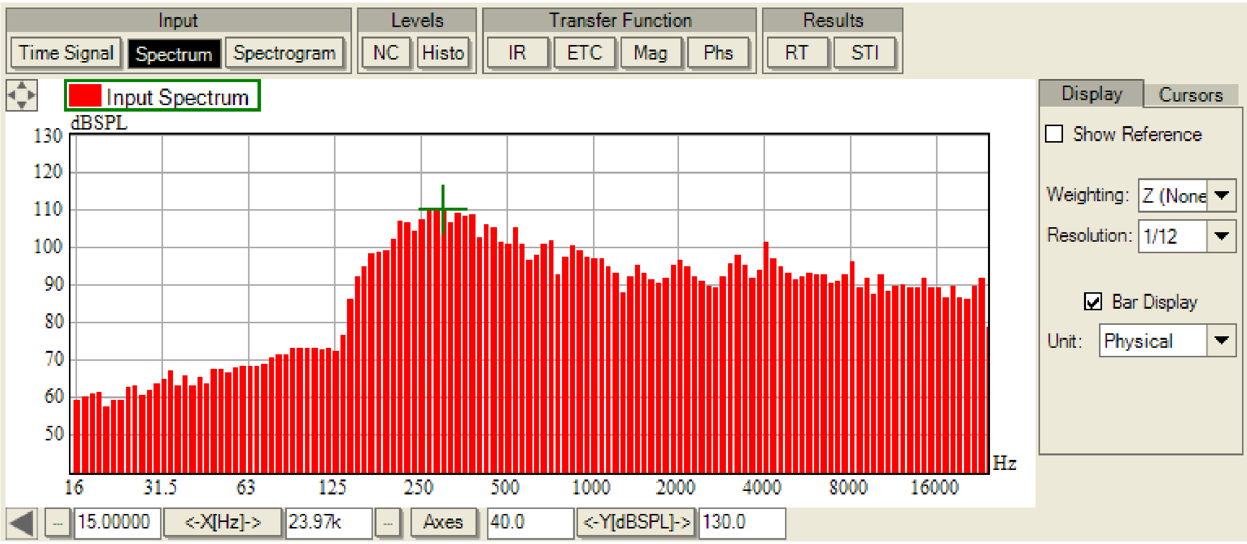

Parameters:

- `Show Reference:` Displays the Reference Channel in addition to the Input Channel.
- `Weighting:` Applies spectral weighting using the following choices according to ISO 61672: `Z (None), A, B, C`.
- `Resolution:` Displays using the following choices for the spectral resolution: `Full, 1/1, 1/3, 1/6, 1/12, 1/24, 1/48, 1/96`. `Full` displays the graph at the frequency resolution of the `FFT Size`. The others integrate the energy in fractional octave bands and then display the data at the selected resolution.
- `Bar Display:` Displays the bands using bars instead of a line.
- `Unit:` Default display is `Digital FS`. If the channel has been calibrated, select `Physical` to display the corresponding Pressure or Voltage.

Spectrogram

This graph shows the combined time signal and spectrum that appears at the selected input.

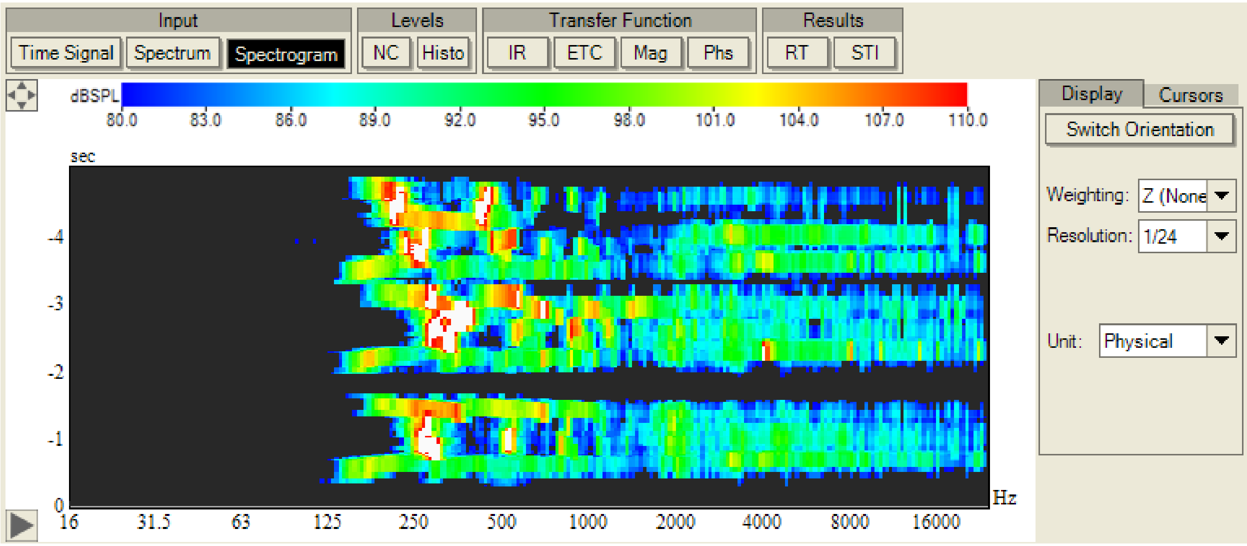

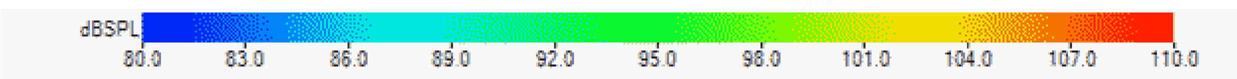

`Color map`: Double-click in the color map shown above to open the Color Map Settings window. In this dialog you can set the minimum and maximum values to be used to display the graph.

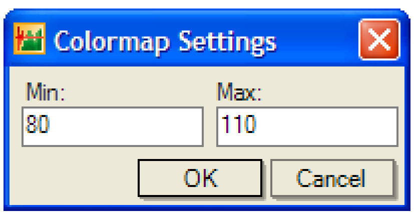

Parameters:

- `Switch Orientation`: By default, the Spectrogram scrolls from the bottom to the top of the graph. Select this button to have it scroll from the right to left side. Select the button again to switch to the original scroll direction.
- `Weighting`: Applies spectral weighting using the following choices according to ISO 61672: `Z (None), A, B, C`.
- `Resolution`: Displays using the following choices for the spectral resolution: `Full, 1/1, 1/3, 1/6, 1/12, 1/24, 1/48, 1/96`. `Full` displays the graph at the frequency resolution of the `FFT Size`. The others integrate the energy in fractional octave bands and then display the data at the selected resolution.
- `Unit`: Default display is `Digital FS`. If the channel has been calibrated, select `Physical` to display the corresponding Pressure or Voltage.

Graphs [Levels]

NC {Noise Criteria}

This graph shows the Noise Criteria (after calibration).

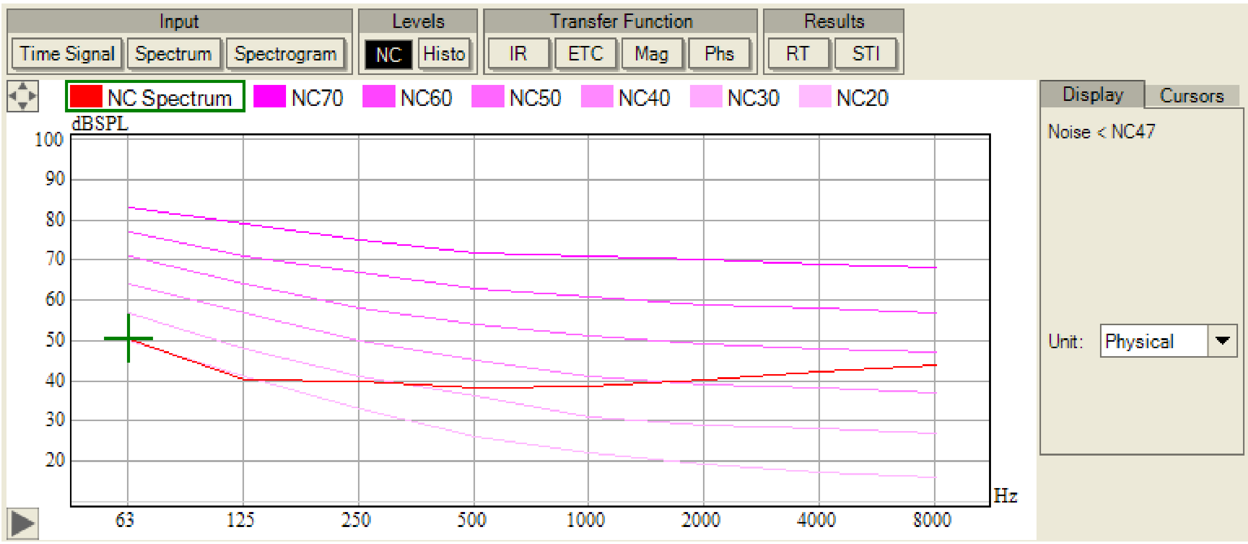

Parameters:

- `Noise < NCxx:` Indicates the lowest Noise Criteria curve that is not exceeded in any octave band interpolated to 1 dB increments. If the input channel has not been calibrated then this will be indicated here also.
- `Unit:` Default display is `Digital FS`. If the channel has been calibrated, `Physical` will display the corresponding Pressure or Voltage.

Histo {Histogram}

SPL {Sound Pressure Level}

This graph shows the `SPL` History.

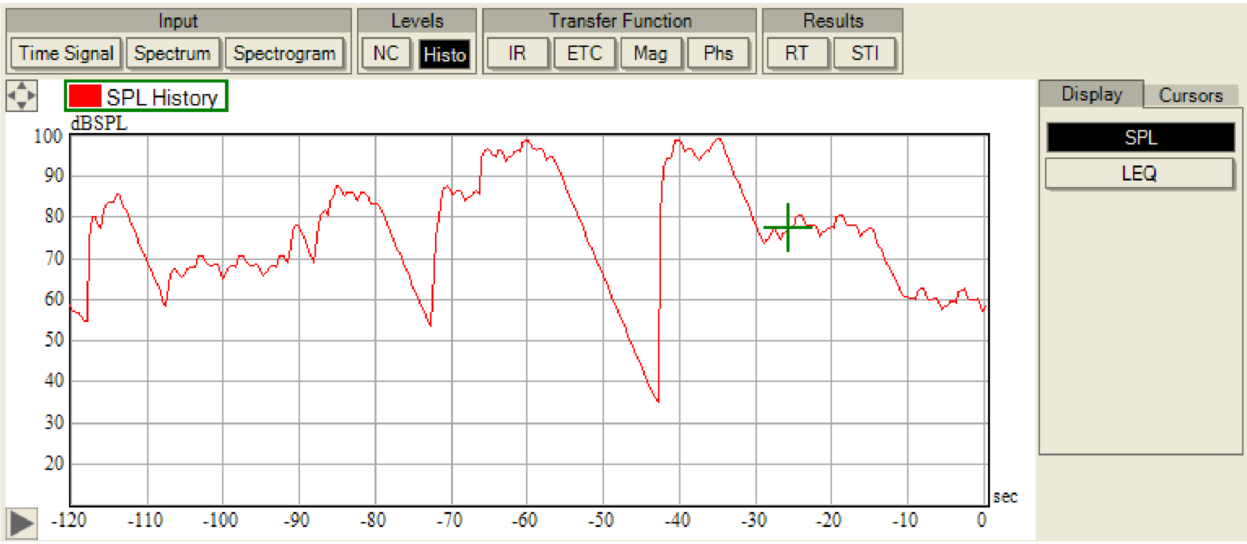

LEQ {Level Equivalent}

This graph shows the `LEQ` History.

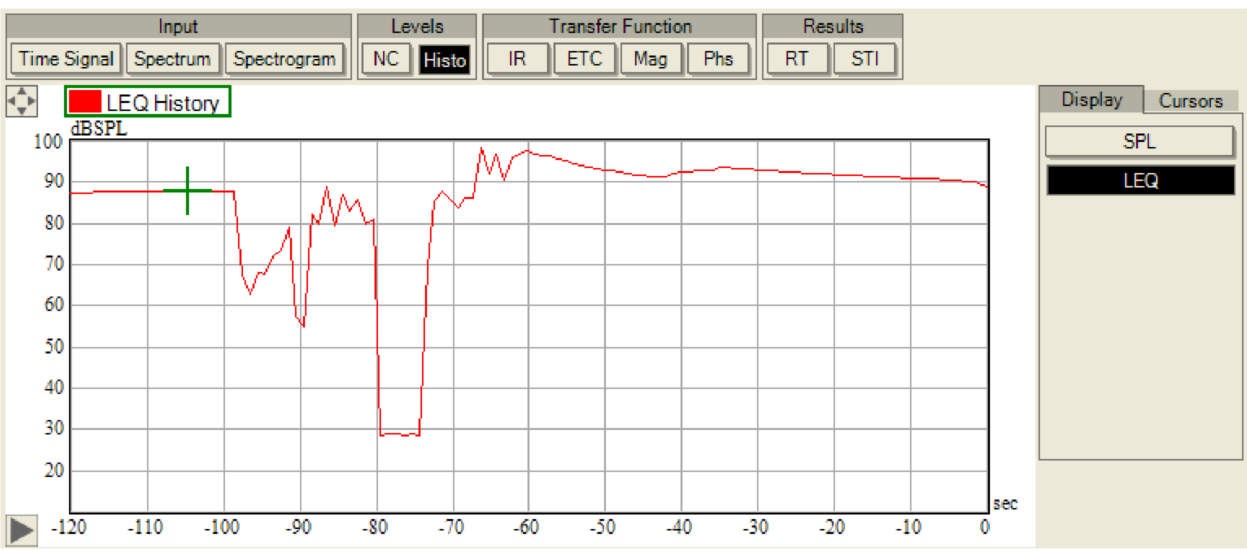

Parameters:

- `SPL:` Select the button to show the `SPL History`.
- `LEQ:` Select the button to show the `LEQ History`.

Graphs [Transfer Function]

Note: All transfer function measurements require an input signal and a reference signal.

IR {Impulse Response}

This graph shows the measured real-time Impulse Response.

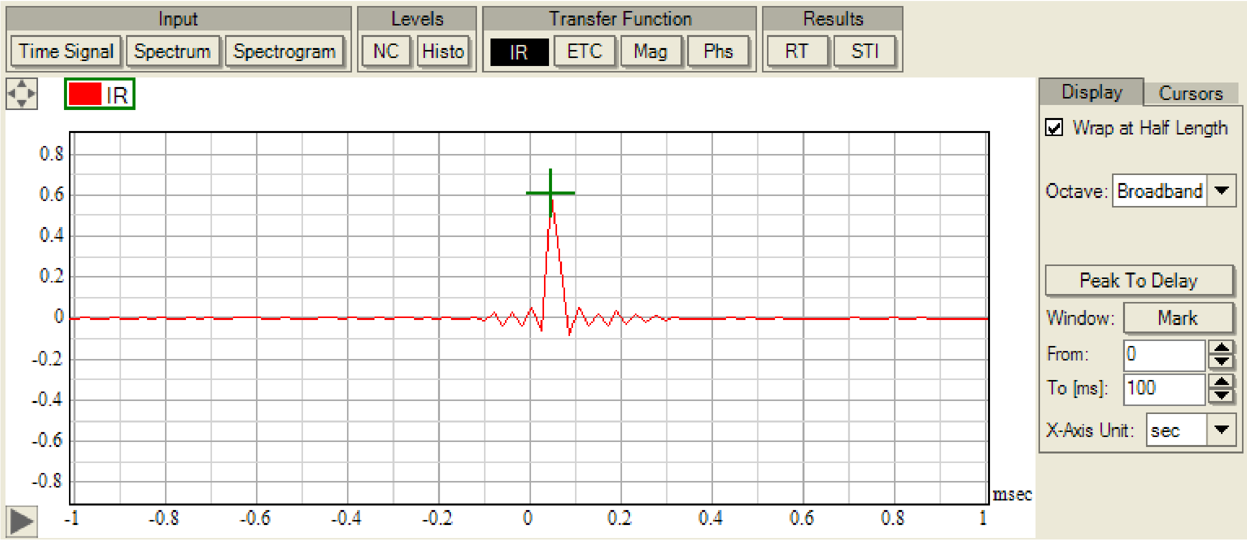

Parameters:

- `Wrap at Half Length:` Shows the end of the time record as negative values and places 0 time in the middle of the display.
- `Octave:` Applies filters to the IR using the following choices: `Broadband, 125, 250, 500, 1000, 2000, 4000, 8000`.
- `Peak to Delay:` Select the button to enter the arrival time of the signal maximum as the delay time for the selected channel. The graph will then shift the signal maximum to time zero.
- `Window:` Select the `Mark` button, then click with the Left Mouse Button to select the start time for the window. Click with the Right Mouse Button to select the stop time for the window. The selected times for the window start and stop can also be directly entered.
- `From:` Enter the start time for the `Window` or use the spin buttons to increment and decrement the value by the length of a sample.
- `To [ms]:` Enter the stop time for the `Window` or use the spin buttons to increment and decrement the value by the length of a sample. The current unit for both the start and stop time is shown in brackets.
- `X-Axis Unit:` Shows the time (`sec` for seconds) or distance (`m` for meters, `ft` for feet) for the current measurement at the cursor position. Click on the arrow to select from a list of alternate units for both cursor display and axis labels.

ETC {Energy Time Curve}

This graph shows the measured real-time Energy Time Curve (log-squared IR).

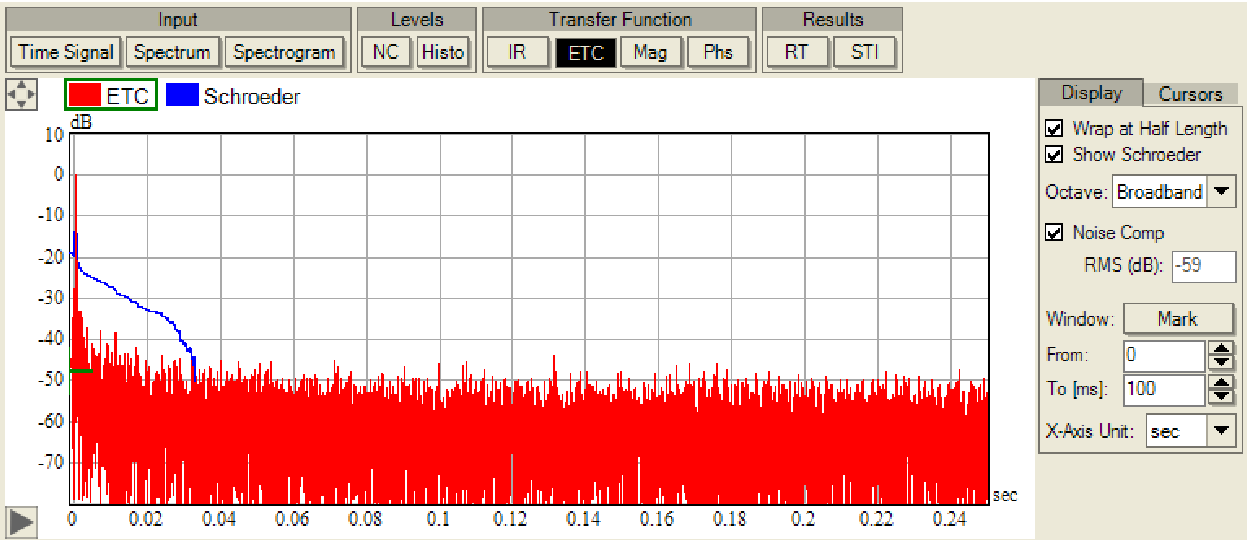

Parameters:

- `Wrap at Half Length:` Shows the end of the time record as negative values and places 0 time in the middle of the display.
- `Show Schroeder:` Displays the Schroeder backward-integration curve.
- `Octave:` Applies filters to the IR using the following choices: `Broadband, 125, 250, 500, 1000, 2000, 4000, 8000`.
- `Noise Comp:` Uses noise compensation to remove the noise from the IR. The `Schroeder` plot will show the effect of the noise compensation. The `RMS (dB)` box shows the calculated noise level.
- `Window:` Select the `Mark` button, then click with the Left Mouse Button to select the start time for the window. Click with the Right Mouse Button to select the stop time for the window. The selected times for the window start and stop can also be directly entered.
- `From:` Enter the start time for the `Window` or use the spin buttons to increment and decrement the value by the length of a sample.
- `To [ms]:` Enter the stop time for the `Window` or use the spin buttons to increment and decrement the value by the length of a sample. The current unit for both the start and stop time is shown in brackets.
- `X-Axis Unit:` Shows the time (`sec` for seconds) or distance (`m` for meters, `ft` for feet) for the current measurement at the cursor position. Click on the arrow to select from a list of alternate units for both cursor display and axis labels.

Mag {Magnitude}

This graph shows the full `Transfer Function Magnitude`.

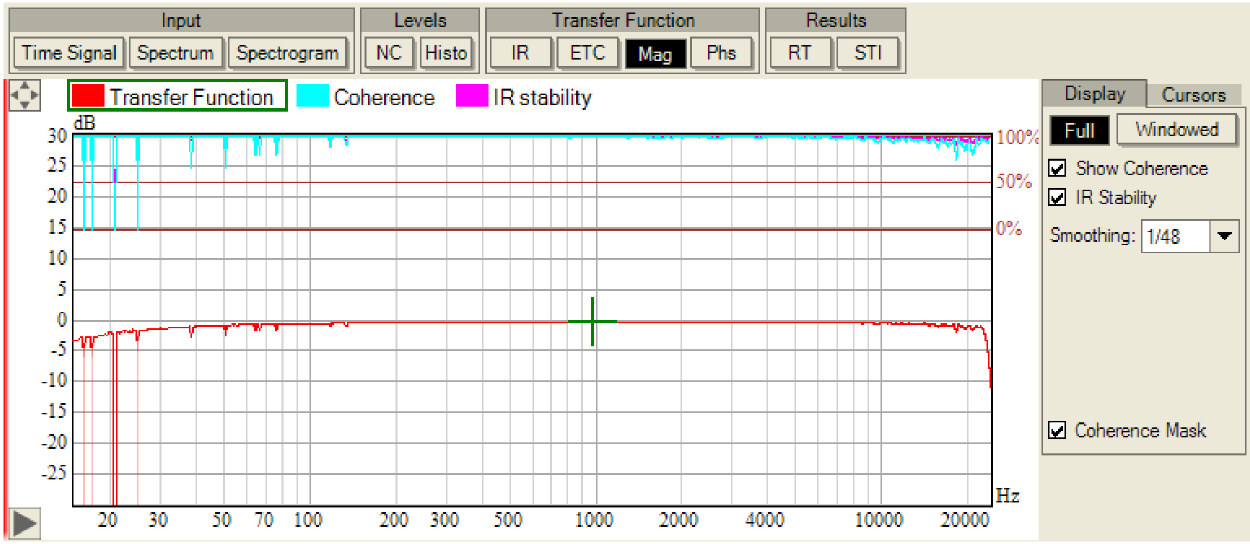

Parameters:

- `Full/Windowed`: By default, the `Full` view displays the spectrum of the entire (non-windowed) IR. The `Windowed` selection displays only the spectrum of the IR as derived with the currently selected window (Select the window location and size in the IR or ETC display).
- `Show Coherence`: Displays the coherence as a function of frequency which is a measure for the (linear) correlation of the signal input and the reference. Uses a graph overlay of the Coherence using 25% of the display (Select the percentage in Options).
- `IR Stability`: Displays the IR stability as a function of frequency which is a measure for the consistency (and repeatability) of the IR over time. Uses a graph overlay of the IR Stability using 25% of the display (Select the percentage in Options).
- `Smoothing`: Displays using the following choices: `Full, 1/1, 1/3, 1/6, 1/12, 1/24, 1/48, 1/96`. `Full` displays the graph at the frequency resolution of the `FFT Size` without any smoothing. The others smooth the energy over the fractional octave bands and then display the data.
- `From`: Enter the start time for the `Window` or use the spin buttons to increment and decrement the value by the length of a sample.
- `To [ms]`: Enter the stop time for the `Window` or use the spin buttons to increment and decrement the value by the length of a sample. The current unit for both the start and stop time is shown in brackets.
- `Coherence Mask`: Displays the magnitude or phase curve as a partially transparent curve dependent on the `Coherence mask threshold [%]` and `Masked transparency [%]` settings in `Options`.

Phs {Phase}

This graph shows the full `Transfer Function Phase`.

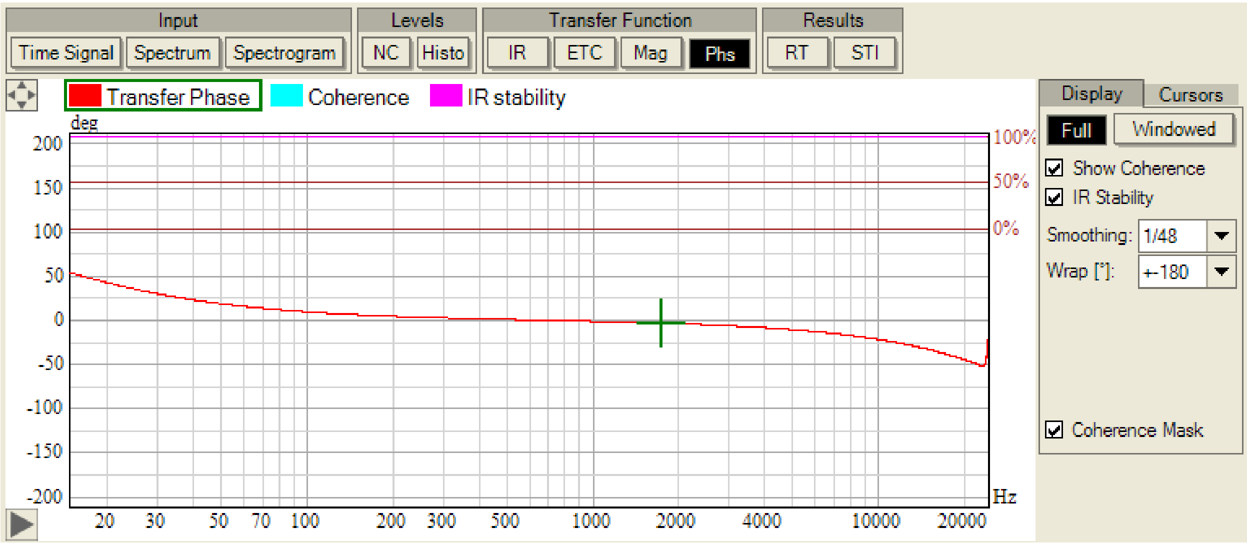

Parameters:

- `Full/Windowed`: By default, the `Full` view displays the spectrum of the entire (non-windowed) IR. The `Windowed` selection displays only the spectrum of the IR as derived with the currently selected window (Select the window location and size in the IR or ETC display).
- `Show Coherence`: Displays the coherence as a function of frequency which is a measure for the (linear) correlation of the signal input and the reference. Uses a graph overlay of the Coherence using 25% of the display (Select the percentage in Options).
- `IR Stability`: Displays the IR stability as a function of frequency which is a measure for the consistency (and repeatability) of the IR over time. Uses a graph overlay of the IR Stability using 25% of the display (Select the percentage in Options).
- `Smoothing`: Displays using the following choices: `Full, 1/1, 1/3, 1/6, 1/12, 1/24, 1/48, 1/96`. `Full` displays the graph at the frequency resolution of the `FFT Size` without any smoothing. The others smooth the energy over the fractional octave bands and then display the data.
- `Wrap [°]`: Displays using the following choices: `+-180, 0-360, +-360, +-540, +-720, None`. `None` displays the graph without any wrapping of the phase values. The others display the phase wrapped at 360 degree increments for the first two, and 720, 1080, or 1440 degrees respectively for the others.
- `From`: Enter the start time for the `Window` or use the spin buttons to increment and decrement the value by the length of a sample.
- `To [ms]`: Enter the stop time for the `Window` or use the spin buttons to increment and decrement the value by the length of a sample. The current unit for both the start and stop time is shown in brackets.
- `Coherence Mask`: Displays the magnitude or phase curve as a partially transparent curve dependent on the `Coherence mask threshold [%]` and `Masked transparency [%]` settings in `Options`.

Graphs [Results]

RT {Reverberation Time}

This graph shows the computed `RT`.

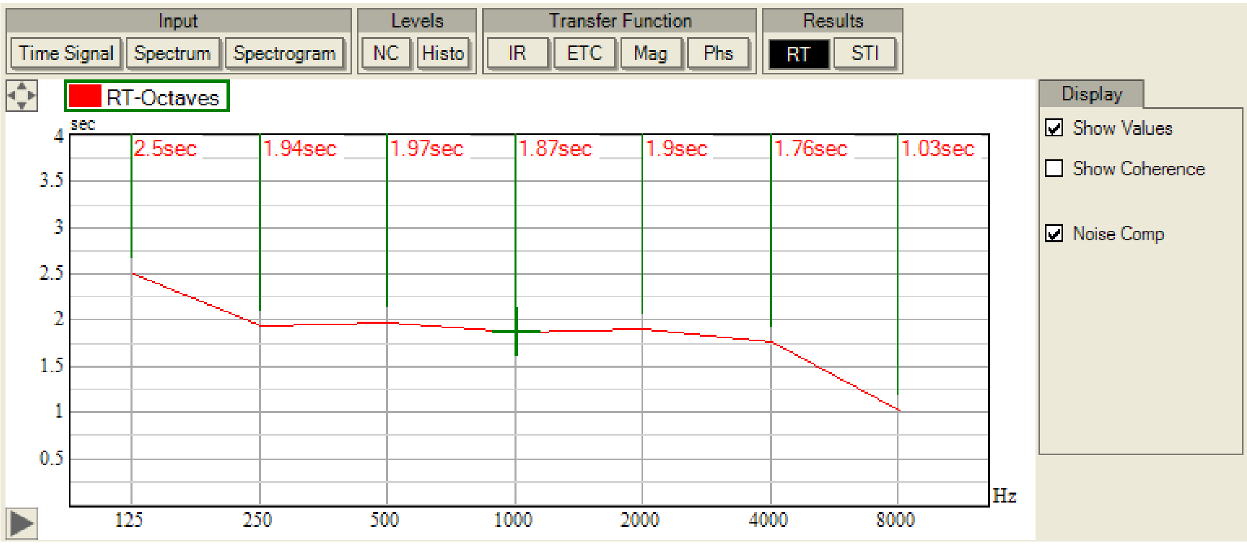

Parameters:

- `Show Values`: Displays a cursor at each of the octave bands showing the value for each of the displayed curves.
- `Show Coherence`: Displays the coherence as a function of frequency which is a measure for the (linear) correlation of signal input and the reference. Uses a graph overlay of the Coherence using 25% of the display (Select the percentage in Options).
- `Noise Comp`: Uses noise compensation to remove the noise from the IR. Note that also the ETC/Schroeder plot will show the effect of the noise compensation.

STI {Speech Transmission Index }

This graph shows the `STI` and `MTI (Modulation Transfer Index).`

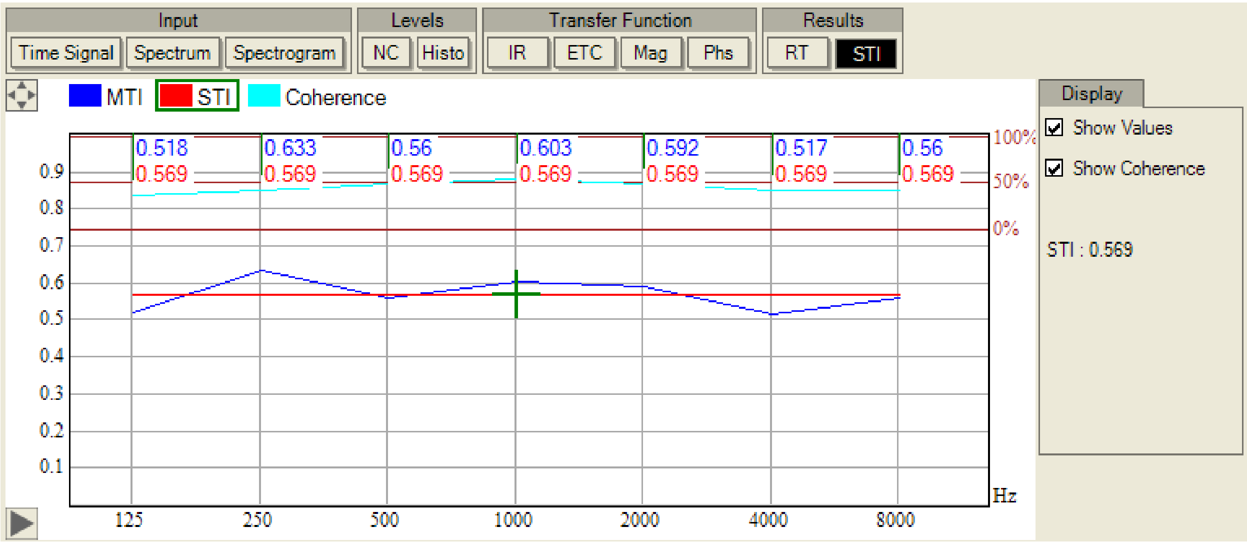

Parameters:

- `Show Values`: Displays a cursor at each of the octave bands showing the value for each of the displayed curves.
- `Show Coherence`: Displays the coherence as a function of frequency which is a measure for the (linear) correlation of signal input and the reference. Uses a graph overlay of the Coherence using 25% of the display (Select the percentage in Options).
- `STI`: Indicates the Speech Transmission Index for the active curve.