

SALIENT SCIENCES



VIDEOFOCUS USER GUIDE

Version 8.0

VideoFOCUS User Guide

Version 8.0

October 2018



Digital Audio Corporation
d/b/a Salient Sciences
1000 Centre Green Way Ste 200
Cary, NC 27513
Phone: +1 (919) 572 6767
Fax: +1 (919) 572 6786

sales@salientsciences.com
www.salientsciences.com

© 2001-2018 DIGITAL AUDIO CORPORATION d/b/a SALIENT SCIENCES.
ALL RIGHTS RESERVED.

Windows, Windows XP, Windows 7, and Windows 10 are registered trademarks of Microsoft Corporation. QuickTime is a registered trademark of Apple Computer, Inc. PhotoShop and Acrobat are registered trademarks of Adobe Systems Incorporated.

TABLE OF CONTENTS

CHAPTER 1	9
Introduction and Tutorial.....	9
Highlights (Common to Both VFSource and VFPro).....	9
Highlights (VFPro Only)	9
User Guide.....	10
Viewing this Guide Electronically	10
Sample Movies	10
VideoFOCUS: A Quick Tour.....	12
Capturing Media	12
Importing Video and Still Images	13
Browsing Video	13
Browsing Session Stills and Frames.....	14
Processing Sessions and Stills.....	14
Exporting Images and Session Movies.....	15
Setting Preferences	15
CHAPTER 2.....	17
Organizing, Capturing, and Importing Media	17
Overview.....	17
Working with Cases	17
Case Properties	18
Archiving of Cases	18
Restoring Archives for Use in VideoFOCUS	18
Capturing Video	19
Capture Window Features	20
Setting the Video Capture Device	21
Configuring Capture Devices	21
Configuring an Analog Capture Device	21
Configuring VideoFOCUS for Screen Capture	22
Configuring a Soundcard for Analog or Screen Capture	23
Importing Media Files.....	24
To Import Standard Multimedia File Formats	24
To Import Proprietary CCTV File Formats.....	25
Setting Video Attributes on Import.....	25
CHAPTER 3.....	27
Browsing Captured and Imported Sessions	27
Session Information	27
Browsing Frames.....	28
Setting the Thumbnail Display.....	29
Setting Frame Size.....	29
Browsing Video	30

Elements of the Player Window	30
Undocking and Resizing the Player Window	31
Zooming and Panning the Player Window	32
Marking Frames	33
Locating Still Frames	33
Printing Frames and Thumbnails	33
Viewing and Editing Session Properties	34
Processing Sessions.....	35
Exporting Sessions	35
Deleting Sessions	35
Labeling Segments (VFPro Only)	35
CHAPTER 4.....	37
Browsing Session Stills and Frames	37
Still Image Information	37
Viewing Still Images.....	38
Resizing the Still Window	38
Zooming and Panning the Still Image Window	38
Viewing a Frame of Video.....	38
Viewing Still Image Properties	39
Renaming Session Stills and Frames	39
Printing Session Stills and Frames	39
Processing Session Stills and Frames.....	39
Exporting Session Stills and Frames	39
Deleting Session Stills and Frames	40
CHAPTER 5.....	41
Processing Videos and Still Images	41
Modify - Spatial Properties.....	42
Crop Frame	42
Change Aspect Ratio	42
Change Frame Size	43
Rotate Frame	43
Modify – Temporal Properties.....	44
Marked Frames Only.....	44
Force Uniform Frame Rate.....	44
Change Speed	44
Split Fields into Frames.....	45
Force Interlace Format.....	45
Modify – Audio Properties.....	46
Sync Adjust	46
Channel Select.....	46
Channel Mix	46

Demultiplex	47
Demultiplex Window	48
Reassigning Frames and Views	48
Review All Frames in a View	50
Reassign Frames into Unselected Views	51
Saving Views	51
Strategies for Demultiplexing	51
Changing the Sensitivity of the Matching Algorithm	51
“Too Many” Views are Created	52
Mismatched Frames Grouped Together in Views	53
Video and Imaging Filters	54
Filter Overview	54
Zooming and Panning	54
Synchronizing Players and Locking	54
Filter Tools	55
Grayscale Filter	55
Sharpen Filter	55
Equalize Filter	55
Blur Filter	56
Stabilize Filter (VFPro Only)	56
Levels Adjust Filter	56
Super-resolution Filter (VFPro Only)	57
Manage Views	61
Combine (VFPro Only)	63
Macro Processing (VFPro Only)	63
Auditing Sessions and Still Images	64
CHAPTER 6	65
Processing Audio	65
Redact Audio	65
Extract Audio	65
Combine Audio and Video	65
Basic Audio Filters	66
Noise Reduction	66
Voice Spectrum Equalization	67
Hum Reduction	68
Automatic Gain	69
QuickEnhance	70

Advanced Audio filters - CARDINAL MiniLab (VFPro Only).....	72
MiniLab Introduction	72
General Plug-in Concepts	72
Filter Chain	75
1CH Adaptive Filter	78
Reference Cancellor.....	82
Spectral Inverse Filter	86
Comb Filter.....	94
X-Pass Filter.....	97
Spectral Subtraction Filter	111
Multi-Band Filter	116
Graphic EQ	125
Gain Stage	131
Parametric Equalizer	137
CHAPTER 7	141
Editing Video	141
CHAPTER 8.....	147
Exporting Media	147
Exporting Still Images	147
Exporting Movie Frames	148
Exporting Session Movies and Audio	149
CHAPTER 9	153
Creating Reports	153
To Create a Report for a Session	153
To Create a Report for a Still Image	153
APPENDIX A	155
Troubleshooting.....	155
Contacting Technical Support.....	155
APPENDIX B	157
VideoFOCUS Registration.....	157
License Information	157



CHAPTER 1

INTRODUCTION AND TUTORIAL

VideoFOCUS provides all the features you'll need to get superior results from video. There are two versions of the product: VFSource, providing basic capabilities, and VFPro, providing additional advanced capabilities such as multi-video synchronization, super-resolution processing, and macros.

HIGHLIGHTS (COMMON TO BOTH VFSOURCE AND VFPRO)

- Directly imports most standard multimedia and many proprietary CCTV file formats.
- Analog and Digital capture: Whether the source of your video is a VHS tape, DVR player, or streaming video off the web, VideoFOCUS will capture it.
- Automatic labeling of “key frames” (or “I” frames) on imported compressed videos.
- Allows user “marking” of important video frames to allow quickly accessing them again later.
- A wide variety of forensics operations and filters are available for images, videos, and audio.
- Simple video and image redaction via the Blur filter.
- Simple audio redaction through selection of segment and replacing with either tone or silence.
- Simple cut, copy, and paste editing of movies and images into final presentable videos.
- Comprehensive reporting function, which details all necessary steps to reproduce any processed or edited session or image and includes any notes added by the user.
- Speed and simplicity: VideoFOCUS is easy to use and install. The intuitive software, with a friendly graphical user interface, can be effectively utilized with no specialized training.
- Portability: VideoFOCUS runs well on most Windows-based laptop and desktop personal computers.

HIGHLIGHTS (VFPRO ONLY)

- Can additionally synchronize multiple videos to a single “MultiView” session.
- In addition to simple marking of video frames, user-defined labels can also be applied to selected segments to allow both quick access and quick selection for editing purposes.
- Provides additional advanced filters such as Super-resolution, Stabilize, and CARDINAL MiniLab audio filtering.
- Can quickly repeat the processing steps for any session or image through the application of automatically-recorded “macro” operations.

USER GUIDE

The user guide provides detailed information about all VideoFOCUS procedures, commands and features. It assumes you have a basic working knowledge of the Windows operating system, including using a mouse, selecting items in menus and dialog boxes, and opening, saving, and closing files. For information about these, please refer to your Microsoft documentation.

VIEWING THIS GUIDE ELECTRONICALLY


This guide is also provided as a Portable Document Format (PDF) file under the Help menu. To view it, you must have Adobe Acrobat Reader installed on your computer.

Note: *Acrobat Reader will automatically install when you attempt to open the user guide for the first time, or you can download it from Adobe.*

The PDF version of this guide includes hypertext links to help you move quickly to related points in the guide or to the Salient Sciences website. Hypertext links are displayed as underlined text.

SAMPLE MOVIES

The sample movies provided with VideoFOCUS allow users to get a head start exploring the features demonstrated in the tutorial presentations. Once a sample movie is selected it will appear in the VideoFOCUS application window. Sample movies can be deleted at any time and reimported from the Help menu, as needed.



- **Quad:** This is an example of a four camera video, divided into quadrants. Users can crop and resize the quadrants.
- **Multiplex1, Multiplex2 and LEAF:** These movies are three examples of multiplexed video. These are useful for learning demultiplex techniques.
- **Bank, Silver Car:** These movies are two examples of video that benefit from the application of super-resolution and masking.
- **Gates, Nice Robbery:** These movies contain an audio track. VideoFOCUS supports import, export and editing of media files containing audio.
- **Stabilization (VFPro only):** This movie is an example of shaky video shot

inside a moving vehicle. Users can apply stabilization filtering to correct this.

- Levels: This movie has low light conditions. Users can use this to experiment with the levels filter.


VIDEOFOCUS: A QUICK TOUR

This tour gives you a quick overview of the basic commands and features of VideoFOCUS. Take a look at the following sections for an overview of how to organize, acquire and process video, or skip ahead to [CHAPTER 2](#) to start capturing video.

1. Creating Cases
2. Capturing and Importing Media
3. Browsing Video
4. Browsing Session Stills and Frames
5. Processing Sessions and Still Images
6. Editing Session Movies
7. Exporting Images and Session Movies
8. Setting Preferences

CAPTURING MEDIA

VideoFOCUS digitizes the incoming video signals through installed capture cards or any source window playing video on the desktop. Use the Capture window to name the captured clips and to start and stop capturing. See [CHAPTER 2](#) for details on capturing video.



IMPORTING VIDEO AND STILL IMAGES


VideoFOCUS can import many proprietary DVR formats as well as standard video formats such as QuickTime, Windows Media, AVI, MPEG and many others. It can also import a wide variety of still image formats. See Importing Media Files for details on importing video.



BROWSING VIDEO

VideoFOCUS stores captured and imported video and displays it in a thumbnail or video player format. Users can toggle between both the thumbnail view and the video player view by clicking the Frames and Player tabs at the top of the application window. Additionally, if users wishes to view the video player and frames simultaneously, the player can be undocked by double clicking on any frame. Closing the player window re-docks it to the main window.


See [Chapter 3](#) for information about Browsing Captured and Imported Sessions.



BROWSING SESSION STILLS AND FRAMES

The Session Stills and Frames window is where you can view the images and stills you have created from any video session. It is also where you can find single images that you have imported into VideoFOCUS.


See [CHAPTER 4](#) for more information about browsing session stills and frames.



PROCESSING SESSIONS AND STILLS


Once sessions and still images are stored in VideoFOCUS, they can be processed further with a wide variety of filters.

See [CHAPTER 5](#) for more information about Processing Sessions and Stills.




EXPORTING IMAGES AND SESSION MOVIES

When you have finished processing your videos and still images, you can export them in the format of your choice. See [CHAPTER 8](#) for more information about exporting images and session movies.



SETTING PREFERENCES

The Preferences dialog allows you to set application-wide default settings, such as language, time zone for cases, and image export formats. This dialog can be accessed by clicking *Edit->Preferences*.





CHAPTER 2

ORGANIZING, CAPTURING, AND IMPORTING MEDIA

OVERVIEW

VideoFOCUS allows you to organize your videos and images into collections called *Cases*. You can have more than one case open in VideoFOCUS at any time. It is easy to open and close cases, move media between cases and archive them for later use.

VideoFOCUS can capture analog video from video tapes, digital video from proprietary video players or online media players where a file is not readily available. VideoFOCUS can also import video files stored in standard file types, single images and image sequences - a number of still frames typically representing a duration of sequential time.

Through the use of FFmpeg and a special software component supplied under license by DME Forensics, Inc., VideoFOCUS includes extensive import support for a wide variety of media formats and codecs. It provides VideoFOCUS with the ability to import video with accompanying audio from a wide range of media file types, including many proprietary DVR formats.

WORKING WITH CASES

Cases are the way you can organize your videos and images into a group related to a particular project. Upon launching VideoFOCUS you will need to make a case into which you can then capture, import, and process media. You may have more than one case open in VideoFOCUS at any time, each of which will appear as a separate “tab” in the application. You can close cases at any time, and the current contents will be automatically saved.

To create a new case

1. Choose *File->Open or Create Case*
2. Give the case a meaningful name, then click the *Create* button

To open a case


1. Choose *File->Open or Create Case*
2. Navigate to the Directory where your Case is stored.
3. Select the case then click the *Open* button

To close a case

1. Select the tab for the case you wish to close
2. Choose *File->Close Case*

CASE PROPERTIES

To provide basic details for reporting purposes and to allow for consistent treatment of local vs. UTC time for timestamping and synchronization purposes within each case, the Case Properties window is provided under *File->Case Properties*. By default, when a new case is first created the Notes field will be empty, the other text fields will be set to the last values used, and the Time Zone will be set to the default value specified under *Edit->Preferences*.



The screenshot shows a window titled "Case Properties: Convenience Store Robbery". It contains several text input fields and a text area. The "Agency" field is filled with "Anytown USA Police Department". The "Examiner" field is filled with "John Smith". The "Case #" field is filled with "NC20171115". The "Title" field is filled with "Convenience Store Robbery". The "Notes" field is a large text area containing the text "Retrieved CCTV footage from the robbery and homicide at the Qwikie Mart convenience store at 555 Main Street, Anytown, USA 12345". At the bottom, there is a "Time Zone" dropdown menu set to "America/New_York", and two buttons labeled "Local" and "UTC" next to it. Below these are "Save" and "Cancel" buttons.

ARCHIVING OF CASES

You may further archive your case into a ZIP file for easy sharing and storage. In order to restore this data into VideoFOCUS, you must unpack the zip file into a directory. This directory will then be seen by VideoFOCUS as an openable case.

To archive a case

1. Choose *File->Archive Case*
2. Select a location
3. Name the Case
4. Click *Save*

RESTORING ARCHIVES FOR USE IN VIDEOFOCUS


In order to restore archived cases into VideoFOCUS, you must unpack the archived ZIP file into a directory. This directory will be seen by VideoFOCUS as a case. Once the ZIP file is uncompressed, follow the instructions on [opening a case](#).

CAPTURING VIDEO

According to which device is selected in the Device menu, users can capture signals through an installed video card (such as Decklink or Microsoft DV) or from a video player via **Screen Capture**. When capturing an incoming video signal, digitizes the signal and stores the video on the computer's hard drive. The stored digital video is called a **Capture Session**.

To open the Capture window, choose *File->Capture*

The bar indicates
free disk space



Incoming signals
are displayed here

Choose the Audio
and Video devices
from which to capture
here

CAPTURE WINDOW FEATURES

The following table describes the features of the Capture window:

Feature	Video Status options
Monitor Area	This area allows you to watch the incoming video signal. The label below indicates whether you are “Monitoring” or “Capturing”.
Alert Messages	This area is below the Signal Monitor. If there is not enough free space for your capture session, or no video signal, VideoFOCUS will display an alert message. This area will also indicate other problems with capture such as slow writes to the disk.
Video Summary	The Video Summary Area displays the amount of available disk space. Blue indicates the amount of data currently residing on the drive. Free space is left blank. When capturing, the data being written to disk is shown in red.
Audio Selector	Allows the selection of the audio capture device.
Video Selector	Allows the selection to capture with an installed capture card, such as a BlackMagic card, or to execute a ScreenCapture.
Configure	Allows the user to select the correct format for a PAL or NTSC signal. See Setting the Video Capture Device
Source	Activates the source window selection window for Screen Capture.
Name	Where a session movie is named before capture.
Duration	Where a user chooses to capture a fixed time in minutes, or to record continuously.
Start/Stop Capture	Initiates and stops a capture.
Cancel Capture	Cancels a capture, eliminating all captured data.

To capture video

1. For analog video make sure your video source is plugged in and powered on. For screen capture, make sure your video player is open on the desktop and you can visible see the video you wish to capture.
2. Select the Video device. If you wish to select Screen Capture as your device, see [Configuring VideoFOCUS for Screen Capture](#).
3. Enter a name for your Capture Session in the Session Name field.
4. Choose the duration for your capture (you can manually start and stop the capture process, or you can capture for a fixed number of minutes. The capture will stop after the duration of time specified).
5. Click Start Capture to begin capturing video.
6. Click Stop Capture when you’ve captured the footage needed.
7. Click Cancel Capture if you wish to abort the capture session. The captured video will not be saved.

SETTING THE VIDEO CAPTURE DEVICE

The devices listed in the menu depend on what video capture and audio hardware is installed in the computer. Choose the device you intend to capture with.

Many digital video recording devices (DVRs) require a proprietary player to display captured video. There is often no way to export a standard video file from these systems. Screen Capture provides a way of capturing this data as it plays on the monitor. Additionally, Screen Capture can be employed to capture streaming video off the web, if desired by choosing the source window to be the Internet browser (Chrome, Mozilla etc.) and selecting the region where the video.

Note: You may find that various DirectShow and Video for Windows devices appear in the Capture Video Dialog. Some of these devices may work with VideoFOCUS if they produce video sessions in a suitable format. Please contact Salient Sciences support for more information on the current list of supported cards - (+1) 919.572.6767 (option 2) or support@salientsciences.com.

CONFIGURING CAPTURE DEVICES

Depending on the Video Device selected, you will need to configure the Video Source and Video Format options as described below.

CONFIGURING AN ANALOG CAPTURE DEVICE

1. For a Blackmagic Video Capture Device:

First be sure to configure the BlackMagic card via the BlackMagic Desktop Video control panel on your Windows machine.

Note: Video Format: Select the item that says UYVY (720x486) if you are capturing NTSC, the video standard in North America and Japan; or UYVY (720x576) if you are capturing PAL, the video standard in Europe, Asia, and Australia. Click Accept.

2. For a DV Video Device: Select a 1394 DV capture device in the Video Device list. Typically, your DV input device will be called "Microsoft DV Camera and VCR"

Note: For 1394 DV capture devices, the playback device (camera or deck) you are using must be connected and powered on when you open the Capture Video dialog. Your playback device needs to output a signal in order to be recognized by VideoFOCUS. Most playback devices automatically output a signal when powered on. If yours does not, you might have to play a tape to output a signal.

Note: Video Format: Select the item that says DVSD (720x480) or DVSD (720x576) and click Accept. VideoFOCUS recognizes two video standards: NTSC (720x480), the standard in North America and Japan; and PAL (720x576), the standard in Europe, Asia, and Australia. When you select a 1394 DV video capture device, VideoFOCUS determines the standard automatically and presents the proper format for selection.


CONFIGURING VIDEOFOCUS FOR SCREEN CAPTURE

1. Launch VideoFOCUS.
2. Once VideoFOCUS has opened, launch the proprietary video player or open the window containing your source video.


Note: You may need to install any proprietary players on your computer first. It is recommended that you run the player from your local disk, not the CD drive, as this will ensure optimal performance from the player.

3. Open the capture panel. Choose Screen Capture in the Device menu and click Source. The Source Windows dialog appears. The drop-down menu lists all the windows present on your desktop.
4. Choose the name of the window you want to capture. Once a window is selected it will be brought to the foreground and display in the source selection window.

Note: In case a selected window does not display in the source selection window, you may want to decrease the Hardware Acceleration for the computer's graphics card. To do this, exit VideoFOCUS, then go to: Start Menu > Control Panels > Display > Settings tab > Advanced button > Troubleshoot tab, and lower the Hardware Acceleration slider. If you are using Windows Media Player, you can lower the Hardware Acceleration within the player itself by opening: Tools > Options > Performance, and lowering the Video Acceleration slider.




DVR Player Window




Source Selection in Capture Media Dialog
Video Portion is being selected with crop tool

5. Use the crop cursor to select only the video portion of the player (most media player windows are comprised of video displayed within a console). Additionally, you may choose a preset pixel size from the drop-down menu or manually enter a custom size using the pixel height and width text fields. If you wish to move the crop box, select the arrow tool to drag it.




Preset Pixel Sizes



Pixel fields and Crop Tools

6. If there is timecode overlaying the video, you may want to select the Remove Overlay Flicker checkbox. This will help ensure that the timecode remains visible at all times in the captured video.



7. Select an Audio Source, if your video contains audio and you wish to capture it. See [Configuring a Soundcard for Analog or Screen Capture](#).
8. Click OK and the selected portion of the source window is displayed in the capture panel.
9. Activate the player to play the desired video segment. Click Start Capture to begin capturing video data. If the player occupies the full screen, the capture panel will go to the background, ensuring no overlap of the target window.
10. Click Stop Capture when you wish to end the capture session.



Double-clicking the VideoFOCUS icon that appears in the lower right corner of the screen will stop video capture.

11. The captured video will appear in the Captured and Imported Sessions list.

CONFIGURING A SOUND CARD FOR ANALOG OR SCREEN CAPTURE

VideoFOCUS will capture audio, in addition to video, if a compatible sound card is available and properly configured. Please contact us if you have any questions as to whether your specific sound card will work with VideoFOCUS - (+1) 919.572.6767 (option 2) or support@salientciences.com.

IMPORTING MEDIA FILES


VideoFOCUS can import most video formats that open in QuickTime or Windows Media Player. Additionally VideoFOCUS can import single image files or a collection of sequential image files and convert them to a video at a specified frame rate. Some, but not all, of the supported image and video formats that VideoFOCUS supports include the following:

Format	Description
QuickTime (.mov)	Video file format developed by Apple Computer.
Windows Media (.asf, .asx, .wmv, .wm)	Developed by Microsoft for streaming and playing back media files, consists of: Advanced Streaming Format (ASF), Active stream redirector (.ASX) Windows Media Video (WMV), Windows Media (WM).
AVI (.avi)	Audio Video Interleave, a Microsoft format for digital audio and video commonly produced by video recording devices.
MPEG (.mpeg, .mpg, .m1v, .m2v)	Includes MPEG-1 and MPEG-2 formats. MPEG-1 is a low-resolution format widely used on the Internet. MPEG-2 is a higher resolution format, most commonly used for DVD.
MPEG-4 (.mp4, .m4v, .mp4v)	Designed to transmit video and images over a narrow bandwidth, the current primary use for MPEG-4 video is Internet streaming media.
Indeo (.i)	A high compression format developed by Intel, commonly used for CD-ROM production.
DivX (.divx)	A popular compression format based on MPEG-4. DivX files can be downloaded over high-speed lines in a relatively short time with minimal loss of video quality. Often used to exchange video files on the Internet.
DV (.dv)	Digital Video (DV) defines both a compression and a tape format. It is popular for consumer and semiprofessional video production.
VOB (.vob)	Short for Video Object. A DVD file that contains the actual Video, Audio, Subtitle, and Menu content.
3G Mobile (.3g2, .3gp)	A low-quality format used in mobile phones to store audio/video files.
Various Proprietary CCTV Formats	Supports various proprietary CCTV DVR file formats.
Most Standard Image Formats	Can also load still image formats such as JPG, BMP, TIF, etc.

TO IMPORT STANDARD MULTIMEDIA FILE FORMATS

1. Select *File->Import Media*. "All Files" is the default file type selection. If you wish to narrow your search to either Video, Images or Audio, select it in the file type drop down. Navigate to and select the file you wish to import.

The Import Media provides a viewer that allows, wherever possible, reviewing the video and images prior to import.



Note: One or a series of videos can also be imported into the application by dragging them from a file explorer window and/or the Desktop.

TO IMPORT PROPRIETARY CCTV FILE FORMATS

1. Select *File->Import Media*.
2. Choose All Files from the file type drop down.
3. You may wish to modify the Accept drop down to narrow the type of media you wish to import from the CCTV file.
4. Navigate to and select the file you wish to import.

Note: In VFPro, when importing a CCTV file with multiple camera views, it will place them all in a single MultiView session. In VFSource, each camera view will be in its own single-view session.

SETTING VIDEO ATTRIBUTES ON IMPORT

VideoFOCUS will allow you to change the attributes of a video (or series of sequential images imported as a video) via the **Movie Attributes** window.


Transcoding

Certain video formats may require transcoding before they can be imported. Transcoding is a process by which the original video is copied and converted to a format that VideoFOCUS can more readily work with. Typically, file formats that may require transcoding are streaming formats (.asx, .asf, .wmv) that do not display video frames unless the clip is playing. This can cause problems if the user wishes to browse through the video frames at random. Transcoding is a way of creating a non- streaming copy of the original video.

To Transcode a Video File

1. Select *File->Import Media*.
2. Navigate to and select the video file you want to import.

3. Select the Options button which will open the **Movie Attributes** window. Click on the **Compression Settings** button, which contains a list of compression formats.



Note: The default compressor is *None*, which uses no compression when converting the file. This setting is preferred if you wish to maintain the quality of the original video, but may result in large files if the original video is several minutes long. In that event, any compressor can be selected from the list to create a smaller file.

4. Click OK to close the Movie Attributes window.
5. Click Open in the Import Media window. The transcode progress bar will appear. Once the transcode is complete, the file will appear in the session list.

Always Transcode Videos

If you wish to always transcode videos with a particular setting on import, click the **Always Transcode** checkbox. The videos will be transcoded with the current settings set in the Compression Settings Dialog.

Override Aspect

The Override Aspect Field allows the user to override the current aspect ratio. Please enter values in the form X:Y where X is the width and Y is the height.

Override Interface Format

Some video may have the incorrect aspect ratio set or it may not be defined at all. If you know the interlace format of the video, you can set it in the Override Interface Format pulldown menu. Options are

- Not Interlaced
- Interlaced (Even)
- Interlaced (Odd)




CHAPTER 3

BROWSING CAPTURED AND IMPORTED SESSIONS

Video can be brought into VideoFOCUS through various methods: by capturing it digitally, by screen capturing from digital media players, by importing self-contained media files, such as those encoded as QuickTime (.MOV) Windows Media (.WMV) or AVI (.AVI) or by importing a set of sequentially-numbered still images that are converted into a video file. The import process can also be initiated by a drag & drop or copy & paste action into or within VideoFOCUS. Videos reside in the **Captured and Imported Sessions** pane in VideoFOCUS.

When you acquire video, VideoFOCUS automatically displays it in the **Captured and Imported Sessions** pane.



A session as it appears in the Captured and Imported Sessions window (Note the “question mark” icon indicating that the video is sourced from the VideoFOCUS help system).

For imported videos that are compressed, as they are scrubbed or played you will notice a blue “Key” indication on the video player for each frame that VideoFOCUS has determined to be a key frame, also known as an “I” frame, meaning that that frame is a complete image from the camera at that moment in time. Frames that are not so designated are *interstitial* compressed frames, also known as “P” and “B” frames, which are not complete frames but are instead more heavily compressed delta frames that are calculated in between the key frames based upon motion tracking. In general, if you are interested in extracting fine details from the video frames such as faces and automobile number plates, you want to look most closely at the key frames than the interstitials.

SESSION INFORMATION

As sessions are acquired, the left column of the window becomes a scrollable list. To select a session, simply find it in the list and click on it. When it is selected, it will become blue in color and an orange highlight will appear. Any frames associated with the session will display as thumbnails, or in a video player, in the right-hand panel of the window. For audio-only sessions, the waveform graph will appear in lieu of video or images.

Original sessions are distinguished from their processed versions by means of icons

over the poster frame, as well as hierarchical visual relationship. For the sample movie sessions that are available under *Help->Sample Movies*, a “question mark” icon will appear over the poster frame to indicate that it comes from the Help system within VideoFOCUS itself. For sessions that have been imported from a file, a “disk” icon appears; for sessions that have been captured, a “camera” icon appears.

As processed versions of original sessions, or “child” sessions, are created, “arrow” icons appear above the poster frames for those sessions and all of their “parent” sessions will visually collapse to a smaller size unless they are explicitly selected by the user for viewing. A “↓” indicates that the session has been produced by processing the parent immediately above it in some way. A “↩” indicates a “branched” session, meaning that there is at least one other child session in the case that has the same immediate parent. Finally, a “↩↓↩” indicates that the child session has been produced from multiple parents using a Combine operation.

Because each session can be unique, session names and creation dates are always displayed in the list, in addition to other information.

The following table describes the items displayed in the session list.

Session Information	
Session Name	The name of the session.
Session Length	The length of the session in minutes and seconds.
Created Date	For capture sessions, the date and time that the session was captured; for imported file sessions, the file creation date will be used or, if metadata is available within the file, the file’s internal time/date stamp will be used.
Audio Format	The sample rate and number of channels of audio.
Interlaced	The interlace status of the movie. “Interlaced (Odd)” means the session is field interlaced with the odd field dominant. “Interlaced (Even)” means the session is field interlaced with the even field dominant. “Not Interlaced” means the session is not interlaced, also called progressive scan. Field dominance refers to which field is scanned first in a frame of video.
Marked Frames	The number of marked frames in the session.
Still Frames	The number of video frames collected as still images, plus any super-resolution stills and snapshots, originating from the session.

BROWSING FRAMES

Once a session is selected, the thumbnails of the individual video frames are displayed in the right panel of the Captured and Imported Sessions pane. Browsing frame thumbnails is one way to navigate through a video.

If you click on an individual frame with which you want to work, you can quickly view it in detail using the video player.



Frame thumbnails display for a selected session when the Frames tab is chosen.


Users can toggle between viewing frames or the video player by clicking the **Frames** and **Player** tabs at the top of the application window. Both views are synchronized, meaning, that if you click on a thumbnail in the Frames view, the Player view will display that frame when selected. Likewise, if you scrub the video player to a particular area, the Frames view updates in the background to display that area, as well.

Additionally, to view the video player and frame thumbnails simultaneously, the player can be undocked by double clicking on any frame. Closing the player window re-docks it to the main window.

SETTING THE THUMBNAIL DISPLAY

The Thumbnail Display is a drop-down menu that lets you choose a granularity for viewing the captured video. The menu can display a range of individual frames, from one frame per minute to every frame. You can also view only marked frames.


This can be a useful feature when working with longer videos, of several minutes in length. For example, by adjusting Thumbnail Display setting (to a setting such as Every Minute or Second), you can browse the video to get a rough overview of it, and then ultimately find a frame of interest using a finer setting (such as Every Second or Frame).



Adjust the Thumbnail Display to get a quick overview of long capture sessions

SETTING FRAME SIZE

The size of the frame thumbnails can be adjusted using the Thumbnail Size slider. The smallest setting is 80x60 pixels, but they can be resized to be much larger, to the full width of the application window, if necessary. To enlarge the size of the frame thumbnails, simply drag the slider to the right. Dragging the slider to the left shrinks them back again.



BROWSING VIDEO

Once a session is selected, users can toggle between viewing it as frames or as a movie by clicking the Frames and Player tabs at the top of the application window in the video player.



A video player displays for a selected session when the Player tab is chosen.

Both views are synchronized, meaning, that if you click on a thumbnail in the Frames view, the Player view will display that frame when selected. Likewise, if you scrub the video player to a particular area, the Frames view updates in the background to display that area, as well.

Additionally, to view the video player and frame thumbnails simultaneously, the player can be undocked by double clicking on any frame. Closing the player window re-docks it to the main window.

To view your movie in the video player

1. Select a session from the session list
2. Choose *View->View Movie*, double-click the thumbnail, or click the Player tab. The video player window opens, displaying the selected frame.
3. Use the playback controls to view the video. You can also use the space key to start and stop playback. The arrow keys frame forward or frame backward.

ELEMENTS OF THE PLAYER WINDOW

The player window has series of controls that can affect the way video is displayed,

scaled and selected.

UNDOCKING AND RESIZING THE PLAYER WINDOW

The player window can be undocked from the main application window by double clicking on it. This creates an independent window with its own menus. The View menu has size options for scaling the player.

The following table describes the items in the View menu.

View Menu	
Half Size - Ctrl+1	Sets the player to display the video at one-half its native size.
Normal Size - Ctrl+2	Sets the player to display the video at its native size.
Actual Pixels - Ctrl+3	Sets the player to display the video at its actual pixel size (a change between "Actual" and "Normal" may not be noticeable unless the video has been field split, demultiplexed, or its display size has been adjusted for non-square pixels).
Maximize - Ctrl+4	Sets the player to fill the entire screen.
Aspect Adjust	Adjusts the width of the player to display a video with the aspect ratio corrected (a change may not be noticeable unless the video has been field split, demultiplexed, or its display size has been adjusted for non-square pixels).

Note: In addition to resizing the player window via the View menu, the player can be resized freely, by dragging the edges of the player window. Also, pressing F11 on the keyboard will always maximize any dialog that contains a player.

To undock the video player

To access the menu settings above, the video player must first be undocked. To undock it, select a session from the session list, then do any of the following:

- Choose View->View Movie
- Double-click any session thumbnail
- Right-click on the docked player, and select Undock from the menu. To redock the player, close the player window.

More About Display Size

Certain operations in VideoFOCUS can alter the pixel dimensions of a video frame, specifically, field splitting and cropping. Following a field splitting operation, only one field, or set of scan lines will be present in a frame. Therefore, an NTSC video frame, which is normally 720x486 pixels, will be 720x243 pixels. Because only half of the horizontal lines are present, the frame without correction would seem squashed if viewed in a player as shown in Figure 1).



Figure 1 - Display of "Actual Pixels" of a field split movie

Viewing material this way can be difficult, for obvious reasons. As a solution, VideoFOCUS compensates for the difference in pixel height by adjusting the video player width, thus providing a display comparable to the original pixel dimensions. The underlying data is not changed in this operation, only the visual display size as shown in Figure 2.




Figure 2 - Display width "Aspect Adjusted" for pixel dimensions of a field split movie.

Note: Movies exported from VideoFOCUS contain information which instructs media players to display our movies this way as well. This can be enabled or disabled at the user's discretion. See [Chapter 8, Exporting Media](#).

ZOOMING AND PANNING THE PLAYER WINDOW

Frames in the player window can be zoomed by rolling the mouse wheel (if you don't have mouse with a wheel, there is currently no other way to zoom). Clicking and dragging the mouse when zoomed in repositions or pans the frame in the context of the window.

To zoom and pan the frame of video in the video player

1. Open the player by double-clicking a frame thumbnail
2. Roll the mouse wheel to zoom in and out
3. Hold down the left mouse button and drag the image to pan it

MARKING FRAMES

The Label Frame, Marked command is used to designate specific frames of interest. This makes it easier to find the same frames later, after browsing elsewhere in the video.

To mark or unmark a frame

1. Select a frame, either a thumbnail or a frame in the video player.
2. Choose *Edit->Label Frame->Marked* to toggle between marked and unmarked. Marked frames are displayed with a solid orange marker.

Note: *Marked frames are displayed in the frame sampler pane with a solid orange marker. A striped orange marker means that the frame itself is not marked, but one that it represents in a finer filter setting is marked. Adjust the filter Browse Video setting to a finer setting such as Every Frame to see the exact marked frames.*



A striped orange marker, compared to a solid orange marker.

LOCATING STILL FRAMES

A frame is considered a still frame only if it has been used to create a super-resolution still image or has been dragged or copied into the **Session Stills and Frames** pane. These frames are displayed with a red marker. To locate any still frames, set the Thumbnail Display menu to *Still Frames*.

To view all Still Frames

1. Set the Session window to Frames view
2. Set the Thumbnail Display menu to *Still Frames*

PRINTING FRAMES AND THUMBNAILS

You can print any frame from the video player window at screen resolution. The print size will adjust to fill the page in landscape format. You can also print the whole window of frame thumbnails at any size the Thumbnail Size slider is set to.

To print a frame

1. Set the Session window to Player view
2. Select a frame in the video player.
3. Choose *File->Print Window* to print the frame.

To print the window of frames

1. Set the Session window to Frames view
2. Choose *File->Print Window* to print all the frames.

To print all Marked Frames

1. Set the Session window to Frames view


2. Set the Thumbnail Display menu to *Marked Frames*
3. Choose *File->Print Window* to print all the marked frames.

Note: *Still frames are displayed in the frame sampler pane with a solid red marker.*

VIEWING AND EDITING SESSION PROPERTIES

To access session properties

1. Select a session. Choose *Edit->Session Properties*.
2. This will open the Session Properties window.
3. If desired, a new name for the session can be entered (by default, sessions are automatically named according to the import filename and any processes that have been applied to it).
4. Optionally, if the session is the original import (as indicated by the “disk” icon over the session thumbnail), the MD5, SHA1, and SHA-256 hash values for the imported file can be viewed. These values will also appear in any report that is generated for this session or any processed sessions that have been produced from it.
5. Optionally, Item Notes can be added to provide a user narrative on what the item is, how it has been processed, what can be seen, etc. These Item Notes will appear in any report that is generated for this session.



PROCESSING SESSIONS

You can perform many operations on existing sessions such as filtering and demultiplexing. Further modifications, such as cropping, adjusting the speed or frame rate, and changing the aspect ratio are also possible. See [CHAPTER 5](#) for more information.

EXPORTING SESSIONS

Any session can be exported as an AVI (.avi), QuickTime (.mov) or Windows Media (.wmv) file. See [CHAPTER 8](#) for more information.

DELETING SESSIONS

If you are running out of space on your video drive, you might need to delete some sessions. When you delete a capture session, it is important to understand that you are deleting everything associated with it, including marked frames and still frames.

Note: *Exported stills or videos derived from a session are not affected if it's deleted.*

To delete a session

1. Select a session in the Browse Video window. You can also shift-select a group of sessions, if you wish to delete more than one at a time.
2. Choose *Edit->Delete Session*.

To Delete All Related Sessions

Many times sessions reference data from one another. For example, if a multiplexed video is digitized and then demultiplexed, any number of saved "views" from this session will reference the original multiplexed movie. Additionally, filtering, editing, or other modifications can create a session that is not self-contained at all, but merely references the original video data. The idea behind Select All Related Sessions is to find any file that references the underlying media on the hard disk. This is a sure way of clearing off space when many sessions reference the same media file and it's not obvious to the user which session is the original.

Note: *Once the media file of a session is deleted from the hard disk it can't be restored, unless it was an imported file to begin with and resides elsewhere, rather than the VideoFOCUS movie storage directory.*

1. Select a session.
2. Choose *Edit->Select All Related Sessions*. All related sessions will be highlighted.
3. Hit the Delete key to delete all the selected sessions.

LABELING SEGMENTS (VFPRO ONLY)

The Label Frame, New command is used to give meaningful names to selected segments of video, thereby "tagging" them for easier access later. For example, every place in the session where a particular car used in a crime appears could be labelled as "Suspect Blue Car", or every place that a particular person appears could be labelled as "Robbery Suspect" or "Victim #1". Labels are universal to each

case, meaning that once created they are available to be applied to segments of ANY session in the case. This makes it easier to find the same segments later, also to quickly select all sessions containing the label of interest.

To apply a new label to a segment of video

1. Scrub to either the beginning or the end of the segment of interest, then hold down the Shift key and scrub to the opposite end of the segment. Note the orange bar that now appears on the scrubber bar, indicating the temporal selection.
2. Choose *Edit->Label Frame->New*, then type in a meaningful name where indicated and, if desired, click the Change button and choose a foreground and background color for the label.
3. Click OK to apply the label, which should now appear at the top left of the video player whenever the playhead is within this labelled segment.

To reuse an existing label on a segment of video

1. Scrub to either the beginning or the end of the segment of interest, then hold down the Shift key and scrub to the opposite end of the segment. Note the orange bar that now appears on the scrubber, indicating the temporal selection.
2. Choose *Edit->Label Frame*, then select the existing label to be applied from the dropdown list. The label should now appear at the top left of the video player whenever the playhead is within this labelled segment.

To select all segments in a session having a particular label

1. With the session selected and visible on the video player, Click on *Edit->Select Times Labelled*, then choose the label of interest from the dropdown menu. All segments having this label within this session will now appear as orange segments on the scrubber bar.
2. To quickly queue the player to the beginning or end of each labeled segment, click either the Next Selection Section (“”) button or Previous Selection Section (“”) button and observe the playhead move instantly to one end of an orange segment on the scrubber bar.

To select all sessions containing a label of interest


1. Click on *Edit->Select All Containing Label*, then choose the label of interest from the dropdown menu. All sessions having segments with this label will now be selected in the Captured and Imported Sessions area.



CHAPTER 4

BROWSING SESSION STILLS AND FRAMES

The Session Stills and Frames window stores any super-resolution stills or video frames derived from Captured and Imported Sessions. For example, if you were to apply super-resolution processing to a video frame, the results are automatically stored in this window. Likewise, any video frame of interest can be copied from a captured or imported session and pasted to this window, saving it for later processing, if needed. It also contains any single images imported directly into the application



This area contains stills made from a super-resolution process, single frames dragged from sessions and any single images imported directly into VideoFOCUS.

As stills and frames are added to the window, they arrange themselves according to the order of the videos in the session list. Stills or frames created from the same video sessions are automatically grouped together as well. Additionally, when a video session is selected above, any super-resolution stills or frames derived from it are given white highlights in the window below, so they can be easily located.

STILL IMAGE INFORMATION

The pixel dimensions of each still image are noted at the bottom of the image thumbnail. Hovering the mouse over the image reveals the name in a tooltip.

The pixel size and naming depends on how the still images were created. For example, if the still image was created by simply copying and pasting a video frame, into the window, the pixel dimensions are those of the original video, and the name will be the same as well, with the frame number added. Still images can be modified however, to change their size and resolution. See [CHAPTER 5](#) for more information.

VIEWING STILL IMAGES

To view a still image, find it in the window and click on it. When it is selected, an orange highlight will appear around it. By double clicking it, or hitting the Enter key when it is selected, it will open it in its own window.

RESIZING THE STILL WINDOW

When still image windows are opened, the View menu has size options for scaling the window. The following table describes the items in the menu.

View Menu	
Half Size - Ctrl+1	Sets the window to display the still at one-half its native size.
Normal Size - Ctrl+2	Sets the window to display the still at its native size.
Actual Pixels - Ctrl+3	Sets the window to display the still at its actual pixel size (a change between "Actual" and "Normal" may not be noticeable unless the source video has been field split, demultiplexed, or its display size has been adjusted for non-square pixels).
Maximize - Ctrl+4	Sets the window to fill the entire screen.
Aspect Adjust	Adjusts the width of the window to display a still with the aspect ratio corrected (a change may not be noticeable unless the source video has been field split, demultiplexed, or its display size has been adjusted for non-square pixels).

ZOOMING AND PANNING THE STILL IMAGE WINDOW

Opened still images can be zoomed by rolling the mouse wheel (if you don't have mouse with a wheel, there is currently no other way to zoom). Clicking and dragging the mouse when zoomed in repositions or pans the frame in the context of the window.

To zoom and pan the frame of video from a still image

1. Open a still image window by double-clicking a still thumbnail
2. Roll the mouse wheel to zoom in and out
3. Hold down the left mouse button and drag the image to pan it

VIEWING A FRAME OF VIDEO

All the still images in the Session Stills and Frames windows are derived from videos present in the session list. It's therefore possible to see images within the context of the video from which they originated.


To view the frame of video from a still image

1. Select a still image thumbnail
2. Choose View->View Movie
3. The video player will open, displaying the appropriate session movie. The video player will be set to the exact frame that was used to create the still image.

VIEWING STILL IMAGE PROPERTIES

To view still image properties

1. Select a still image.
2. Choose *Edit->Image Properties*.
3. This will open the Image Properties window.



RENAMING SESSION STILLS AND FRAMES

To rename a still image

1. Select a still image. Choose *Edit->Image Properties*.
2. This will open the Image Properties window.
3. In the Image Properties window, enter a new name, click OK.

PRINTING SESSION STILLS AND FRAMES

You can print any still image window at screen resolution. The print size will adjust to fill the page in landscape format. You can also print the whole window of Session Stills and Frames thumbnails.

To print a still image

1. Select a still image thumbnail.
2. Right-click and select *Print Image*.

To print the Session Stills and Frames window

1. Click once into the window so it has focus.
2. Choose *File->Print Window* to print all the still image thumbnails.

PROCESSING SESSION STILLS AND FRAMES

You can perform many operations on existing still and frames. Filtering and other modifications, such as cropping, and changing the aspect ratio are also possible.


EXPORTING SESSION STILLS AND FRAMES

Any still that has been created or modified can be exported as a JPEG, BMP or TIFF file. See [Exporting Still Images](#) for more information.

DELETING SESSION STILLS AND FRAMES

To delete a still image or frame.

1. Select a image thumbnail. You can also shift-select a group of images, if you wish to delete more than on at a time.
2. Choose *Edit->Delete Image*, or hit the Delete key.



CHAPTER 5

PROCESSING VIDEOS AND STILL IMAGES

Any video or still image in VideoFOCUS can be processed. Processing can be applied to any video selected in the Captured and Imported Sessions area, or to any image selected in the Session Stills and Frames area.

The Modify menu contains various operations that can be applied to both sessions and still images. The following tables describe the options in the Modify panel.

Modify - Spatial Properties

Crop Frame	Allows the user to crop the video to any size. The crop marquee initially appears constrained to the aspect ratio of the original frame, but can be resized freely by dragging the corners.
Change Aspect Ratio	Allows the user to correct or change the aspect ratio of a video or still image.
Change Frame Size	Allows the user to change the frame size of a video or still image, while preserving the original height and width ratio.
Rotate Frames	Allows the user to rotate a video or still by a specified amount. Negative values are allowed.

Modify - Temporal Properties

Marked Frames Only	Allows the user to create a modified new session that is made up exclusively of marked frames from the current session. The number of marked frames in the current session is indicated in a label below the checkbox.
Force Uniform Frame Rate	Allows the user to assign a new frame rate to the video (the frame-rate of the original file is not affected). This is a useful feature for speeding up or slowing down the motion of any kind of video file or standardizing the frame rate of a time-lapse video.
Change Speed	Allows the user to alter the speed of a clip by entering a percentage value above or below the value of its current speed. Entering a negative number will play the clip in reverse.
Split Fields Into Frames	Splits the two fields comprising each video frame into their own individual frames. This option is useful for video that has been field multiplexed or in any situation where there is a very large disparity between fields.
Force Interlace Format	This is used to reassign the interlace status of a video. This may need to be done with imported videos which have the incorrect field dominance assigned. Negative values are allowed.

MODIFY - SPATIAL PROPERTIES

Spatial options are available when modifying both video and still images. Videos are stored in the Captured and Imported Sessions window. Still images are stored in the Session Stills and Frames window.

CROP FRAME


If you are working with surveillance video that is divided into four quadrants, cropping can isolate each quadrant. Or, in the case of a still image, you can crop out any non-essential details.

To modify a session or still image by cropping

1. Select a video from the session list, or a still image from the Session Stills and Frames window.
2. Choose *Processing->Modify*
3. Choose the Crop Frame checkbox
4. Position and scale the marker to the area of the frame that you want to crop (placing the cursor in the middle of the crop box will position it, and placing it on a corner will scale it).
5. Name the modified session or still image.
6. Click Create Session or Still. Modified videos will appear in Captured and Imported Sessions window. Modified stills will appear in the Session Stills and Frames window.
7. You may create more than one cropped video or image from an original.
8. When done, click Done to exit the Modify panel.

CHANGE ASPECT RATIO

Aspect Ratio refers to the ratio of width to height of a picture. Standard definition television screens use a 1.33:1 aspect ratio (also known as 4:3). High definition television use a 16:9 (or 1.78:1) aspect ratio, which is a wider picture, similar to what is seen in a movie theater.



To modify a session or still image by changing the aspect ratio

1. Select a video from the session list, or a still image from the Session Stills and Frames window.
2. Choose *Processing->Modify*
3. Choose the Change Aspect Ratio checkbox (the current aspect ratio values

will appear).

4. Enter a new value to either text field.
5. Name the modified session or still image.
6. Click Create Session or Still. Modified videos will appear in Captured and Imported Sessions window. Modified stills will appear in the Session Stills and Frames window.
7. Click Done to exit the Modify panel.

CHANGE FRAME SIZE

Changing the frame size is a way of altering the pixel dimensions of a video or still image. When entering a new height or width, the other value changes accordingly to preserve the original height to width ratio. To modify a session or still image by changing the frame size. Typically this is only done on export to fit the movie to a particular needed format.

1. Select a video from the session list, or a still image from the Session Stills and Frames window.
2. Choose *Processing->Modify*
3. Choose the Change Frame Size checkbox (the current height and width values will appear).
4. Enter a new value to either text field.
5. Name the modified session or still image.
6. Click Create Session or Still. Modified videos will appear in Captured and Imported Sessions window. Modified stills will appear in the Session Stills and Frames window.
7. Click Done to exit the Modify panel.

Note: *The Constrain Pixel Shape checkbox will retain the current relative dimensions of the original object.*

ROTATE FRAME

Rotate frame will rotate the movie or image by the specified degree amount. Please note that negative values are allowed.

To modify a session or still image by rotating the frame

1. Select a video from the session list, or a still image from the Session Stills and Frames window.
2. Choose *Processing->Modify*
3. Choose the Rotate Frame checkbox.
4. Enter a degree value for rotation (negative values are allowed)
5. Name the modified session or still image.
6. Click Create Session or Still. Modified videos will appear in Captured and Imported Sessions window. Modified stills will appear in the Session Stills and Frames window.
7. Click Done to exit the Modify panel

MODIFY - TEMPORAL PROPERTIES

These options are only available when modifying video. They are not present when modifying still images.

MARKED FRAMES ONLY

Marking a frame is a simple way of bookmarking a frame of interest so it can be easily retrieved at a later time. Additionally, this feature can be used as a simple way of creating a new video comprised of a few selected frames.

To modify a session from marked frames

1. Select a video from the session list.
1. Select a frame, either by looking at a thumbnail view or in the video player.
2. Choose *Edit->Label Frame->Marked* to toggle between marked and unmarked.
3. Mark all the frames that you wish to include in your new session.
4. Choose *Processing->Modify*.
5. Choose the Marked Frames Only checkbox.
6. Name the modified session or still image.
7. Click Create Session.
8. A new session will be created that contains only those frames.
9. Click Done to exit the Modify panel.

FORCE UNIFORM FRAME RATE

This assigns a new frame rate to the video (the frame-rate of the original file is not affected). This is a useful feature for speeding up or slowing down the motion of any kind of video file or standardizing the frame rate of a time-lapse video.

To modify a session by forcing a uniform frame rate

1. Select a video from the session list.
2. Choose *Processing->Modify*.
3. Choose the Force Uniform Frame Rate checkbox.
4. The original frame rate is displayed. Raise the frame rate to speed up a clip. Lower the frame rate to slow it down.
5. Name the modified session or still image.
6. Click Create Session.
7. A new session will be created with the designated frame rate.
8. Click Done to exit the Modify panel.

CHANGE SPEED

This changes the speed of a video by entering a percentage value. For example, entering a value of 50% will slow the video to half its original speed, while entering a value of 200% will speed it up twice as fast. Use a negative number to play the video in reverse. While the speed of the video is changed, the frame rate of the video is not affected, as when using Force Uniform Frame Rate.

Note: When changing the speed of a video containing an audio track, the audio is not scaled with the video. The audio is muted in speed and frame rate adjusted videos.

To modify a session by changing speed

1. Select a video from the session list.
2. Choose *Processing->Modify*.
3. Choose the Change Speed checkbox.
4. The original percentage rate is displayed.
5. Raise the percentage rate to speed up a clip, or lower the percentage rate to slow it down.
6. Name the modified session or still image.
7. Click Create Session.
8. A new session will be created with the designated percentage rate.
9. Click Done to exit the Modify panel.

SPLIT FIELDS INTO FRAMES

Analog video frames are comprised of two interlaced video fields. The ghosting effects of field interlacing may not be obvious except when rapidly moving objects have been caught on video, or if a video has been field multiplexed. Splitting each field into a separate video frame can eliminate this effect, making the video easier to work with. See the section on [De-Interlacing](#) for a more detailed discussion of field interlacing.

To modify a session by splitting fields into frames

1. Select a video from the session list.
2. Choose *Processing->Modify*.
3. Choose the Split Fields into Frames checkbox.
4. Name the modified session or still image.
5. Click Create Session.
6. A new session will be created with the fields split into frames.
7. Click Done to exit the Modify panel.

FORCE INTERLACE FORMAT

On rare occasions when analog video is imported (most digitally derived media is not field interlaced) an incorrect field dominance may be assigned. This option is available to correct that, in necessary.

To modify a session by forcing the interlace format

1. Select a video from the session list.
2. Choose *Processing->Modify*.
3. Choose the Force Interlace Format checkbox.
4. Choose one of the three options (if you don't know the correct format, you may have to try more than one to get it right).
5. Name the modified session or still image.
6. Click Create Session.
7. A new session will be created.
8. Click Done to exit the Modify panel.

MODIFY - AUDIO PROPERTIES

These options are only available when modifying video with audio or an audio session. They are not present when modifying still images or a video session without audio.

SYNC ADJUST

This adjusts the selected audio channel's start time. This is useful to correct synchronization issues between the audio and video playback.

To modify a session with Sync Adjust.

1. Select a video or audio session from the session list.
2. Choose *Processing->Modify*
3. Choose Audio Properties.
4. Check one or more of the channel boxes under Sync Adjust.
5. Adjust the start time.
6. Name the modified session.
7. Click Create Session. Modified videos or audio sessions will appear in Captured and Imported Sessions window. Click Done to exit the Modify panel

CHANNEL SELECT

This allows an audio channel to be either retained (selected) or removed from the video session.

To modify a session with Channel Select.

1. Select a video or audio session from the session list.
2. Choose *Processing->Modify*
3. Choose Audio Properties.
4. Check to remove one or more of the channel(s) under Channel Select.
5. Name the modified session.
6. Click Create Session. Modified video or audio sessions will appear in Captured and Imported Sessions window. Click Done to exit the Modify panel

CHANNEL MIX

This allows you to either create a Mono (50/50) mix of all the audio channels or a Stereo balanced mix. The Stereo balance mix is useful when the level of one channel is lower than the other and needs compensation.

To modify a session with Channel Mix.

1. Select a video or audio session from the session list.
2. Choose *Processing->Modify*
3. Choose Audio Properties.
4. Check the Stereo Balance box and make your slider adjustment or check the Mono box.
5. Name the modified session.
6. Click Create Session. Modified video or audio sessions will appear in Captured and Imported Sessions window. Click Done to exit the Modify panel.


DEMULPLEX

Multiplexed video is often a compilation of rapidly changing camera views. Additionally, camera views may be allocated to separate video fields, resulting in frames with mismatched, overlapping images. The Demultiplex feature uses a pixel-matching algorithm to sort through multiplexed camera views and organize them into individual video streams. After this process, these views can be reviewed, edited and saved.

To demultiplex a session

1. Select a multiplexed video from the session list.
2. Select *Processing->Demultiplex*.
3. Click Process Selected Views.

While processing, a pixel-matching algorithm compares and matches the pixels among every frame in the video. The results are displayed in the Found Views list. The number of found views are dependent on how successful the pixel matching algorithm is distinguishing one frame from another. Anomalies such as tape static, noise or camera panning can sometimes cause more views to appear than are actually present on the tape. These problems can be corrected, or compensated for, however. See [Strategies for Demultiplexing](#) for more information.



Found Views are displayed in a list. The views can be selected, reviewed in the player, and edited before being saved.

DEMULTIPLEX WINDOW

The following table describes the items in the Demultiplex window.

Demultiplex Window	
View Name	A default name is assigned to each camera view. This name can be changed if the view is saved as a Session.
Number of Frames in View	The number of video frames contained in the view.
Player Area	A video player in which individual views can reviewed and edited if needed. The video player review area contains a play button, a frame forward button, a frame backward button, mark-in and mark-out buttons and a video scrubber.
Algorithm Sensitivity Slider	Adjusts the percentage of the pixels in a frame that must match another in order to consider them from the same camera. A smaller percentage setting generally results in the creation of fewer views. A larger percentage will likely result in the creation of more views.
Reassign Frames Into Views	Takes all the frames from selected views and redistributes them into the unselected views.

To view the frames in a Demultiplex View

1. Select a view in the Demultiplex panel
2. Use the buttons and scrubber in the video player to play or scroll through the view.


REASSIGNING FRAMES AND VIEWS

Since the demultiplex feature works by a frame-matching algorithm, there may be occasional frames which are omitted from a view because of irregularities or may be assigned to the wrong view. Individual frames can be reassigned by using cut and paste methods, or drag and drop reassignment.

You can reassign individual frames, group of frames or combine entire views in together. You can also take an existing view and split it into more views by adjusting the sensitivity slider to a higher percentage value and processing it once more.

To reassign single frames by Cut/Paste

1. Display the frame in the video player area.
2. Select *Edit->Cut*.
3. Select the View you want to move it to.
4. Select *Edit->Paste*.



Frames in views can be edited by accessing the cut and paste options under the Edit menu.

Note: All frames pasted within the views will be ordered in their proper chronological sequence.

To move single frames by Drag/Drop

Individual frames can also be reassigned by using drag and drop from the video player.

1. Display the frame in the video player area.
2. Drag the selected frame from the video player area to the desired View Icon on the left.

Note: All frames dropped within the views will be ordered in their proper chronological sequence.

To move a region of frames


A group of frames can be moved from one view to another by using the video player.

1. Display the first frame of the region in the videoplayer.
2. Hit the Mark In Button
3. Display the last frame of the region in the videoplayer.
4. Hit the Mark Out Button
5. Select *Edit->Cut*
6. Select the View you want to move the region to
7. Select *Edit->Paste*

You can also combine entire views, comprised of many frames, together.

To combine entire views by Drag/Drop

1. Select a view in the Demultiplex panel.
2. Drag the view and drop it onto the view that you want to merge it with.



Consolidate views that match by dragging and dropping one on the other.

To combine entire views by Cut/Copy

1. Select the view you want to move.
2. Select *Edit->Cut*
3. Select the View you want to move it to.
4. Select *Edit->Paste*

To create New Views

Sometimes, you may want to isolate a series of frames into their own view. To do so you can cut them from an existing view and then use *Edit->Paste New*.

1. Select the View or Frames you want to move.
2. Select *Edit->Cut*
3. Select *Edit->Paste New*

REVIEW ALL FRAMES IN A VIEW

Once you have reassigned frames, you should review all the frames in the view to make sure that you have everything you need. The best way to ensure that you are seeing every frame is to play the view in the video player. If you quickly scrub through the movie with the video scrubber, you may miss misassigned frames, especially for long views.

REASSIGN FRAMES INTO UNSELECTED VIEWS

There may be times when you want to take a view or group of views and reassign them into the remaining views. This will typically happen after you have run a first demultiplex on a session. Sometimes there will be a set of views of a small number of frames at the end of the list. These were defined to be their own distinct cameras. You can force them into one of the existing views by using the reassign option.

1. Select the view(s) you want to reassign.
2. Click the check box "Reassigning frames into Unselected Views".
3. Hit the Process Selection button.

The views where these frames have been assigned will be highlighted in blue.

SAVING VIEWS

Once you have a view from which you would like to create stills, save it to the main window.

To save a view

1. Select a View in the Demultiplex panel.
2. Select *File->Save*.

To save All Views

You can also save all the views from the Demultiplex panel at once.

1. Select *File->Save All*

STRATEGIES FOR DEMULTIPLEXING

Depending on the length and quality of the video data, demultiplex results may vary. For example, since the found views are based on pixel-matching, not from decoding the original multiplex signal, more views may be found than correspond to the actual number of cameras on the tape.

After the initial processing completes, browse through the views in the list to assess how well the process worked. The views will be arranged according to the largest quantity of frames per view, in descending order.

Check for any mismatched frames in the views. Use the video player scrubber to browse through each one, playing them if necessary. Views at the bottom of the list may contain very few frames, or possibly only one frame. When this occurs, it's possible that the views are comprised of frames with a high degree of video noise, or similar anomalies, that prevented a match with other frames in the video.

Additionally, it is generally easier to demultiplex shorter videos, rather than longer ones, comprised of several minutes.


CHANGING THE SENSITIVITY OF THE MATCHING ALGORITHM

The sensitivity slider sets the percentage of data that must match between frames in order to assign them to the same view. In most cases, demultiplexing will perform well using the default slider position (the default is set to group frames according to 40% of matching pixel data). It is recommended to use the default slider setting when processing for the first time.

A smaller percentage number means that a smaller area of the frame must match

in order for the two frames to be considered the same view. A higher percentage means that a larger area of the frame must match to be considered the same view.

Moving the slider to a lower percentage number (left) usually results in fewer views. This means if frames are similar, but don't necessarily match, they may be grouped into the same view. Moving the slider to a higher percentage number (right) may result in more views, as it forces the algorithm to be more discriminating. Increased sensitivity to anomalies such as static, camera noise or panning are likely to cause more views to be found.



Adjust the sensitivity slider to set the percentage of data that must match between frames to be considered the same camera view

Varying results of the demultiplex process may appear as follows:

“TOO MANY” VIEWS ARE CREATED

The demultiplexer may create dozens of views, even though the number of actual cameras could have been limited to 5 or 10. Because the number of found views is a result of the percentage of matching pixel data in individual frames, more views are created when the demultiplexer recognizes discernable differences in that data.

There are different ways to correct this:

Manual consolidation of views

This is generally the easiest way of re-grouping views that are recognizably from the same camera. Drag and drop the views that match on top of one another until all the matching views are consolidated.

Reassign frames into existing views

The Reassign checkbox performs this process automatically, distributing frames to other existing views. The views that have been changed by this process will appear in the list highlighted in blue text. However, it may take longer for the computations to perform this task than a manual drag and drop would. It's generally more effective to use this feature with larger groups of views too difficult to consolidate manually. Once running this process, you will need to check the resulting views for consistency.

Lowering the slider setting

The demultiplexing process can be run again on the original session with a lower slider setting. This will produce fewer views, however, another effect of lowering the frame-matching sensitivity could be the presence of mismatched frames within individual views, thus prompting the need for more demultiplexing.

MISMATCHED FRAMES GROUPED TOGETHER IN VIEWS

It's possible some mismatched frames will be grouped together in views. This can be corrected by demultiplexing the newly created views themselves. Highlight the views to be demultiplexed. Move the slider further to the right, thereby increasing the sensitivity of the frame-matching algorithm, and process again. The new views will appear in the list highlighted in blue text. These views can then be combined with any other matching views in the list.


It may be simpler in some cases, when relatively few frames are misplaced, to manually drag the frames from the video player to the appropriate views in the list. Additionally, after examining the video closely, you may want to delete any views or frames that contain irrelevant data, thus simplifying the regrouping process.

VIDEO AND IMAGING FILTERS

Each of the filters listed under the Processing menu has specific function designed to improve the quality or modify videos and still images. After applying the effects of a particular filter, users can take the modified video or still image and apply additional filter passes to combine their effects.

FILTER OVERVIEW

Each filter has two displays. The left is the Input display, representing the original video or still image. The right is the Output display, representing the filtered results. When filtering videos, the input and output displays additionally function as players.



Input display is on the left, Output display is on the right. The Output display is showing a zoom and pan effect.

ZOOMING AND PANNING

The Input and Output displays in the filter panels can be zoomed by rolling the mouse wheel. (if you don't have mouse with a wheel, you can use Ctrl+ and Ctrl- on your keyboard). Clicking and dragging the mouse when an image is zoomed repositions or pans the image in the context of the display.

SYNCHRONIZING PLAYERS AND LOCKING

The arrow buttons between the players synchronize the views between the input and output displays. Which is synchronized to which depends on the direction of the arrow clicked. This also locks the input and output players to keep them synchronized. Clicking on the unlock button will decouple the two players



Clicking the top arrow syncs the Input display to the Output display, and the bottom arrow, vice-versa.

FILTER TOOLS

The tools in the upper right corner allow the user to select a sub-portion of any frame as a region of interest to apply the filter. If the Tracking button is activated, and the area under the region of interest contains motion, the filter attempts to track it, and the region will follow it as it moves frame to frame. If no region is selected, the filter is applied to the entire frame area.

Once any filter adjustments have been made, the video, or individual frames, can be immediately saved, by clicking the respective “Save Image” or “Save Video” buttons.

GRAYSCALE FILTER

The Grayscale filter converts color video or images to grayscale mode. Users can choose to do a straight RGB conversion, or designate which color channel (Red, Green or Blue) to convert through the drop-down menu. The results that looks best will depend on the color values in the original video or image.

To use the Grayscale filter

1. Select a session or still image.
2. Choose *Processing->Filter->Video->Grayscale*
3. Choose Grayscale or a different color channel from the drop-down menu.
4. Name the modified session or still image.
5. Click Save.
6. Close to exit the filter

SHARPEN FILTER

The Sharpen filter increases the level of definition in the edges of the entire image, or a region of interest. It can improve image quality, but it can't restore clarity to severely blurred images or video.

To use the Sharpen filter

1. Select a session or still image.
2. Choose *Processing->Filter->Video->Sharpen*
3. Make adjustments to the Weighting Value slider.
4. Name the modified session or still image.
5. Click Save.
6. Close to exit the filter.

EQUALIZE FILTER

The Equalize filter finds the brightest and darkest values in the entire image, or the selected region of interest, then remaps them so the brightest value represents white and the darkest value represents black. The intermediate values are then equalized evenly throughout the grayscale. Though it may yield superior image quality, especially in the case of dark video or still images, it may not necessarily make smooth or natural looking adjustments to the brightness and contrast levels in a session or still image.

To use the Equalize filter

1. Select a session or still image.
2. Choose *Processing->Filter->Video->Equalize*
3. Adjustments are made automatically to the entire image or the selected region of interest.
4. Name the modified session or still image.
5. Click Save.
6. Close to exit the filter.

BLUR FILTER

The blur filter softens selected pixel areas of the entire image, or a selected region of interest, reducing detail. Combined with the built-in tracking capability, it can be used to obscure faces in videos (if there is not excessive motion).

To use the Blur filter

1. Select a session or still image.
2. Choose *Processing->Filter->Video->Blur*
3. Make adjustments to the Pixel Radius slider.
4. Name the modified session or still image.
5. Click Save.
6. Close to exit the filter.

STABILIZE FILTER (VFPRO ONLY)

The Stabilize filter minimizes the shakiness that can result from hand-held video. It works by utilizing an algorithm that considers parameters of Translation and Rotation throughout a series of video frames.

To use the Stabilize filter

1. Select a session.
2. Choose *Processing->Filter->Video->Stabilize*
3. Make adjustments to the Translation and Rotation sliders
4. Test the effectiveness of the filter by playing the sequence, making further adjustments, if necessary.
5. Name the modified session.
6. Click Save.
7. Close to exit the filter.

LEVELS ADJUST FILTER

The Levels Adjust filter allows you to make adjustments to the brightness, contrast, and mid-tones of a session or a still image. When the filter is opened, two histograms are displayed. The Input Range histogram provides a picture of the pixel value distribution in the original session or still. The Output Range histogram changes reflecting any slider adjustments.

To use the Levels Adjust filter

1. Select a session or still image.
2. Choose *Processing->Filter->Video->Levels Adjust*

3. Make adjustments to the Input, Output and Gamma sliders.
4. Test the effectiveness of the filter by playing the sequence, making further adjustments, if necessary.
5. Name the modified session.
6. Click Save.
7. Close to exit the filter.

SUPER-RESOLUTION FILTER (VFPRO ONLY)

Super-resolution is a technique that combines multiple frames of video and can thereby eliminate extraneous noise and produce images with a greater resolution than the original video frames from which they came. The resulting stills can then be exported in BMP, TIFF or JPEG format.

To create a Super-resolution still

1. Select a frame from a video session. This will be the frame for your still.
2. Choose *Processing->Filter->Video->Super-resolution*.
3. Draw mask(s), if needed.
4. Adjust the number of Frames Used and Noise Reduction, if needed.
5. Click the Save Image button.
6. VideoFOCUS creates the still and places it in the Session Stills and Frames area.

STILL CREATION OVERVIEW

The Super-resolution process, though it can create superior looking stills from video, cannot extract details that are not present in the video or otherwise invisible to the naked eye.

The process of creating super-resolution stills involves separate steps.

- An alignment step which "lines up" sequential frames.
- A rendering step that uses the information from those aligned frames to create a cleaner, higher resolution image.

Four important aspects of a video sequence impact the creation process: the Frames Used, Motion Correction, Noise Reduction, and De-interlacing.

FRAMES USED

The "frames used" slider controls the number of frames in the neighborhood. When creating a still, information is taken from every frame in the neighborhood. The displayed center frame is combined, pixel by pixel, with the frames in the neighborhood on either side. This combination can correct off-color pixels caused by static or other video defects and depending on the content, motion and quality of the source video, boost the resolution of the output image.

By default, the frame neighborhood includes the center frame and two surrounding frames (3 frames used). An increased size frame neighborhood contains more image information and more opportunities to improve image quality. However, if the video changes a great deal from frame to frame, a large frame neighborhood may result in a blurry image.

You should avoid including frames if they contain:

- A scene change
- Sudden motion, either by the camera or the subject
- Major video defects

To adjust the number of frames used by Super-resolution

Adjust the "Frames Used" slider to select the number of frames. This slider is logarithmic and will only supply an odd number of frames (to ensure you always have a center frame).

MOTION CORRECTION

Virtually all video sequences contain some movement, whether it is movement of the subject being filmed, or camera effects like zooming and panning. With the "Apply Motion Correction" checkbox enabled, the Super-resolution filter will attempt to analyze and align the frames to compensate for the motion.

Leaving this control unchecked will give no motion correction and the resulting image will simply be an image that has frame averaging applied.

NOISE REDUCTION

The Noise Reduction slider is only enabled when Motion Correction is checked. Increasing the Noise Reduction slider will result in a smoother, less noisy image, but will possibly have reduced resolution. Decreasing the slider will result in a sharper image with more resolution, but could increase speckle noise and even induce dropouts in the image.

DE-INTERLACING

Primarily, interlacing is only encountered while working with analog video. Interlaced video frames are comprised of two fields - an "even" field having only the even-numbered scan lines of the image, and an "odd" field having only the odd-numbered scan lines of the image. Each field is scanned at a different moment in time; consequently, a single video frame displays two separate moments in time overlaid.

NTSC, the video standard in North America and Japan, captures 30 video frames per second. Therefore, it records 60 fields per second, or 60 images. Because the even field of each frame is scanned first, NTSC is "even-field dominant". Each image consists of only half of the scan lines of the complete picture at a given time, as illustrated below (see Figure 1)



Figure 1 - A moving ball with the fields separated.

When the fields are combined into their respective frames, an object may look like it is in two places at the same time. This happens when there is a good deal of motion in a video sequence. (see Figures 2 and 3)




Figure 2 - The moving ball with the fields combined into their respective frames.



Figure 3 - A moving car displaying the same interlace effect.

De-interlacing is the act of separating the even and odd fields into their own frames, thereby producing a video at twice the frame rate of the original but half the line resolution. This process can produce a more satisfactory still image but may, as a result, introduce a jagged, “stairstep” look to diagonal lines.

When a session is marked as Interlaced, the Super-resolution filter will automatically split the fields into frames before processing. This method separates the video fields and treats them as “frames” (the representative number of frames will effectively double, however, each frame will consist of only one field). See the section on the Modify menu option [Split Fields into Frames](#) for more information.

CREATING MASKS

Video occasionally displays superimposed text or graphics. Additionally, sequences may consist of moving objects or exhibit a large amount of camera motion. These conditions can interfere when aligning video frames. Masking lets you mask out unimportant areas to exclude them from the frame alignment computation (they will appear in the final still, but can be cropped out using the Crop Frame feature in the Modify menu).


Masks in the Super-Resolution dialog are *inverted masks* - meaning, the mask is applied outside the defined area.

Items you should typically mask out include subtitles, logos or other superimposed graphics, or objects that move differently from the subject in the video.

The mask additionally tells the Super-Resolution filter where to apply the process. Only objects contained within the mask area will be processed.

To draw masks

1. In the Super-resolution dialog box, click on one of the drawing tools in the upper- left of the dialog.
2. Use the draw tools to draw mask objects that define the inverted masked areas on the left-hand image (the original frame). When initially drawn or subsequently selected by the user, a masked area is represented by an orange outline. When not selected, the outline color changes to green.



3. If the mask needs to be moved, simply place the mouse cursor within the mask area and drag it by holding down the left button. To resize the mask, place the mouse at either corner and drag as needed to increase or decrease the size.

To delete masks

When a mask object is selected (outline is orange instead of green), you can delete it by clicking on the "X" button.


MANAGE VIEWS

When media is imported into VideoFOCUS, the application will attempt to determine the time stamp of the media. If the media does not contain a time stamp, the file's "Last Modified" time stamp is used. This time stamp is displayed in the lower right corner of the video player and can be used to time-align multiple videos.

The Manage Views window allows you to adjust the time stamp of a video and can be used on either a single view session or a MultiView session.

To adjust the time stamp of a single session

1. Select the session you wish to adjust
2. Select *Processing->Manage Views*
3. Adjust the play head to a position where the time and date stamp is known
4. Adjust the Date, Time, and Time Region under Frame Timing to match the known time stamp
5. Click Save to save the new, adjusted session




To adjust the time stamp of a MultiView session

1. Select the MultiView session you wish to adjust
2. Select *Processing->Manage Views*
3. Select the view you wish to adjust by clicking on the view. An orange box will appear around the view. If you wish to adjust all views, either do not select any view or select all views by typing Ctrl+A. To select more than one view, but not all views, click each desired view while holding down the Ctrl key.
4. The time range of the selected view(s) will show up in the playhead slider area

as an orange bar. Adjust the playhead to a position where the time stamp is known. **The playhead must be placed somewhere along this orange bar to be within the current timeframe of the selected view and enable the date/time for that view to be adjusted.**

5. Adjust the Date, Time, and Time Region under Frame Timing to match the known time stamp
6. Repeat steps 3-5 for other views as necessary
7. Click Save to save the new, time-aligned MultiView session

Frequently with DVR imports that consist of segmented files, the same camera view will show up as separate views typically because they were imported as separate files. These views can easily be joined together into one view.



To join multiple views into one view

1. Select the MultiView session you wish to adjust
2. Select *Processing->Manage Views*
3. Select the views you wish to join by clicking on the views using the Ctrl key. An orange box will appear around the selected views.
4. Click the Join button. The views will collapse into a single view. If their time stamps overlap, the video that starts with the later timestamp will be the one that is displayed. If there is a gap in time between the views, a black frame will be inserted for the duration of the gap.
5. If you wish to undo a Join, select the joined view and click Split. *This can only be done in the same Manage Views session that created the Join. Once the new session is saved, a Split can no longer be performed.*
6. Click Save to save the new MultiView session.

COMBINE (VFPRO ONLY)

Separate sessions can be combined into a single MultiView session to view them simultaneously in a time-aligned fashion. If there are gaps in time between two views, this time will be collapsed and a time offset overlay will appear to indicate the duration between the current playhead position and the start/end of the view.

To combine multiple sessions into a MultiView session

1. Select the desired sessions in the session list by clicking each one with the Ctrl key pressed.
2. Select *Processing->Combine* and a new MultiView sessions will be created.
3. Apply *Manage Views* if needed to perform any time adjustments and/or joins.

MACRO PROCESSING (VFPRO ONLY)

Complex sequences of filters can easily be copied and applied to other sessions by means of the *Copy Macro* and *Apply Macro* menu functions.

To replicate all processing steps applied to one session to another session


1. Select the session having all of the desired processing steps by clicking it in the session list.
2. Right-click over the session and select *Copy Macro Starting At* in the pop-up menu, then choose the bottom-most session (which would be the original import session) from the dropdown menu at the right.
3. Select the other session(s) to which the macro is to be applied by clicking the first one, then holding down the Ctrl key and clicking any additional sessions if desired.
4. Right-click over one of the sessions and choose *Apply Macro* from the pop-up menu to apply the same processing to these sessions.
5. Observe that new child sessions have been created with the same macro applied to each one.

AUDITING SESSIONS AND STILL IMAGES


As new versions of videos and images are created by the processing, VideoFOCUS keeps track of all operations made along the way, and these can be easily viewed in an expandable Audit tree. The Audit tree provides a machine-level representation of all steps leading up to the present state of the selected session or still. By following the noted operations, the user has the basic steps needed to recreate or demonstrate the changes, if required.

To view an Audit tree of a session or still image

1. Select any session or still image.
2. Choose *View->View Audit*.



This Audit tree indicates a Levels Adjust filter was applied to create this session, and provides the precise settings needed to recreate it.



CHAPTER 6

PROCESSING AUDIO

Any audio in VideoFOCUS can be processed. Processing can be applied to an entire audio session.

REDACT AUDIO

Redact Audio can be used to suppress individual words or other sounds that you do not wish to be audible in the exported media file.

To redact an individual word or other sound in a session

1. Turn on the audio signal visualization in the player by right-clicking over the player and selecting the *Audio Signal* option in the pop-up menu.
2. Scrub to either the beginning or the end of the individual word or sound, then hold down the Shift key and scrub to the opposite end of the segment. Note the orange bar that now appears on the scrubber bar, as well as the orange segment of the audio waveform, indicating the temporal selection.
3. Right-click over the player and choose *Redact Audio* from the pop-up menu, then click either *Silence* or *Tone* in the dropdown menu to make the redaction.
4. Selected segment of the audio waveform should now be replaced with either silence or tone as specified.

EXTRACT AUDIO

Extract Audio simply extracts the audio from a video session and saves it as a separate Audio session.

To extract the audio track of a video session to a separate audio session

1. Select a video session with audio.
2. Choose *Processing->Extract Audio*.
3. A new audio session will be created in your session list.

COMBINE AUDIO AND VIDEO


VideoFOCUS allows you to combine two audio sessions or combine an audio session with a video session.

To combine an audio session with another audio or video session

1. Select two audio session, or select an audio session and a video session.
2. Choose *Processing->Combine*.
3. A new combine session will be created in your session list.

BASIC AUDIO FILTERS

NOISE REDUCTION



The Noise Reduction filter is used to reduce both Predictable and Random broadband noises occupying the same wide range of frequencies as desired voice signals. The Predictable stage employs an adaptive algorithm that automatically recognizes consistent temporal patterns in the audio signal to attenuate repetitive noise signals such as motors, fans, and engines, plus time-based acoustical effects such as reverberation. Conversely, the Random stage utilizes a spectral subtraction algorithm that automatically measures the frequency pattern of any noise signals present during pauses in the speech, with this pattern then subtracted from the audio signal at all times to provide a general “de-noising”; this algorithm is especially effective against persistent Gaussian, or random, noises such as hiss and static. With either stage, the degree of noise Attenuation can be adjusted to any value between 0% (no effect) and 100% (maximum effect), to tailor the tradeoff between sound quality and noise reduction in the resulting output audio.


The 200Hz highpass filter (HPF) is used to reduce low-frequency noises. The filter attenuates all sound energy below 200Hz, which is considered to contain no signals contributing to speech intelligibility, and leaves all higher frequencies unaffected.

To use the Noise Reduction filter

1. Select a video session with audio or select an audio session.
2. Choose *Processing->Filter->Audio->Noise Reduction*.
3. Make adjustments to the Predictable and/or the Random sliders.
4. Test the effectiveness of the filter by playing the sequence, making further adjustments, if necessary.
5. You can also click on the Render button to view the output audio wave.

6. Name the modified session.
7. Click Save.
8. Close to exit the filter.

VOICE SPECTRUM EQUALIZATION




The Voice Spectrum Equalizer is used to reshape the spectrum of the output audio. The sliders allow attenuation of the audio energy in a series of adjacent frequency bands, with each band centered at the specified frequency. The lower the slider bar is moved, the more attenuation is applied in that frequency band. The Preset dropdown box can be used to quickly adjust the slider array as a whole to one of three presets - Allpass, lowpass or highpass.

To use the Voice Spectrum Equalization filter

1. Select a video session with audio or select an audio session.
2. Choose *Processing->Filter->Audio->Voice Spectrum Equalization*.
3. Make adjustments to the Frequency sliders.
4. Test the effectiveness of the filter by playing the sequence, making further adjustments, if necessary.
5. Name the modified session.
6. Click Save.
7. Close to exit the filter.

HUM REDUCTION



The Hum Reduction filter employs a special comb filtering algorithm that attenuates harmonically-structured noises such as 50/60Hz electrical mains hum. Such hum can occur on any recording device, analog or digital, that is AC-powered, and is typically the result of close proximity of the microphone and audio circuitry to power s, fluorescent lighting, or other equipment that may be emanating strong electromagnetic fields. The filter's fundamental frequency can be manually adjusted via the Frequency control; alternatively, the Auto-Tracking feature can be enabled to automatically identify the frequency of any hum present in the incoming audio and adjust the Frequency control accordingly. Precise adjustment of the Frequency setting is critical, as recorder speed and audio sampling rate variations often result in the embedded hum varying slightly from the nominal 50/60Hz frequency. Additionally, the degree of noise Attenuation can be adjusted to any value between 0% (no effect) and 100% (maximum effect), to tailor the tradeoff between sound quality and hum reduction in the resulting output audio


The 200Hz highpass filter (HPF) is used to reduce low-frequency noises. The filter attenuates all sound energy below 200Hz, which is considered to contain no signals contributing to speech intelligibility, and leaves all higher frequencies unaffected.

To use the Hum Reduction filter

1. Select a video session with audio or select an audio session.
2. Choose *Processing->Filter->Audio->Hum Reduction*.
3. Make adjustments to the slider and check the desired boxes.
4. Test the effectiveness of the filter by playing the sequence, making further adjustments, if necessary.
5. You can also click on the Render button to view the output audio wave.

6. Name the modified session.
7. Click Save.
8. Close to exit the filter.

AUTOMATIC GAIN




The AGC (Automatic Gain Control) is a dynamic level processor that should only be applied once all other audio processing is complete. The AGC applies gain as needed to maintain a good audio level (approximately -12dB). The gain applied by the AGC is limited by the Max Gain setting; regardless of the audio level, no gain greater than Max Gain will be applied. The Release Time controls how quickly the AGC will respond to decreases in input signal level. The shorter the Release Time, the more quickly the AGC will react.

To use the Automatic Gain filter

1. Select a video session with audio or select an audio session.
2. Choose *Processing->Filter->Audio->Automatic Gain*.
3. Make adjustments to the sliders.
4. Test the effectiveness of the filter by playing the sequence, making further adjustments, if necessary.
5. Name the modified session.
6. Click Save. Close to exit the filter.

QUICKENHANCE




QuickEnhance is an easy-to-use audio processing plug-in designed for use in environments that support the AAX or VST plug-ins. QuickEnhance provides a set of speech clarification tools, primarily for use in forensic applications. QuickEnhance comes equipped with the following features:

- Auto-Normalizer
- 200 Hz High-pass Filter
- 1000-point Spectrum Analyzer (with maximum peak indicator)
- Hum Filter (adjustable base frequency with Auto-Tracking feature)
- Predictable Noise Reduction Filter (Deconvolver)
- Random Noise Reduction Filter (Broadband filter)
- 20-Band Graphic Equalizer

- Automatic Gain Control

To use the QuickEnhance

1. Select a video session with audio or select an audio session.
2. Choose *Processing->Filter->Audio->QuickEnhance*.
3. Click on Settings and the QuickEnhance Dialog will appear, make necessary adjustments.
4. Test the effectiveness of the filter by playing the sequence on the output player, making further adjustments, if necessary.
5. You can also click on the Render button to view the output audio wave.
6. Name the modified session.
7. Click Save. Close to exit the filter.



ADVANCED AUDIO FILTERS - CARDINAL MINILAB (VFPRO ONLY)

MINILAB INTRODUCTION

The entire suite of CARDINAL MiniLab plug-ins includes the following:


- 1-CH Adaptive Filter
- Comb Filter
- Gain Stage
- Graphic Equalizer
- Multi-Band Filter
- Parametric Equalizer
- Reference Cancellor Filter
- Spectral Inverse Filter
- Spectral Subtraction Filter
- X-Pass Filter (6 filter set)
- Filter Chain (runs combinations of any of the above)

These plug-ins are all available through the Filter Chain interface in VideoFOCUS. The Filter Chain plug-in provides the ability to put up to 8 MiniLab plug-ins in series and render the cumulative effect in one pass of the audio.

The CARDINAL MiniLab suite of plug-ins can be accessed in VideoFOCUS by selecting the Filter Chain option in the Audio Filter menu.

GENERAL PLUG-IN CONCEPTS

SPECTRUM ANALYZER



Spectrum Analyzer

Most plug-ins have an integrated spectrum analyzer to display frequency content of the audio. Most will look similar to the figure above.

The **Spectrum Analyzer** displays the frequency content of the input and/or output audio. The frequency axis of the **Spectrum Analyzer** plot goes from left to right, with the lowest frequency at the leftmost side and the highest frequency at the rightmost side. The "energy" (loudness) axis of the **Spectrum Analyzer** plot

goes from top to bottom, with stronger (louder) frequencies indicated by higher peaks on the plot.

For example, a signal consisting of a single tone will appear as a single peak, located at the tone frequency. A "white noise" signal, which contains equal amounts of every frequency in the signal bandwidth, will appear as an (approximately) flat line across the entire frequency range.

The spectrum analyzer can display the **Input** audio, the **Output** audio, or both. To enable the **Input** and/or **Output** trace, click the corresponding LED button so that the indicator "light" is on. The **Input** audio is the signal before the filter is applied. The **Output** audio is the signal after the filter is applied. When both signals are shown, each signal is indicated by a line, using yellow for the **Input** and blue for the **Output**.

The **Spectrum Analyzer** provides a cursor marker to help in identifying specific frequency values. At any time when the **Spectrum Analyzer** is activated, clicking in the graph area displays a red vertical marker at the frequency location clicked. The frequency (in Hz) of the location clicked is displayed in red text on the bottom left of the graph. The input level (in dB) is given in yellow text and the output level (in dB) is given in blue text.

The **Track Peak** feature displays a green vertical line in the graph display at the strongest (loudest) frequency. When this feature is enabled, the maximum peak value (in dB) and frequency (in Hz) is displayed in green text on the bottom right of the graph. To enable this indicator, click the button beside the **Track Peak** text so that the indicator "light" turns green.

Use the **Averager Mode** control to select what type of averaging is done to the signals displayed. **No Average** applies no averaging and the traces will move very rapidly. **Exponential** applies averaging to the trace values based on the **Num Averages** value. **Peak Hold** will hold the maximum value for each data point of the trace. Trace data will never leak down in this mode. This is useful for determining peak energy for a given frequency.


Use the **Num Averages** control to increase or decrease the number of averages applied to the frequency spectrum. A small **Num Averages** value allows a more accurate snapshot of the spectrum at a given time, but the trace values will change rapidly, making longer-term spectral characteristics difficult to see. Larger **Num Averages** values result in a smoother spectral plot that represents stable frequency characteristics well but does not accurately show rapidly time-varying signal characteristics. This control only has an effect when **Exponential** is selected for the **Averager Mode**.

The **Channel Select** control allows you to choose which channel of audio you wish to see displayed in the spectrum analyzer window. You can choose **Left Only**, **Right Only** or **Mix**. The **Mix** selection sums the two audio channels together before performing the FFT analysis. **Right Only** and **Mix** are only available for stereo

audio files.

The analyzer also has three **Zoom** controls: **Zoom In**, **Zoom Out** and **Zoom Full**. When zooming, the graph will always try to center the graph on wherever the red user marker is located. Each level of zoom halves the current frequency span of the display. For instance, if the current sample rate is 44100Hz, the first level of zoom will cause the analyzer to display 11025Hz instead of 22050Hz worth of spectral data. **Zoom Full** will always cause the analyzer to reset back to the full bandwidth of the signal.

COEFFICIENT DISPLAY



Coefficient Display

A few filters contain a Coefficient Display, particularly the 1-CH Adaptive Filter and Reference Canceller Filter. This graph can be useful in setting up these filters.

The Coefficient Display displays the impulse response (filter coefficients) of the filter. Vertical scaling of the filter's coefficients for display is accomplished by clicking on the **Vertical Zoom** combo box control. Supported zoom factors range from **1x** to **32x**.

The **Track Peak** feature displays a green vertical line in the graph display at the coefficient with the largest (absolute) value. When this feature is enabled, the maximum peak value and coefficient number is displayed in green text on the bottom right of the graph. To enable this indicator, click the button beside the **Track Peak** text so that the indicator "light" turns green.

The **Channel Select** control allows you to choose which channel's filter coefficients you wish to see displayed in the coefficient display window. You can choose **Left Only**, or **Right Only**. **Right Only** is only available for stereo audio files.

SAMPLE RATE CONVERSION



Sample Rate Conversion

All CARDINAL MiniLab plug-ins contain a resampling feature that utilizes a true Whittaker-Shannon Interpolation to resample the incoming audio at 16kHz. This method eliminates any distortion components commonly found in cruder resampling techniques (such as linear interpolation).

Resampling the audio down at 16kHz sample rate primarily does two things for forensic audio:

9. Bandlimits the audio to the forensic voice spectrum (200 - 5000Hz)

10. Allows for larger filter sizes with less computational requirements. For example, a 1024-point filter at 48kHz will require roughly 3x the CPU resources that the same 1024-point filter will require at 16kHz.


Note: After processing, the plug-in resamples the audio back to the original sample rate to return it to the host editor environment. If you process a 44.1kHz WAV file, you will still have a 44.1kHz WAV file in the end, but any information above 8kHz will have been eliminated.

FILTER CHAIN

OVERVIEW OF THE FILTER CHAIN

The Filter Chain plug-in provides the ability to put up to 8 MiniLab plug-ins in series and render the cumulative effect in one pass of the audio. Normally, each plug-in must be brought up separately in the editor and rendered on the audio independent from one another. This process can be time consuming and inefficient if you wish to apply more than one plug-in. With the Filter Chain, it is possible to create a chain of plug-ins, modify their settings and listen to their cumulative effect before rendering the audio just once.

The spectrum analyzer displays the unmodified input from the audio editor and the total resulting output audio after the last stage of the filter chain.




Filter Chain Window

To use the FilterChain

1. Select a video session with audio or select an audio session.
2. Choose *Processing->Filter->Audio>FilterChain*.
3. Choose from 10 different MiniLab filters from the drop downs, edit and make the necessary changes.
4. Test the effectiveness of the filter by playing the sequence on the output player, making further adjustments, if necessary.
5. You can also click on the Render button to view the output audio wave.
6. Name the modified session.
7. Click Save. Close to exit the filter.

FILTER CHAIN CONTROLS

Filter Select



Filter Select Controls

The Filter Select controls are a set of 8 combo boxes that allow you to select from all the available CARDINAL MiniLab plug-ins that are installed on your system. Some filters are only available if the audio presented matches their criteria (e.g. Reference Cancellor is only available for stereo audio files).

The button next to each filter select control enables the filter. The audio always runs through each filter, however if the filter is disabled, its audio output is discarded. This allows adaptive filters to keep their solution in sync with the audio playback, but prevents the output of the filter to continue down the chain.

Filter Control Button Group



Filter Control Group

The Filter Control Group of buttons allows the user to move, delete and edit the corresponding plug-in.

The Move button  will shift the plug-in up and down in the chain without altering its settings.

The Delete button  will remove the corresponding plug-in from the chain and shift any filters below it up to fill the gap.

The Insert button  will insert a blank filter slot and shift the existing plug-ins down (or up, if no room).

The Edit button  brings up the graphical interface for the corresponding plug-in. You can have multiple plug-in windows open at one time.

Bypass All



Bypass All

The **Bypass All** control will disable all the filters in the chain. This is useful to easily hear before and after the filter chain processing to see if you are achieving the desired result.

Store Button



Store Button

The **Store** button will save the entire filter chain configuration, including all filter parameters, into a Filter Chain Settings (FCS) file for recall later. These files can also be used by the to apply the settings to a large number of audio files in batch mode.

Recall Button



Recall Button

The **Recall** button will restore the entire filter chain configuration from a Filter Chain Settings (FCS) file that was preserved using the **Store** button. Using this feature will wipe out the current Filter Chain configuration.

Clear Button



Clear Button

The **Clear** button will clear all the adaptive filters in the chain. It does not remove any filters from the chain.

1CH ADAPTIVE FILTER


OVERVIEW OF THE 1CH ADAPTIVE FILTER

The 1CH Adaptive filter is used to automatically cancel predictable and convolutional noises from the input audio. Predictable noises include tones, hum, buzz, engine/motor noise, and, to some degree, music. Convolutional noises include echoes, reverberations, and room acoustics.

The 1CH Adaptive Filter is a forensic plug-in that has been specifically designed to work with VST host audio editing systems.

The 1CH Adaptive Filter comes equipped with the following features:


- Adjustable Filter Size (up to 8192 taps)
- Adjustable Prediction Span (up to 32768 samples)
- Adjustable Adapt Rate (with optional Auto-Normalize feature)
- Conditional Adaptation
- Selectable Filter Output
- Coefficient Display (with max peak indicator)
- Dual-trace spectral analysis (with max peak indicator and averaging)



1CH Adaptive Filter Main Window

FILTER CONTROLS

PREDICTION SPAN



Prediction Span

Sets the number of samples in the prediction span delay line. Prediction span is indicated both in samples and in milliseconds. Shorter prediction spans allow maximum noise removal, while longer prediction spans preserve voice naturalness and quality. *A prediction span of 2 or 3 samples is normally recommended.*

FILTER SIZE




Filter Size

Sets the number of FIR filter taps in the adaptive filter. Filter size is indicated both in taps (filter order) and in milliseconds. The maximum filter size is 8192 taps. Small filters are most effective with simple noises such as tones and music. Larger filters should be used with complex noises such as severe reverberations and raspy power hums. *A nominal filter size of 512 to 1024 taps is a good overall general recommendation.*

CAUTION: Large filter sizes (> 2048 taps) will require large computing resources to maintain real-time audio processing. You may begin to hear skips in the audio during preview if your computer cannot keep up with the processing requirements. However, during render the audio will not contain any skips but may take longer than real-time to process the file.

ADAPT RATE




Adapt Rate

Used to set the rate at which the adaptive filter adapts to changing signal conditions. An adapt rate of 1 provides very slow adaptation, while an adapt rate of 5884 provides fastest adaptation. *A good approach is to start with an adapt rate of approximately 100-200 to establish convergence, and then back off to a smaller value to maintain cancellation.*

Larger adapt rates should be used with changing noises such as music; whereas, smaller adapt rates are acceptable for stable tones and reverberations. Larger adapt rates sometimes affect voice quality, as the filter may attack sustained vowel sounds.

When Auto Normalize is turned on, the specified adapt rate is continuously scaled based upon the input signal level. This scaling generally results in faster filter convergence without greatly increasing the risk of a filter crash. *It is recommended that Auto Normalize be enabled for most speech signal processing.*

CONDITIONAL ADAPTATION




Conditional Adaptation

For advanced users only. Novice users should keep Conditional Adaptation set to **Always**. The threshold setting has no effect when **Always** is selected.

Conditional Adaptation allows the adaptive filter to automatically Adapt/Freeze based upon signal bargraph levels. This can be very useful in situations where there are pauses or breaks in the speech being processed.

Hint: Conditional adaptation is useful for maintaining adaptation once the filter has converged. Recording environment factors such as air temperature and motion in the room can cause the signal characteristics to change over the course of a recording. For this reason, simply freezing the filter once convergence is reached may mean that noise cancellation will degrade over time. Instead of freezing the filter, use Conditional Adaptation. First allow the filter to converge in **Always** mode, and then select **If Normal Output < Threshold**, and adjust the threshold by observing the bargraph levels during pauses in speech

Click on the **Clear** button if you desire the filter to completely readapt based upon the new Conditional Adaptation settings.



Filter Output

Used to optionally listen to the "rejected" audio that is being cancelled by the adaptive filter. *Normal should almost always be selected*, but the Predicted setting can be useful when configuring the filter, allowing the user to hear exactly what is being removed by the filter.

ADAPT BUTTON

Used to enable or disable filter adaptation. When Adapt is on, the filter adapts according to its settings. When Adapt is off, the filter freezes at its current state and will not adapt any further regardless of the other settings.

CLEAR FILTER

Used to reset the coefficients of the 1CH Adaptive Filter. Clearing a filter is

useful when the audio characteristics change dramatically, so that the filter can readapt to a new, clean solution. Clearing is also useful in the case of a filter “crash,” when the filter coefficients diverge to an unstable state, usually in response to a large and abrupt change in the signal coupled with a fast adapt rate.

REFERENCE CANCELLER


OVERVIEW OF THE REFERENCE CANCELLER FILTER

The Reference Cancellor adaptive filter is used to automatically cancel from the Primary channel any audio which matches the Reference channel. For example, the Primary channel may be microphone audio with desired voices masked by radio or TV noise. The radio/TV interference can be cancelled in real-time if the original broadcast audio, usually available from a second receiver, is simultaneously recorded to the Reference channel.

The Reference Cancellor Filter is a forensic plug-in that has been specifically designed to work with VST host audio editing systems. Since it is an inherently two- channel filter, it is only available for stereo audio files. The Primary channel **MUST** be on the left channel, while the Reference channel consequentially **MUST** be on the right channel. Future updates of this plug-in will allow these channels to be switched if desired.

The Reference Cancellor Filter comes equipped with the following features:


- Adjustable Filter Size (up to 8192 taps)
- Adjustable Delay (up to 32768 samples, either Pri or Ref input)
- Adjustable Adapt Rate (with optional Auto-Normalize feature)
- Conditional Adaptation
- Adjustable Reference Gain
- Selectable Filter Output
- Coefficient Display (with max peak indicator)
- Dual-trace spectral analysis (with max peak indicator and averaging)



Reference Canceller Filter Main Window

FILTER CONTROLS

DELAY




Delay

Sets the number of audio samples by which the selected channel should be delayed. Adjusting the Delay allows the alignment of the Primary and Reference channels to be adjusted. Minimum Delay is 1 sample, but can be set to as high as 32768 samples.

Specifies whether the delay line is to go into either the Primary channel (**Delay Pri**) or the Reference channel (**Delay Ref**). *For most applications, a slight delay (typically 5 msec) is placed in the Primary channel*, For applications with long distances between the microphone and radio/TV, a delay in the Reference channel may be required. Extreme caution should be exercised when using reference channel delay; allowing the reference to lag the target noise in the primary signal will result in poor cancellation.

FILTER SIZE




Filter Size

Sets the number of filter taps in the adaptive filter. Filter size is indicated both in taps (filter order) and in milliseconds. The maximum filter size is 8192 taps. *Normally, larger filters sizes are used in the Reference Canceller adaptive filter.*

CAUTION: Large filter sizes (> 2048 taps) will require large computing resources to maintain real-time audio processing. You may begin to hear skips in the audio during preview if your computer cannot keep up with the processing requirements. However, during render the audio will not contain any skips but may take longer than real-time to process the file.

ADAPT RATE




Adapt Rate

Used to set the rate at which the adaptive filter adapts to changing signal conditions. An adapt rate of 1 provides very slow adaptation, while an adapt rate of 5884 provides fastest adaptation. *A good approach is to start with an adapt rate of approximately 100-200 to establish convergence, and then back off to a smaller value to maintain cancellation.*

When Auto Normalize is turned on, the specified adapt rate is continuously scaled based upon the input signal level. This scaling generally results in faster filter convergence without greatly increasing the risk of a filter crash. *It is recommended that Auto Normalize be enabled for most speech signal processing.*

CONDITIONAL ADAPTATION



Conditional Adaptation


For advanced users only. Novice users should keep Conditional Adaptation set to **Always**. The threshold setting has no effect when **Always** is selected.

Conditional Adaptation allows the adaptive filter to automatically Adapt/Freeze based upon signal bargraph levels. This can be very useful in situations where there are pauses or breaks in the speech being processed.

Hint: Conditional adaptation is useful for maintaining adaptation once the filter has converged. Recording environment factors such as air temperature and motion in the room can cause the signal characteristics to change over the course of a recording. For this reason, simply freezing the filter once convergence is reached may mean that noise cancellation will degrade over time. Instead of freezing the filter, use Conditional Adaptation. First allow the filter to converge in **Always** mode, and then select **If Normal Output < Threshold**, and adjust the **Threshold** by observing the bargraph levels during pauses in speech

Click on the **Clear** button if you desire the filter to completely readapt based upon the new Conditional Adaptation settings.


REFERENCE GAIN



Reference Gain

Ref Gain is used to add gain to the reference channel audio if necessary. *To achieve good cancellation, it is important that the reference audio be at least as loud as the noise it is intended to cancel from the primary audio.*

FILTER OUTPUT



Filter Output

Used to optionally listen to various audio outputs in order to optimize the filter.

- **Normal** - Listen to the resulting primary channel of audio after the reference audio has been cancelled. Normal should almost always be selected before rendering the final audio.
- **Rejected** - Listen to the audio that is being subtracted from the primary channel of audio. This is useful to determine if any target voice is being accidentally targeting and removed.

- Primary - Listen to the raw input primary channel of audio. This is useful to make before and after comparisons.
- Reference - Listen to the raw input reference channel of audio. This is useful to hear precisely what is contained in the reference audio.

In all cases, whatever output is selected, it will be placed on both the left and right channels of output.

FREEZE FILTER

Used to enable or disable filter adaptation. When Freeze is off, the filter adapts according to its settings. When Freeze is on, the filter never adapts regardless of the other settings.

CLEAR FILTER

Used to reset the coefficients of the Reference Canceller Filter. Clearing a filter is useful when the audio characteristics change dramatically, so that the filter can readapt to a new, clean solution. Clearing is also useful in the case of a filter “crash,” when the filter coefficients diverge to an unstable state, usually in response to a large and abrupt change in the signal coupled with a fast adapt rate.

SPECTRAL INVERSE FILTER

OVERVIEW OF THE SPECTRAL INVERSE FILTER


The Spectral Inverse Filter (SIF) has two modes: an adaptive mode and a manual mode. Thus it is actually a combination of the Adaptive Spectral Inverse Filter (ASIF) and the traditional Spectral Inverse Filter (SIF). Both are equalization filters that readjust the spectrum to match an expected spectral shape. It is especially useful when the target voice has been exposed to spectral coloration (i.e. muffling, hollowness, or tinniness), but it can also be used to remove bandlimited noises. The ASIF is much like the SIF, except it continually updates the spectral solution, whereas the SIF only updates the solution when it is “built”.

The filter maintains an average of the signal’s spectrum and uses this information to implement a high-resolution digital filter for correcting long-term spectral irregularities. The goal of the filter is to reshape the overall spectral envelope of the audio, not to respond to transient noises and characteristics.


Several user controls are available for refinement of SIF operation. The user can specify the expected spectrum so that the output audio is reshaped to a flat, pink or voice-like curve. An adapt rate setting controls the update rate for the spectral average, which in turn determines how quickly the filter responds to changes in the input audio. Upper and lower limit controls allow the user to specify the range over which equalization is applied, and a mode setting controls whether frequencies outside the equalization range are attenuated or left unaffected. The amount of spectral correction is adjustable using the Filter Amount control. The user can enable the auto-gain functionality to ensure that the output audio level is maintained at approximately the same as the input audio level. If the user

disables the auto- gain, an output gain slider is available to manually boost the level of the output signal.

As an aid to visualizing the filter operation, the user can view the input and output audio traces as well as the filter coefficient trace.




Spectral Inverse Filter Main Window (in Adaptive Mode)



Spectral Inverse Filter Main Window (in Manual Mode)

FILTER CONTROLS

FILTER OPERATION



Filter Operation


In this block, the user can select the operational mode of the filter. If the filter is being used to correct spectral coloration, the **Equalize Voice** mode should be selected. If the filter is being used to remove bandlimited noise, the **Attack Noise** mode should be selected.

Note: The **Filter Operation** mode selection only affects the behavior of the filter

outside the range selected by the upper and lower limits. In **Equalize Voice** mode, the frequency ranges outside the limits are attenuated. In **Attack Noise** mode, the frequency ranges outside the limits are left unaffected (subject to a transition region near the limits). If the auto gain is disabled and the manual gain is set to 0 dB, frequencies outside the limits and transition regions will be unaffected. However, if gain is applied, the gain will be reflected over the entire frequency range. See the section on **Upper and Lower Voice Limits** for more information on selecting the range.

Note: Changing the filter operation mode does not require an adaptation period to arrive at a “good” solution. Because a full average spectrum is maintained regardless of the mode setting, the new mode takes effect instantaneously in both the output audio and the display traces. However, since the auto gain adapts based on the actual applied filter with operational mode taken into account, there may be some adaptation time required to reach a stable auto gain value after the mode is changed.

OUTPUT SHAPE



Output Shape

In this block, the user can select the target spectral shape that the filter attempts to achieve. The SIF has an inherent spectral flattening effect on the audio. The selected spectral shape is applied to further reshape the audio spectrum. The following output shapes are available:

Flat - no additional shaping after ASIF flattening

Pink -- 3 dB per octave rolloff above 100 Hz is applied in addition to ASIF flattening

Voice - 6 dB/octave rolloff above and below 500 Hz in addition to ASIF flattening

Note: Changing the output shape does not require an adaptation period to arrive at a “good” solution. Because a full average spectrum is maintained regardless of the output shape setting, the new output shape takes effect instantaneously in both the output audio and the display traces. However, since the auto gain adapts based on the actual applied filter with the shaping curve taken into account, there may be some adaptation time required to reach a stable auto gain value after the shaping curve is changed.

FILTER AMOUNT

This setting controls the degree to which the SIF can affect the signal, with 0% corresponding to no filtering and 100% corresponding to full filtering. In general, it is best to use the minimum Filter Amount setting that produces the desired result. When **Equalize Voice** mode is used, a lower **Filter Amount** can reduce artifacts that result from a fast adapt rate, so the **Filter Amount** can be used to help strike a balance between responsiveness and stability. When **Attack Noise** mode is used to reduce bandlimited noise, a lower **Filter Amount** setting will often be a better choice to prevent the elevation of background noises.

Note: Changing the **Filter Amount** setting does not require an adaptation period to arrive at a “good” solution. Because a full average spectrum is maintained regardless of the setting, the new filter amount setting takes effect instantaneously in both the output audio and the display traces. However, since the auto gain adapts based on the actual applied filter with filter amount taken into account, there may be some adaptation time required to reach a stable auto gain value after the filter amount is adjusted.

OUTPUT GAIN AND AUTO GAIN




Output Gain and Auto Gain

These controls provide two options for adjusting the level of the SIF output. When **Auto Gain** is enabled, the SIF automatically monitors the input and output levels and applies a gain value that matches the output level to the input level. When **Auto Gain** is disabled, the user can use the **Output Gain** setting to specify the amount of boost applied to the SIF output. If **Auto Gain** is enabled, then the **Output Gain** slider is ignored.

The **Auto Gain** is an adaptive value whose rate of change depends on the same **Adapt Rate** slider setting that controls filter coefficient averaging. This means that when the filter response changes rapidly and dramatically, the auto gain will take some time to “catch up” to these changes. In particular, the output audio may clip when user settings are changed in a ways that have a boosting effect, such as switching from a pink to a flat shaping curve, adjusting the filter amount, or increasing the size of the ASIF region in Equalize Voice mode so that some frequencies that had been heavily attenuated are now present. While these settings changes will take effect immediately, the Auto Gain may take some time to adapt to the change. For this reason, when the user expects to be making many changes in the settings, it is often better to disable Auto Gain and instead choose a manual gain setting that avoids clipping.

LOWER AND UPPER VOICE LIMITS



Lower and Upper Voice Limit

The **Lower** and **Upper Voice Limits** allow the user to specify the frequency range, or “SIF region,” over which the SIF is applied. Two orange markers on the **Spectrum Analyzer** indicate where the lower and upper voice limits are located. Viewing audio on the display trace while manipulating the markers is an easy way to identify where your SIF region limits should fall.

In **Equalize Voice** mode, the SIF region is typically chosen to be the range over which speech frequencies are found. Setting a **Lower Limit** above 300 Hz or an **Upper Limit** below 3000 Hz is not recommended in equalize voice mode, as intelligibility may suffer. When in **Equalize Voice** mode, all frequencies outside the SIF region are assumed to be non-speech and are therefore attenuated.

In **Attack Noise** mode, the SIF region is typically chosen to “bracket” the bandlimited noise as closely as possible. Frequencies outside the ASIF region will be “passed through,” i.e. there will be little or no effect outside the SIF region except for a narrow transition band between the SIF region and the passbands.

Note: Changing the Voice Limits does not require an adaptation period to arrive at a “good” solution. Because a full average spectrum is maintained regardless of the Voice Limit settings, the new Voice Limits will take effect instantaneously in both the output audio and the display traces. However, since the auto gain adapts based on the actual applied filter with voice limits taken into account, there may be some adaptation time required to reach a stable auto gain value after the limits are changed.

ADAPTIVE MODE



Adaptive Mode

In **Adaptive** mode, the above figure is displayed and gives the user access to the **Adapt**, **Clear** and **Adapt Rate** controls. This puts the plug-in in the ASIF mode.

When the **Adapt** button is lit green, the ASIF is adapting in response to incoming audio. When the button is grayed, the ASIF response is frozen.


The **Clear** button allows the user to re-initialize the ASIF response and restart adaptation.

Note: After a **Clear** operation or after re-enabling adaptation, there will be an adaptation period while the filter adapts to the current input signal. The length of this adaptation period depends on the **Adapt Rate** control setting.

The **Adapt Rate** control allows the user to select the rate of adaptation for the spectral average on which the ASIF response is based. The spectral averager uses an exponential average of the form $H_{i+1} = (\alpha)(X_{i+1}) + (1 - \alpha)(H_i)$. The value shown in the display box corresponds to the averaging constant α in the exponential average. The lower the adapt rate value, the slower the filter will respond to changes in the input audio.

Note: “Fast response” sounds like a good thing, so it can be tempting to set the adapt rate to a high value. However, the goal of the ASIF is not to remove transient noises, but rather to reshape the long-term spectral envelope of the signal. If the adapt rate is too fast, the filter will respond too quickly to transient audio characteristics, which will produce artifacts in the output audio and will prevent the filter from settling on a good average solution. *For this reason, most applications will work best with adapt rates at the low end of the available range.* If you hear tonal artifacts that come and go in the output audio, or if the filter trace display coefficients seem to be changing rapidly, you probably need to reduce the adapt rate.

MANUAL MODE



Manual Mode

In **Manual** mode, the above figure is displayed and gives the user access to the **Run Analyzer**, **Clear** and **Build** controls. This puts the plug-in in the SIF mode.


The **Clear** button is used to zero the averager memory and causes the averaged spectrum to be recalculated anew.

The **Run Analyzer** button allows the user to start or stop the update of the averaged spectrum.

The **Build** button builds the spectral inverse filter based on the original input audio spectrum and the SIF control settings. Once the filter build is complete the calculated spectral inverse filter curve will be displayed as a green trace in the Filter Display area.

Hint: Before clicking the **Build** button, it is recommended that the spectrum analyzer be stopped (using the **Run Analyzer** button) to allow experimentation with the control settings for the same input spectrum.

DISPLAY



SIF Display

The display trace is used to view the filter input and output audio and the SIF filter response. The input audio is always shown in yellow, the output trace in

blue and the filter trace in green. The **Lower** and **Upper Voice Limits** are shown in red.

The controls for this display are identical to the controls for the Spectrum Analyzer, except that the analyzer is always in the Exponential averaging mode (hence, no choice is given to change the averaging mode) and the zooming function is disabled.

COMB FILTER


OVERVIEW OF THE COMB FILTER

The Comb filter is used to remove, or "notch out", harmonically related noises (noises which have exactly equally-spaced frequency components), such as power-line hum, constant-speed motor/generator noises, etc., from the input audio. The filter response consists of a series of equally-spaced notches which resemble a hair comb, hence the name "Comb filter".


Adjust the **Comb Frequency** to the desired spacing between notches (also known as "fundamental frequency"). Set the **Notch Limit** to the frequency beyond which you do not want any more notches. Set the **Notch Depth** to the amount in dB by which noise frequency components are to be reduced.

Normally, the Notch **Harmonics** option will be set to **All**, causing frequencies at all multiples of the Comb Frequency (within the Notch Limit) to be reduced. However, certain types of noises have only the odd or even harmonic components present. In these situations, set the Notch Harmonics option to either **Odd** or **Even**.

The **Auto Tracking** feature allows the filter to automatically hone in on the exact fundamental frequency and continually track it even if it shifts slightly.




Comb Filter Main Window - Manual Mode



Comb Filter Main Window - Auto Tracking

FILTER CONTROLS

COMB FREQUENCY



Comb Frequency

Specifies fundamental frequency in Hertz of comb filter. Notches are generated at multiples, or harmonics, of this frequency. When **Auto Tracking** is turned on, this control does not appear.

NOTCH LIMIT



Notch Limit

Specifies frequency in Hertz above which no notches are generated. Minimum Notch Limit is half the **Comb Frequency**, while maximum Notch Limit depends upon the sample rate.

NOTCH DEPTH



Notch Depth

Specifies the depth of notches that are generated. Notch Depth is adjustable from 0 dB to 120 dB in 1 dB steps.

NOTCH HARMONICS



Notch Harmonics

Specifies whether notches will be generated at All, Odd, or Even multiples, or harmonics, of the Comb Frequency. If, for example, the Comb Frequency is set to 60.0 Hz, then selecting **All** will generate notches at 60 Hz, 120 Hz, 180 Hz, 240 Hz, 300 Hz, etc. Selecting **Odd** will generate notches at 60 Hz, 180 Hz, 300 Hz, etc.

Selecting **Even** will generate notches at 120 Hz, 240 Hz, 360 Hz etc.

AUTO TRACKING



Auto Tracking

Enables **Auto Tracking** for the Comb filter. When enabled, the **Comb Frequency** control is no longer available and the filter will determine the fundamental frequency to use.

When disabled, the last calculated frequency will then be loaded into the **Comb Frequency** control.

NOTE: It is advised that you adjust the **Comb Frequency** manually to get it as close as possible to the desired fundamental frequency before enabling **Auto Tracking**.


X-PASS FILTER

OVERVIEW OF THE X-PASS FILTER

The X-Pass Filter provides six different bandlimiting filters in a single plug-in. The following filters are included:

Filter	Application
Lowpass	Decrease the energy level of high frequency noises
Highpass	Decrease the energy level of low frequency noises
Bandpass	Decrease the energy level of noises above and below a passband region
Bandstop	Decrease the energy level of noises within a stopband region
Notch	Remove a narrow-band noise
Slot	Isolate a narrow-band signal

Only one filter can be operational at a time. While they all share similar controls, they perform different functions (as detailed below). Whichever filter you select, its transfer function (or envelope) is displayed in the Spectrum Analyzer as a light gray trace.



X-Pass Main Window

MAIN PLUG-IN WINDOW

The **Filter Type** control selects which of the six filters is currently engaged. Only one filter is operational at a time. If you desire to apply more than one of the filters on a given piece of audio, you must either use multiple instances of the plug-in in the Filter Chain or render the audio multiple times (using a different filter each time).

When a filter is selected, that filter's controls are displayed in the main window. This window provides access to all the functionality of the plug-in. A spectrum analyzer is provided as an aid in determining the characteristics of the target audio. Details of the plug-in are described in the following sections.


LOWPASS FILTER

The Lowpass filter is used to decrease the energy level (lower the volume) of all signal frequencies above a specified Cutoff Frequency, thus reducing high-frequency noises, such as tape hiss, from the input audio. The Lowpass filter is sometimes called a "hiss filter."

The Cutoff Frequency is usually set above the voice frequency range so that the voice signal will not be disturbed. While listening to the filter output audio, the Cutoff Frequency can be incrementally lowered from its maximum frequency until

the quality of the voice just begins to be affected, achieving maximum elimination of high-frequency noise.

The amount of volume reduction above the Cutoff Frequency can further be controlled by adjusting the Stopband Attenuation setting (maximum volume reduction is 120dB). The slope at which the volume is reduced from normal (at the Cutoff Frequency) to the minimum volume (specified by Stopband Attenuation) can also be controlled by adjusting the Transition Slope setting.



Lowpass Filter Controls

FILTER CONTROLS

CUTOFF FREQUENCY



Cutoff Frequency

Specifies frequency in Hertz above which all signals are attenuated. Frequencies below this cutoff are unaffected. Maximum Cutoff Frequency depends upon the Sample Rate. Cutoff Frequency can be adjusted in 1 Hz steps.


TRANSITION SLOPE



Transition Slope

Specifies slope at which frequencies above the Cutoff Frequency are rolled off in dB per octave. Sharpest roll off occurs when Transition Slope is set to maximum, while gentlest roll off occurs when Transition Slope is set to minimum. Sharp rolloffs may cause the voice to sound hollow but will allow more precise removal of high frequency noises. Note that the indicated value changes depending upon the Cutoff Frequency and Sample Rate.

STOPBAND ATTENUATION



Stopband Attenuation


Specifies amount in dB by which frequencies above the Cutoff Frequency are ultimately attenuated. Stopband attenuation is adjustable from 0dB to 120dB in 1 dB steps.

HIGHPASS FILTER

The Highpass filter is used to decrease the energy level (lower the volume) of all signal frequencies below a specified Cutoff Frequency, thus reducing low-frequency noises, such as tape or acoustic room rumble, from the input audio (The Highpass filter is sometimes called a "rumble filter").

The Cutoff Frequency is usually set below the voice frequency range (somewhere below 300 Hz) so that the voice signal will not be disturbed. While listening to the filter output audio, the Cutoff Frequency, initially set to 0 Hz, can be incrementally increased until the quality of the voice just begins to be affected, achieving maximum elimination of low-frequency noise.


The amount of volume reduction below the Cutoff Frequency can further be controlled by adjusting the Stopband Attenuation setting (maximum volume reduction is 120dB). The slope at which the volume is reduced from normal (at the Cutoff Frequency) to the minimum volume (specified by Stopband Attenuation) can also be controlled by adjusting the Transition Slope setting.



Highpass Filter Controls

FILTER CONTROLS

CUTOFF FREQUENCY



Cutoff Frequency

Specifies frequency in Hertz below which all signals are attenuated. Frequencies above this cutoff are unaffected. Minimum Cutoff Frequency is 0 Hz (no frequencies attenuated), while the maximum Cutoff Frequency depends upon the Sample Rate. Cutoff Frequency can be adjusted in 1 Hz steps.


TRANSITION SLOPE



Transition Slope

Specifies slope at which frequencies below the Cutoff Frequency are attenuated in dB per octave. Sharpest attenuation occurs when Transition Slope is set to maximum, while gentlest attenuation occurs when Transition Slope is set to minimum. Note that the indicated value changes depending upon the Cutoff Frequency and Sample Rates.

STOPBAND ATTENUATION



Stopband Attenuation

Specifies amount in dB by which frequencies below the Cutoff Frequency are ultimately attenuated.


BANDPASS FILTER

The Bandpass filter is used to decrease the energy level (lower the volume) of all signal frequencies below a specified Lower Cutoff Frequency and above a specified Upper Cutoff Frequency, thus combining the functions of a Lowpass and Highpass filter connected in series into a single filter. The signal region between the Lower Cutoff Frequency and the Upper Cutoff Frequency is called the passband region. The Bandpass filter is useful for simultaneously reducing both low-frequency rumble and high-frequency hiss.

The Lower Cutoff Frequency is usually set below the voice frequency range (somewhere below 300 Hz) so that the voice signal will not be disturbed. While listening to the filter output audio, the Lower Cutoff Frequency, initially set to 0 Hz, can be incrementally increased until the quality of the voice just begins to be affected, achieving maximum elimination of low-frequency noise.

The Upper Cutoff Frequency is usually set above the voice frequency range (somewhere above 3000 Hz) so that the voice signal will not be disturbed. While listening to the filter output audio, the Upper Cutoff Frequency, initially set to its maximum frequency, can be incrementally lowered until the quality of the voice just begins to be affected, achieving maximum elimination of high-frequency noise.


The amount of volume reduction outside the passband region can further be controlled by adjusting the Stopband Attenuation setting (maximum volume reduction is 120dB). The slope at which the volume is reduced from normal (at each Cutoff Frequency) to the minimum volume (specified by Stopband Attenuation) can also be controlled by adjusting the Transition Slope setting.



Bandpass Filter Controls

FILTER CONTROLS

LOWER CUTOFF FREQUENCY




Lower Cutoff Frequency

Specifies frequency in Hertz below which all signals are attenuated. Frequencies between this cutoff and the Upper Cutoff Frequency are unaffected. Minimum Lower Cutoff Frequency is 0 Hz, while the maximum Lower Cutoff Frequency is 10 Hz below the Upper Cutoff Frequency. Lower Cutoff Frequency can be adjusted in 1 Hz steps.

NOTE: The Lower Cutoff Frequency can never be set higher than 10 Hz below the Upper Cutoff Frequency.

UPPER CUTOFF FREQUENCY



Upper Cutoff Frequency

Specifies frequency in Hertz above which all signals are attenuated. Frequencies between this cutoff and the Lower Cutoff Frequency are unaffected. Minimum Upper Cutoff Frequency is 10 Hz above the Lower Cutoff Frequency, while the maximum Upper Cutoff Frequency depends upon the Sample Rate. Upper Cutoff Frequency can be adjusted in 1 Hz steps.

NOTE: The Upper Cutoff Frequency can never be set lower than 10 Hz above the Lower Cutoff Frequency.


TRANSITION SLOPE



Transition Slope

Specifies slope at which frequencies below the Lower Cutoff Frequency and above the Upper Cutoff Frequency are attenuated in dB per octave. Sharpest attenuation occurs when Transition Slope is set to maximum, while gentlest attenuation occurs when Transition Slope is set to minimum. Note that the indicated value changes depending upon the Cutoff Frequency and Sample Rates. Also, note that the Lower and Upper Transition Slopes always have different values; this is because the frequency width of an octave is proportional to Cutoff Frequency.

STOPBAND ATTENUATION



Stopband Attenuation


Specifies amount in dB by which frequencies below the Lower Cutoff Frequency and above the Upper Cutoff Frequency are ultimately attenuated.

BANDSTOP FILTER

The Bandstop filter is used to decrease the energy level (lower the volume) of all signal frequencies above a specified Lower Cutoff Frequency and below a specified Upper Cutoff Frequency. The signal region between the Lower Cutoff Frequency and the Upper Cutoff Frequency is called the stopband region. The Bandstop filter is useful for removing in-band noise from the input signal.

The Lower Cutoff Frequency is usually set below the frequency range of the noise, while the Upper Cutoff Frequency is set above the frequency range of the noise. While listening to the filter output audio, the Lower and Upper Cutoff Frequencies can be incrementally adjusted to achieve maximum elimination of noise while minimizing loss of voice.


The amount of volume reduction in the stopband region can further be controlled by adjusting the Stopband Attenuation setting (maximum volume reduction is 120dB). The slope at which the volume is reduced from normal (at each Cutoff Frequency) to the minimum volume (specified by Stopband Attenuation) can also be controlled by adjusting the Transition Slope setting.



Bandstop Filter Controls

FILTER CONTROLS

LOWER CUTOFF FREQUENCY




Lower Cutoff Frequency

Specifies frequency in Hertz below which no signals are attenuated. Frequencies between this cutoff and the Upper Cutoff Frequency are attenuated. Minimum Lower Cutoff Frequency is 0 Hz, while the maximum Lower Cutoff Frequency is 10 Hz below the Upper Cutoff Frequency. Lower Cutoff Frequency can be adjusted in 1 Hz steps.

NOTE: The Lower Cutoff Frequency can never be set higher than 10 Hz below the Upper Cutoff Frequency.

UPPER CUTOFF FREQUENCY



Upper Cutoff Frequency

Specifies frequency in Hertz above which no signals are attenuated. Frequencies between this cutoff and the Lower Cutoff Frequency are attenuated. Minimum Upper Cutoff Frequency is 10 Hz above the Lower Cutoff Frequency, while the maximum Upper Cutoff Frequency depends upon the Sample Rate. Upper Cutoff Frequency can be adjusted in 1 Hz steps.

NOTE: The Upper Cutoff Frequency can never be set lower than 10 Hz above the Lower Cutoff Frequency.

TRANSITION SLOPE




Transition Slope

Specifies slope at which frequencies above the Lower Cutoff Frequency and below the Upper Cutoff Frequency are attenuated in dB per octave. Sharpest attenuation occurs when Transition Slope is set to maximum, while gentlest attenuation occurs when Transition Slope is set to minimum. Note that the indicated value changes depending upon the Cutoff Frequency and Sample Rates. Also, note that the Lower

and Upper Transition Slopes always have different values; this is because the frequency width of an octave is proportional to Cutoff Frequency.

STOPBAND ATTENUATION



Stopband Attenuation

Specifies amount in dB by which frequencies above the Lower Cutoff Frequency and below the Upper Cutoff Frequency are attenuated.

NOTCH FILTER


The Notch filter is used to remove, or "notch out", a narrow-band noise, such as a tone or a whistle, from the input audio with minimal effect to the remaining audio. The Notch filter works best with stable noise sources which have constant frequency; if the frequency of the noise source varies, the 1-Channel Adaptive filter is recommended.

To properly utilize the Notch filter, you will first need to identify the frequency of the noise; this is best done using the Spectrum Analyzer window.

Initially set the Notch Depth to 120 dB and the Notch Width to the narrowest possible value. Next, set the Notch Frequency to the noise frequency. Fine adjustment of the Notch Frequency may be necessary to place the notch precisely on top of the noise signal and achieve maximum reduction of the noise. This is best done by adjusting the Notch Frequency up or down 1 Hz at a time while listening to the Notch filter output on the headphones.

Often, the noise frequency will not remain absolutely constant but will vary slightly due to modulation, recorder wow and flutter, and acoustic "beating." Therefore, you may need to increase the Notch Width from its minimum setting to keep the noise within the notch.


For maximum noise reduction, set the Notch Depth to 120dB. It is best to adjust the Notch Depth up from 120 dB until the tone is observed, then increase the depth 5 dB.



Notch Filter Controls

FILTER CONTROLS


NOTCH FREQUENCY



Notch Frequency

Specifies frequency in Hertz which is to be removed from the input audio. Minimum Notch Frequency is 10 Hz, while maximum Notch Frequency depends upon the Sample Rate. Notch Frequency is adjustable in 1 Hz steps.

NOTCH WIDTH




Notch Width

Width of the generated notch in Hertz.

NOTE: Notch Width varies with the Sample Rate.

NOTCH DEPTH



Notch Depth

Depth of the notch that is generated.

SLOT FILTER

The Slot filter is used to isolate, or "slot", a single-frequency signal, such as a tone or a whistle, in the input audio, attenuating all other audio. This is the exact opposite of the Notch filter function.


NOTE: The Slot filter has very little use in speech enhancement applications; the main value is in isolating other types of signals that are non-speech in nature.

To properly utilize the Slot filter, you will first need to identify the frequency of the signal to be isolated; this is best done using the Spectrum Analyzer window.

Once the frequency of the signal has been identified, initially set Stopband Attenuation to 120 dB and the Slot Width to the narrowest possible value. Next, set the Slot Frequency to the signal frequency. Fine adjustment of the Slot Frequency may be necessary to place the slot right on top of the signal. This is best done by adjusting the Slot Frequency up or down 1 Hz at a time while listening to the Slot filter output on the headphones.

Usually, the signal frequency will not remain constant but will vary slightly due to modulation, recorder wow and flutter, and acoustic "beating". Therefore, you may need to increase the Slot Width from its minimum setting to avoid having the signal move in and out of the slot.

To optimize background noise reduction for your application, set the Stopband Attenuation to 120dB. If, however, you wish to leave a small amount of the background noise mixed in with the isolated signal, adjust the Stopband Attenuation to the desired value.



Slot Filter Controls

FILTER CONTROLS SLOT FREQUENCY



Slot Frequency

Specifies frequency in Hertz which is to be enhanced in the input audio. Minimum Slot Frequency is 10 Hz, while maximum Slot Frequency depends upon the Sample Rate. Slot Frequency is adjustable in 1 Hz steps.

SLOT WIDTH




Slot Width

Width of the generated slot in Hertz.

NOTE: Slot Width varies with the Sample Rate.

STOPBAND ATTENUATION



Stopband Attenuation

Specifies amount in dB by which frequencies other than the Slot Frequency are attenuated.

SPECTRAL SUBTRACTION FILTER

OVERVIEW OF THE SPECTRAL SUBTRACTION FILTER

The Spectral Subtraction filter is a frequency-domain filter that implements automatic noise reduction over the entire frequency spectrum. It operates by continually measuring the spectrum of the input signal and attempting to identify which portions of the signal are voice and which portions are non-voice (or noise). All portions determined to be noise are used to continually update a **noise estimate** calculation; this is used to calculate the equalization curve that needs to be applied to the input signal to reduce each band's energy by the amount of noise energy calculated to be in that band.

The net result is an output signal that has all non-voice signals reduced in level as much as possible, thereby “polishing” the enhanced voice signal as much as possible prior to final equalization and AGC.

The Spectral Subtraction filter has two modes: the **Noise Reducer** mode and the **Noise EQ** mode.


Operation in the Noise EQ mode is governed by 20 control sliders, each representing a frequency band. Adjusting the control sliders allows the user to precisely control the amount of noise reduction being applied within each of 20 distinct groups of frequency bands, offering much more precise control of the spectral subtraction than is available in the Noise Reducer mode, though it does take more time to setup.

In the Noise Reducer mode all 20 control sliders are linked together and are adjusted using the first slider. Adjusting the master slider allows the user to precisely control the amount of noise reduction being applied to all frequency bands; the greater the value, the more aggressive the operation of the Noise Reducer.


Because large amounts of noise reduction invariably create audible “birdy noise” artifacts in the output audio due to the nature of adaptive frequency-domain processing, the user should always try to minimize the amount of noise reduction being applied to achieve the best balance between maximal noise reduction and minimal audible artifacts.

Finally, for convenience an Output Gain control and Output level bargraph are provided to enable the user to adjust the processed output signal to maximum level

for better listening and recording.




Spectral Subtraction Main Window - Noise Reducer Mode



Spectral Subtraction Main Window - Noise EQ Mode

FILTER CONTROLS

FILTER CONTROLS



Filter Controls

The **Filter Controls** combo box allows the user to select either the **Noise EQ** mode or the **Noise Reducer** mode (displayed as N. Reducer). When switching from one mode to the other, the previous mode's controls are remembered and preserved.


CLEAR



Clear Button

Used to clear the spectral subtraction solution currently in memory and restart the algorithm from scratch.

FILTER OUTPUT GAIN



Filter Output Controls


The **Output Gain** control allows user to apply between 0 and 30dB of makeup gain

to the processed output signal to maximize the signal level prior to final equalization, AGC, and listening/recording. The associated bargraph shows the actual output signal level after the gain has been applied.

NOISE EQ CONTROLS

The following sections apply only to the Noise EQ mode.

Slider Controls



Noise EQ Slider Controls

The slider controls are used to specify the amount of noise reduction that the spectral subtraction attempts to apply to the input signal within each of 20 separate groups of frequency bands. Within each band, adjustment range is 0 (no attenuation) to 100% (maximal attenuation) in 1% increments.

Zero All



Zero All Button

This button instantly moves the slider controls for all bands to 0%, defeating the entire equalizer. This is a useful feature when it is desired to reset all sliders from scratch.

Normalize



Normalize Button

This button instantly shifts all slider controls up together until the top slider is at 0%. After normalization, the relative positioning of the sliders remains the same. This allows the filter to implement the desired equalization curve with minimum signal loss.

Maximize All



Maximize All Button

This button instantly moves the slider controls for all bands to 100%, maximizing the attenuation for all bands. This is a useful feature when it is desired to quickly adjust the sliders such that only a few bands are passed with all others rejected.

All Up 1



All Up 1 Button

This button shifts all sliders up 1%, while maintaining the desired curve. This is useful when the desired shape is found, but you want the filter to be less aggressive as a whole. No slider will be allowed to go less than 0%.

All Down 1



All Down 1 Button

This button shifts all sliders down 1%, while maintaining the desired curve. This is useful when the desired shape is found, but you want the filter to be more aggressive as a whole. No slider will be allowed to go more than 100%.

Store / Recall




Store and Recall Buttons

These buttons allow the user to store and recall a slider configuration to a user-specified disk file that will not be lost when the computer is turned off.

NOISE REDUCER CONTROLS

The following sections apply only to the **Noise Reducer** mode.

Slider Control



Noise Reducer Slider Controls

When in **Noise Reducer** mode, all the slider controls are linked to the first slider control. This slider is used to specify the amount of noise reduction that the spectral subtraction attempts to apply to the input signal across all frequency bands. Adjustment range is 0 (no attenuation) to 100% (maximal attenuation) in 1% increments.

MULTI-BAND FILTER

OVERVIEW OF THE MULTI-BAND FILTER

The Multiple Notch filter is used to remove, or “notch out,” single-frequency noises such as tones or whistles with minimal effect on signal frequencies other than the notch frequency. Single notches can be added one at a time and configured individually. Also, notch “groups” can be added to cancel many harmonically related frequencies at once.

The Multiple Notch filter is synthesized from a frequency-domain representation of the desired notch profile. An inverse FFT builds FIR coefficients based on the frequency-domain representation. For this reason, the notches in this filter are “square” notches rather than “V” notches. Square notches mean that frequencies very close to the specified center frequency will be cancelled along with the center frequency. However, the square notches also mean that the Multiple Notch filter is able to tolerate moderate variances in the specified frequency such as those caused by “wow and flutter” effects. (Filters that use “V” notches include the Notch filter, the Comb filter, and the Parametric Equalizer.)

To properly utilize the Multiple Notch filter, you will first need to identify the noise frequencies. The easiest way to do this is to use the included spectrum analyzer. Once the noise frequencies have been identified, add a notch for each frequency.

Notches are defined by three values: the notch frequency, the notch width, and

the notch depth. The notch frequency is simply the frequency at which the notch should be centered. The notch width defines the desired width of the square notch in Hz, and the notch depth defines the desired depth in dB.

Often, tonal noises include not only the fundamental frequency, but also harmonic multiples of that frequency. Instead of requiring the addition of an individual notch for each harmonic, the CARDINAL Multiple Notch filter allows the addition of Notch Groups to cancel harmonically related tones in a single action. A Notch Group is defined in relation to its Base Notch. The Base Notch is defined with a frequency, width, and depth just like a single notch.

Frequency, width, and depth of all other notches in the group will be calculated based on these parameters.

OVERVIEW OF MULTIPLE-SLOT


The Multiple Slot filter is used to isolate, or "slot" single-frequency noises such as tones or whistles in the input audio, attenuating all other audio. This is the exact opposite of the Multiple Notch filter function. Single slots can be added one at a time and configured individually. Also, slot "groups" can be added to isolate many harmonically related frequencies at once.

The Multiple Slot filter is synthesized from a frequency-domain representation of the desired slot profile. An inverse FFT builds FIR coefficients based on the frequency-domain representation. For this reason, the slots in this filter are "square" slots rather than "V" slots. Square slots mean that frequencies very close to the specified center frequency will be cancelled along with the center frequency. However, the square slots also mean that the Multiple Slot filter is able to tolerate moderate variances in the specified frequency such as those caused by "wow and flutter" effects.


To properly utilize the Multiple Slot filter, you will first need to identify the noise frequencies. The easiest way to do this is to use the included spectrum analyzer. Once the noise frequencies have been identified, add a slot for each frequency.

Slots are defined by three values: the slot frequency, the slot width, and the slot gain. The slot frequency is simply the frequency at which the slot should be centered. The slot width defines the desired width of the square slot in Hz, and the slot gain defines the desired amplitude in dB.

Often, tonal noises include not only the fundamental frequency, but also harmonic multiples of that frequency. Instead of requiring the addition of an individual slot for each harmonic, the CARDINAL Multiple Slot filter allows the addition of Slot Groups to cancel harmonically related tones in a single action. A Slot Group is defined in relation to its Base Slot. The Base Slot is defined with a frequency, width, and gain just like a single slot. Frequency, width, and gain of all other slots in the group will be calculated based on these parameters.




Multi-Band Main Window - Multiple Notch Mode



Multi-Band Main Window - Multiple Slot Mode

FILTER CONTROLS

FILTER MODE




Filter Mode

The **Filter Mode** combo box allows the user to select either the **Multiple Notch** mode or the **Multiple Slot** mode. When switching from one mode to the other, the previous mode's controls are remembered and preserved.

MULTIPLE NOTCH FILTER CONTROLS

Notch Frequency



Notch Frequency

Specifies frequency in Hertz which is to be removed from the input audio. Minimum Notch Frequency is 10 Hz, while maximum Notch Frequency depends upon the Sample Rate. Notch Frequency is adjustable in 1 Hz steps. Also sets the Base Frequency for a Notch Group.


Notch Width



Notch Width

Width of the generated notch in Hertz.


Notch Depth



Notch Depth

Depth of the generated notch in dB.


Notch Spacing



Notch Spacing

Notch Spacing defines where the other notches in the group are to be placed. If the Base Notch frequency is F , and the spacing is set to S , then notches will be placed at frequencies F , $F+S$, $F+2S$, $F+3S$, etc.

Width Factor



Width Factor

Width Factor defines how wide the group notches should be. Frequency variations often occur as a percentage of the frequency, so the variation width in Hz is much larger at high frequencies. The Width Factor defines a percentage width up to a maximum of 1.9%, and each notch will be at least the width defined by that percentage. For instance, if a notch group has width factor = 0.015, and one of the notches in that group is at 1000 Hz, then the width of the 1000 Hz notch will be at least $1000 \times 0.015 = 15$ Hz.

NOTE: The frequency-domain representation used to build the Multiple Notch filter has an inherent minimum notch width. Especially at the lower frequency notches, the width specified by the Width Factor will often fall below that minimum width, in which case the minimum width is used. For this reason, the effect of the Width Factor control may only be visible at the higher frequency notches.

Depth Factor



Depth Factor

Depth Factor defines how deep the group notches should be. Many harmonic tonal noises have a “ $1/f$ ” volume profile, where the lower harmonics are strong and higher harmonics are progressively weaker. The Depth Factor controls the depth taper of the notches so that the notch depth can parallel the harmonic strength profile. The base notch always has the specified Notch Depth, while subsequent notches taper to smaller depths as frequency increases. The higher the Depth Factor, the more gradual the taper. A Depth Factor of 0.0 produces the most severe taper and means effectively that there are no harmonics at all. A Depth Factor of 1.0 means that notches have uniform depth at the Base Notch depth setting.

Upper Limit



Upper Limit

Upper Limit defines how many notches there are in the group. If a harmonic tonal noise only extends up to a certain frequency, it may be undesirable to notch out all multiples of the base frequency when only a few are needed. In this case, set the Upper Limit just above the highest frequency where a notch is desired; notches will be added up to that limit, and no notches will be added above the limit.

Add Notch



Add Notch Button

Adds a new single notch at the frequency indicated in the Notch Frequency box and by a marker on the visualization axis. The notch is added with default settings, and the user is presented with controls to adjust the frequency, width, and depth of the notch.

Add Group



Add Group Button

Adds a new notch group with its base notch at the frequency indicated in the Notch Frequency box and by a marker on the visualization axis. The notch group is added with default settings, and the user is presented with controls to adjust the frequency, width, depth, notch spacing, depth factor, width factor, and upper limit of the notch group.

DTMF Button



DTMF Button

Inserts 8 pre-defined frequencies that make up the **Dual-tone multi-frequency (DTMF)**. The version of DTMF used for telephone **tone dialing** is known by the trademarked term **Touch-Tone**, and is standardized by ITU-T Recommendation.

Remove Button



Remove Button

Removes the currently selected notch or notch group from the filter.

Remove All Button



Remove All Button

Removes all notches and notch groups from the filter.

Store Button



Store Button

Saves the filter's current configuration to a disk file.

Recall Button



Recall Button

Loads a previously saved filter configuration from a disk file.

MULTIPLE SLOT CONTROLS


Slot Frequency



Slot Frequency

Specifies frequency in Hertz which is to be enhanced in the input audio. Minimum Slot Frequency is 10 Hz, while maximum Slot Frequency depends upon the Sample Rate. Slot Frequency is adjustable in 1 Hz steps.


Slot Width



Slot Width

Width of the generated slot in Hz.


Slot Gain



Slot Gain

The gain of the base slot, in dB.


Stopband Depth



Stopband Depth

Specifies amount in dB by which frequencies other than the Slot Frequency are attenuated.


Slot Spacing



Slot Spacing

Slot Spacing defines where the other slots in the group are to be placed. If the Base Slot frequency is F , and the spacing is set to S , then slots will be placed at frequencies F , $F+S$, $F+2S$, $F+3S$, etc.

Width Factor



Width Factor

Width Factor defines how wide the group slots should be. Frequency variations often occur as a percentage of the frequency, so the variation width in Hz is much larger at high frequencies. The Width Factor defines a percentage width up to a maximum of 1.9%, and each notch will be at least the width defined by that percentage. For instance, if a slot group has width factor = 0.015, and one of the slots in that group is at 1000 Hz, then the width of the 1000 Hz slot will be at least $1000 \times 0.015 = 15$ Hz.

NOTE: The frequency-domain representation used to build the Multiple Slot filter has an inherent minimum slot width. Especially at the lower frequency slots, the width specified by the Width Factor will often fall below that minimum width, in which case the minimum width is used. For this reason, the effect of the Width Factor control may only be visible at the higher frequency notches.

Gain Factor



Gain Factor

Gain Factor defines how much gain is applied to the group slots. Many harmonic tonal noises have a “ $1/f$ ” volume profile, where the lower harmonics are strong and higher harmonics are progressively weaker. The Gain Factor controls the gain taper of the slots so that the slot gain can parallel the harmonic strength profile.

The base slot always has the specified Slot Gain, while subsequent slots taper to smaller gains as frequency increases. The higher the Gain Factor, the more gradual the taper. A Gain Factor of 0.0 produces the most severe taper and means effectively that there are no harmonics at all. A Gain Factor of 1.0 means that slots have uniform gain at the Base Slot gain setting.

Upper Limit



Upper Limit

Upper Limit defines how many slots there are in the group. If a harmonic tonal noise only extends up to a certain frequency, it may be undesirable to slot out all multiples of the base frequency when only a few are needed. In this case, set the Upper Limit just above the highest frequency where a slot is desired; slots will be added up to that limit, and no slots will be added above the limit.

Add Slot



Add Slot Button

Adds a new single slot at the frequency indicated in the Slot Frequency box and by a marker on the visualization axis. The slot is added with default settings, and the user is presented with controls to adjust the frequency, width, and gain of the slot.

Add Group



Add Group Button

Adds a new slot group with its base slot at the frequency indicated in the Slot Frequency box and by a marker on the visualization axis. The slot group is added with default settings, and the user is presented with controls to adjust the frequency, width, gain, slot spacing, gain factor, width factor, and upper limit of the slot group.

DTMF Button



DTMF Button

Inserts 8 pre-defined frequencies that make up the **Dual-tone multi-frequency (DTMF)**. The version of DTMF used for telephone **tone dialing** is known by the trademarked term **Touch-Tone**, and is standardized by ITU-T Recommendation.

Remove Button



Remove Button

Removes the currently selected slot or slot group from the filter.

Remove All Button



Remove All Button

Removes all slots and slot groups from the filter.

Store Button



Store Button

Saves the filter's current configuration to a disk file

Recall Button



Recall Button

Loads a previously saved filter configuration from a disk file.

GRAPHIC EQ

OVERVIEW OF THE GRAPHIC EQUALIZER

The Graphic EQ provides two different equalizers in a single plug-in - the **Graphic EQ** and **Hi-Res Graphic**.


Only one equalizer can be operational at a time.

Overview of 20-Band Graphic EQ

The 20-band Graphic Equalizer is an easy-to-use linear-phase FIR digital filter that is used to reshape the spectrum of the final output signal. Reshaping is accomplished with twenty vertical scroll bars (also called "slider" controls) which adjust the attenuation of each frequency band. These controls are very similar to the slider controls found on analog graphic equalizers found on many consumer stereo systems, and thus should be very familiar to even the novice user.

However, unlike analog graphic equalizers, this digital equalizer has some very powerful additional capabilities. For example, the **Normalize** button allows the user to instantly move all slider controls up until the top slider is at 0dB. The **Zero All** button instantly sets all the sliders to 0dB, while the **Maximize** button instantly sets all the sliders to -100dB. The **All Down 1dB** button instantly moves all sliders down in 1dB increments, while the **All Up 1dB** button moves all sliders up in 1dB increments. None of these functions are available in an analog graphic

equalizer! Notice also that the 20 sliders are spread across the selected Bandwidth and that the frequency spacing is optimized for voice processing.



Graphic EQ Main Window - Graphic EQ Mode

Additionally, since a computer with a disk drive operates the equalizer, a **Store** and **Recall** capability is available. This allows the user to store commonly-used slider configurations in disk memories so that they can be instantly recalled later whenever they are needed, without having to manually adjust the slider controls.


Overview of Hi-Res Graphic EQ

In some applications, it may be necessary to precisely reshape the spectrum of input audio prior to passing it through successive filter stages. For example, if the audio is from a microphone which has an unusual frequency response curve (for example, a microphone acoustically modified as a result of concealment), a compensation filter that reshapes the audio to a normal spectral shape might be desirable.

The Hi-Res Graphic Filter is essentially a 460-band graphic equalizer; however, instead of having 460 separate slider controls, it allows the user to precisely draw the desired filter shape on the computer screen, using the mouse, with as much or as little detail as desired. Once the filter shape has been drawn, a linear-phase digital filter is constructed in the PC and transferred to the external processor.

The **Normalize** button allows the user to shift the entire filter curve up until the highest point is at 0dB.


A **Store** and **Recall** capability is also provided to allow the user to store commonly-used filter shapes to disk memories so that they can be recalled later.



Graphic EQ Main Window - Hi-Res Graphic Mode

EQUALIZER CONTROLS

Equalizer Mode




Equalizer Mode

The **Equalizer Mode** combo box allows the user to select either the **Graphic EQ** mode or the **Hi-Res Graphic** mode. When switching from one mode to the other, the previous mode's controls are remembered and preserved.

GRAPHIC EQ CONTROLS

Slider Controls



Slider Controls

The twenty vertical scroll bar "slider" controls are used to set the frequency response of the equalizer. Each slider can set the gain of its frequency band to any value between 0dB and -100 dB in 1dB steps.

Zero All



Zero All Button

This button instantly moves the slider controls for all bands to 0dB, defeating the entire equalizer. This is a useful feature when it is desired to reset all sliders from scratch.

Normalize Button



Normalize Button

This button instantly shifts all slider controls up together until the top slider is at 0dB. After normalization, the relative positioning of the sliders remains the same. This allows the digital equalizer to implement the desired equalization curve with minimum signal loss.

Maximize All



Maximize All Button

All Up 1



All Up 1 Button

This button shifts all sliders up by 1dB from their current position; no slider, however, will be allowed to go above 0dB. This button allows the user to shift the entire equalizer curve up so that there will be room to move one or more sliders down relative to the others.

All Down 1



All Down 1 Button

This button shifts all sliders down by 1dB from their current position; no slider, however, will be allowed to go below -100dB. This button allows the user to shift the entire equalizer curve down so that there will be room to move one or more sliders up relative to the others.

HI-RES GRAPHIC CONTROLS

Hi-Res Graphic Mini-Tutorial

The smoothing curve is graphed by the user using *control points*. These control points are seen in Figure 110: Graphic EQ Main Window - Hi-Res Graphic Mode as large circles on the graph. Control points represent a point on the curve where the slope of the line changes. Users can manipulate these control points in one of three ways:

- Add a control point
- Delete a control point
- Move a control point

To add a control point, simply click on the graph where you want it to be. The control point will immediately appear and you will hear the audio change immediately.

To delete a control point, click on an existing control point (it will turn red) and then click the **Delete** button (except the first and the last points, they cannot be deleted). This will remove the control point and the curve will snap back between the control points on either side.

To move a control point, click on an existing control point (it will turn red) and drag it with the mouse. Control points can only be moved vertically, which adjusts the gain at that point. Control points cannot be moved horizontally in an attempt to change the frequency at which the control point exists.

Store Button



Store Button

This button allows the user to store the current configuration to a user-specified disk file that will not be lost when the computer is turned off.

Recall Button



Recall Button

This button allows the user to recall a previously stored configuration from any of the saved disk files previously generated using the **Store** button.

Normalize Button



Normalize Button

The **Normalize** button allows the user to shift the entire filter curve up until the highest point is at 0dB.

Clear All



Clear All Button

The **Clear All** button will clear all graph control points and reset the filter to a pass-through.

Delete Button



Delete Button

The **Delete** button will delete the current control point (highlighted in red) in the graph.

GAIN STAGE

OVERVIEW OF THE GAIN STAGE

The Gain Stage provides two different gain stages in a single plug-in - the **Automatic Gain Control (AGC)** and **Limiter/Compressor/Expander (LCE)**.

Only one gain stage can be operational at a time.

Overview of AGC

The Automatic Gain Control automatically attempts to boost low-level output signals to a peak reference level (-18dB bargraph level) by gradually increasing output signal gain over a specified Release Time interval until either the proper level or Maximum Gain has been reached. This compensates for near party/far party conversations and for losses in signal level which may have occurred during the enhancement process. If the output signal levels are at or above the -12 dB reference level, the AGC will have no effect.

Overview of LCE

The Limiter/Compressor/Expander (LCE) is a three-section signal level processor allowing manipulation of the overall dynamic range of a signal. The LCE is typically used to correct for near-party/far-party or quiet talker scenarios.


The three sections correspond to three types of level processing available - limiting, compression, and expansion. Limiting is applied to the loudest levels in a signal.

Compression is the middle region, and expansion is applied to the quietest levels.

- In the Limiting region, the output signal level is “damped” to the Limiting Threshold level. When the input signal level is in the Limiting region, attenuation is applied to keep the output level from exceeding the specified Limit Threshold.
- In the Compression region, levels are adjusted so that output signal level changes are smaller than their corresponding input signal level changes. Thus, the LCE decreases the dynamic range of the signal for levels in the Compression region. As an example, a 2:1 compressor would produce an output level change of only 10 dB when the input signal changes by 20 dB. Compression is often used to correct near-party/far-party level differences, boosting the lower-level far-party speech relative to the louder near-party speech. Compression also eases listening, especially for noisy audio. Compressors are generally preferred over AGCs since input signal level differences are more closely preserved.


In the Expansion region, levels are adjusted so that output signal level changes are larger than their corresponding input signal level changes. Thus, the LCE increases the dynamic range of the signal for levels in the Expansion region. Expansion is the opposite of compression. For example, a 1:3 expander would produce an output level change of 30 dB when the input signal changes by 10 dB. A 1:2 expansion would restore a signal's dynamic range following a 2:1 compression. Expansion is also used to attenuate objectionable low-level background noise that is below the voice level.

Here is an example LCE curve. In this example, the Limiting Threshold is set at -20dB, and the Compression Threshold is set at -60dB. The Compression Ratio is 2:1, and the Expansion Ratio is 1:3.




Example LCE Curve

In each section, the LCE modifies the amplitude of the signal using a variable-gain digital amplifier. The amplitude is a rectified and smoothed version of the signal waveform, as measured by a real-time digital envelope detector. So, in the figure above, the “Input Level” actually refers to the smoothed level envelope rather than the sample-by-sample instantaneous input level. The operation of the envelope detector is governed by the Attack Time, Release Time, and Lookahead controls.




Gain Stage Main Window - AGC Mode



Gain Stage Main Window - LCE Mode

GAIN CONTROLS

Gain Mode




Gain Mode

The **Gain Mode** combo box allows the user to select either the **AGC** mode or the **LCE** mode. When switching from one mode to the other, the previous mode's controls are remembered and preserved.

AGC CONTROLS


Maximum Gain



Maximum Gain

Maximum Gain specified how much gain the AGC can apply in its attempt to bring the output signal up to the desired level. The greater the Maximum Gain, the lower the signal that can be brought up to the threshold level. The Maximum Gain range is 0-100dB. For most near-party/far-party applications, around 10dB is recommended. Settings greater than 10dB may elevate background noise to an objectionable level during pauses in speech. A "soft AGC" using of 5dB is often useful even when large voice level differences are not present.

Release Time




Release Time

Release Time controls how quickly the LCE will respond to decreases in input signal level. The shorter the Release Time, the more quickly the AGC will react. For most voice applications, a release time of about 200 milliseconds is recommended. Release Time settings less than 200 milliseconds may result in annoying "pumping" sounds as the AGC changes gain during rapid-fire conversations.

LCE CONTROLS

Compression Threshold




Compression Threshold

The level above which compression is applied to the signal. The specified

compression ratio is applied to the input signal whenever the input level is between the Compression Threshold and the Limit Threshold.


Limiting Threshold



Limiting Threshold

The level above which the signal is damped. For instance, if the Limit Threshold is -20dB, all signal levels above -20dB will be attenuated to -20dB.


Expansion Ratio



Expansion Ratio


Specifies the amount of expansion to be applied to the signal when the input level falls in the Expansion Region. The Expansion Ratio is expressed as a ratio 1:N. Jumps in the output signal are N times larger than their corresponding jumps in the input signal. For example, with an Expansion Ratio of 1:3, a 10 dB jump in input level becomes a 30dB jump in output level.

Compression Ratio



Compression Ratio


Attack Time



Attack Time

Controls how quickly the LCE responds to increases in input signal level. For a more peak-sensitive processor, use a short Attack Time. For a more average-sensitive processor, use a longer Attack Time. For most speech applications, a fast Attack Time of 2-5 milliseconds is recommended.

Release Time




Release Time

Controls how quickly the LCE responds to decreases in input signal level. Short Release Times (<100 milliseconds) can create an annoying "pumping"

artifact as the level detector is too responsive to intra-syllabic pauses. Long Release Times (>500 milliseconds) may fail to respond to breath group pauses and exchanges between speakers. For most speech applications, a Release Time of 200-400 milliseconds is recommended.

Look Ahead Time



Look Ahead Time

Lookahead controls the alignment of the envelope detector with the output signal. Since the envelope is a smoothed version of the signal waveform, level changes in the envelope will lag corresponding changes in the signal itself. The applied LCE gain depends on the envelope level, so the same lag is reflected in the applied gain.

The Lookahead control adjusts an internal delay that compensates for this lag. The larger the Lookahead setting, the earlier the gain adjustments will be shifted. For most speech applications, a Lookahead of 1-5 milliseconds is recommended.

Reset Button



Reset Button

The **Reset** button resets all the LCE parameters to their default values, which are:

Parameter	Default Value
Limiting Threshold	-10dB
Compression Threshold	-50dB
Compression Ratio	1.5:1
Expansion Ratio	1:1.2
Attack Time	50ms
Release Time	200ms
Look Ahead Time	5ms

Store Button



Store Button

This button allows the user to store the current configuration to a user-specified disk file that will not be lost when the computer is turned off.

Recall Button



Recall Button

This button allows the user to recall a previously stored configuration from any of the saved disk files previously generated using the **Store** button.

PARAMETRIC EQUALIZER

OVERVIEW OF THE PARAMETRIC EQUALIZER


The Parametric Equalizer consists of a variable number of IIR filter stages, connected in series, which can be used for boosting (or peaking) and cutting (or nulling) portions of the input signal's frequency spectrum. Each stage is described by a center frequency, a frequency width, and a boost/cut amount, and can be configured independently. A common application of the parametric equalizer is to construct a precision notch filter which will perform nulling of the input signal at the specified center frequencies.

In the Current Stage block, the eight available stages can be selected one at a time to adjust their individual configurations. Individual stages can be toggled between Active and Inactive. An active stage is applied to the audio, while an inactive stage is bypassed. When a stage is made inactive, its settings are preserved.

HINT: It is often helpful to activate only one stage at a time when adjusting the stage settings. Then, once satisfactory settings have been found for each individual stage, all stages can be activated for audio processing.

When multiple stages are in use, their effects can overlap so that the overall signal level is reduced or boosted more than expected. For this reason, an output gain control is available as part of the Parametric Equalizer, allowing the user to compensate for overall level changes that may result from Parametric Equalizer filtering. (Advanced users may note that many


Parametric EQ filters provide an input attenuation control so that fixed-point saturation can be avoided. Since CARDINAL uses a floating-point implementation, saturation is not a concern, so only output level adjustment is provided.)



Parametric Equalizer Main Window

FILTER CONTROLS


Current Stage



Current Stage

The buttons in the **Current Stage** box allow the user to select the current equalizer stage for adjustment. When a stage is selected, its settings populate the **Center Frequency**, **Width Factor**, **Boost/Cut** and **Stage Active** controls.

Center Frequency



Center Frequency

The frequency at which the current stage's boost/cut region is centered.


Width Factor



Width Factor

A factor controlling the width of the current stage's boost/cut region.

Boost/Cut



Boost/Cut

The amount of boost or cut to be applied by the current stage.

Stage Active



Stage Active

If the indicator is lit, the current stage is being applied to audio. If the indicator is dark, the current stage is bypassed.


Reset All



Reset All

Returns all stages back to their default settings, which has no effect on the audio (Center Frequency: 0Hz, Boost/Cut: 0dB, Stage Inactive).

Output Gain



Output Gain

Amount of gain or attenuation applied to the audio after all active parametric EQ stages have been applied.

Show Inactive



Show Inactive

Displays an orange trace on the graph to indicate what the transfer curve looks like *if* all the stages are enabled.

Store Button



Store Button


This button allows the user to store the current configuration to a user-specified disk file that will not be lost when the computer is turned off.

Recall Button



Recall Button

This button allows the user to recall a previously stored configuration from any of the saved disk files previously generated using the **Store** button.



CHAPTER 7

EDITING VIDEO

New video sessions can be created by selecting segments from other sessions and executing simple cut, copy and paste operations. In addition to editing movies together, you can edit movies and still images together. These new sessions can then be exported from VideoFOCUS in many different formats including QuickTime (.mov) files, AVI (.avi) files, Windows Media (.wmv), MPEG-4, etc.

Video editing in VideoFOCUS is non-destructive, meaning that the original media is never modified.

EDITING OVERVIEW

Editing begins by creating a new blank editable session into which other session movies, sections of session movies, still frames or titles can be added.


Frames and groups of frames can be selected from session movies either in the frame thumbnail view or the video player view. The user only has to select a session and browse through the video to find any frames of interest. Once located, they can be easily selected, and copied. They are then pasted into a new editable session created by the user.

All actual editing takes place in the new editable session. In the editable session, the user can paste, copy and cut frames (there is no cutting or pasting of frames in regular session movies, only copying). Once editing is finished, the new session is saved and appears in the Captured and Imported Sessions window. The user then has the option of leaving it open to add or subtract further sequences, or closing it. When the editable session is closed, all unsaved changes are lost.

STARTING THE EDITING PROCESS

CREATING A NEW BLANK EDITABLE SESSION


1. Select *File->New Session*.



A new blank session will be created into which whole sessions, collections of frames, images and titles can be added.

CREATING A NEW EDITABLE SESSION FROM AN EXISTING SESSION

1. Select the session you wish to edit in the Captured and Imported Sessions Window
2. Choose *Processing->Edit*



An editable copy of the original session is created into which whole sessions, collections of frames, images and titles can be added. The original session is never changed but a new edited version is created upon save.

ADDING MEDIA TO THE NEW EDITABLE SESSION

Once you have created a new blank session or an editable version of an existing session movie, you can begin to add other media to it.

Media can be dragged and dropped onto the new session object. It will be inserted wherever the video playhead is located.

After saving a session, editing can continue as long as the new session window stays open. Any number of successive sequences from other videos can be added or subtracted, as needed. For instance, you could choose another video from the session list, select and copy a group of frames in the same way, and paste those into the same video.

TO ADD A WHOLE SESSION MOVIE

1. Move the video playhead to the location at which you would like to insert a different session movie.
2. Select the session movie you want to add in the Captured and Imported Sessions Window.
3. Drag the session movie into the new session window. The session movie will be added at that location.

4. Click Save. The new session appears in the session list. Close the new session window.

Note: *Session movies can be added to a new editable movie by either dragging and dropping or by copying and pasting using CTRL-C and CTRL-V or the corresponding menu items.*

TO ADD A SECTION OF A SESSION MOVIE FROM THE VIDEO PLAYER

1. Select the Player tab at the top of the application window so the video player displays.
2. Find the first frame in the sequence to be copied.
3. Click on the scrubber, and hold the shift key down.
4. Drag the playhead slider to the last frame in the sequence to be copied.
5. As the slider is dragged, the region between the first and last frames displays an orange selection highlight.
6. Choose *Edit->Copy* (Ctrl-C) to copy the frames.
7. Choose *File->New Session* (Ctrl-N) to create a new session, if one is not open.

TO ADD A SECTION OF A SESSION MOVIE USING THUMBNAIL VIEW

Selecting frame sequences in the thumbnail view is just as easy as selecting from the video player.

1. Select the Frames tab at the top of the application window so the thumbnails display.
2. Find the first frame in the sequence to be copied.
3. Click on the frame, and hold the shift key down.
4. Find the last frame in the sequence to be copied.
5. Click on the frame, and the region between the first and last frames displays an orange selection highlight.
6. Choose *Edit->Copy* (Ctrl-C) to copy the frames.
7. Choose *File->New Session* (Ctrl-N) to create a new session.
8. In the new session window, choose *Edit->Paste* (Ctrl-V) to paste the sequence.
9. Click Save. The new session appears in the session list. Close the new session window.
10. In the new session window, Choose *Edit->Paste* (Ctrl-V) to paste the sequence.
11. Click Save. The new session appears in the session list. Close the new session window, if desired.

TO ADD A STILL IMAGE INTO A SEQUENCE

Editing a still image into a sequence is also as easy.

1. Find the still image to be copied.
2. Click on the image so it's selected.
3. Choose *Edit->Copy* (Ctrl-C) to copy the image.
4. Choose *File->New Session* (Ctrl-N) to create a new session, if one is not open.

5. In the new session window, Choose *Edit->Paste* (Ctrl-V) to paste the still image.
6. Click Save. The new session appears in the session list. Close the new session window, desired.

TITLING

Titles can easily be added, edited, located and deleted in an editable session.

TITLING ACTIONS

The following titling menu actions are available in player window when editing a session.

New Creates a new blank title at the current playhead position.

Edit Edits an existing title (if a title is at the playhead position)

Delete Deletes an existing title (if a title is at the playhead position)

Go Positions the playhead to the location of existing titles within the session The 'New' and 'Edit' actions bring up a title editor.

ADJUSTING TITLE APPEARANCE

Just as in editing with a typical text editor title text can be entered directly into the editor and will appear at the cursor location. The following elements control the title's appearance.

Justify Controls whether the text will appear on the left, center, or right

Size Controls the size of the text

Duration The number of seconds the title will appear in the session

To adjust existing text the text must be first selected (in the typical way for a document editor) before the 'Justify' or 'Size' values are changed.



FINE TUNING FRAME SELECTION


For a more precise way of reviewing a movie or selecting a frame sequence, there are further options.

SCALE SLIDER

As some videos may be several minutes in length, the scale slider in the lower left corner of the video player window allows the user to adjust the video player timeline to display more, or fewer, frames while moving the playhead slider.



When moving the scale slider away from its default position on the left, a Selected Area slider appears. When moving the playhead slider, you will see whatever is in the selected region. To change that selected region, you can move the Selected Area slider to another region. To reset to the original state, move the Scale Slider back to its default position on the left.



Scale Slider zoomed in

Selected Area


MARK IN/OUT BUTTONS

After zooming into a sequence, it may no longer be practical to shift-select sequences of frames, because it may not be possible to easily scrub the playhead all the way to the desired in or out point. In this case, it's best when locating the beginning of a sequence to mark it, by hitting the Mark In button, and then locating and marking the end by hitting the Mark Out button. The end result is that all the frames in between will be selected and ready for copying. These buttons are found on both ends of the video player controls.

Mark In



Mark Out



CHAPTER 8

EXPORTING MEDIA


You can export still images, videos, and audio from VideoFOCUS to any accessible drive with sufficient space.

EXPORTING STILL IMAGES

After creating a filtered, modified or super-resolution still, it can be exported as a BMP, JPEG or TIFF file to archive or distribute it. It can also be opened in the operating systems default program for images.

TO EXPORT A STILL IMAGE

1. Select a still image in the Session Stills and Frames window.
2. Choose *File->Export Image* to open the Export Image dialog box.



3. Select a file format, BMP, JPG, or TIF.
4. Choose a path to send your file to a location of your choice.
5. Select the *Aspect Correct* option if you want to ensure that the exported image is rendered to the same aspect ratio as it appears in VideoFOCUS.
6. If you also want a PDF report produced alongside the image file, select the *Export Report* option.
7. Enter a name for the file, and click Export.

EXPORT IMAGE DIALOG BOX OPTIONS

The following table describes the options in the Export Image dialog box.

Export Image dialog options	
Look in	Choose a location to save the exported file.
File name	Enter a name for the file. The default value used is the name of the image.
Files of type	Choose a file format: • BMP • JPEG • TIFF
Aspect Correct	Check if you want to ensure that the saved file is aspect corrected to match its appearance in VideoFOCUS; otherwise, particularly for images extracted from interlaced video files, the height may be squashed.
Export Report	Check if you want to have a PDF report saved alongside the image file to detail how it was produced. The report includes MD5, SHA1, and SHA-256 hashes of the associated file(s).

OPEN STILLs IN SYSTEM DEFAULT IMAGING SOFTWARE

You can send any image in the Sessions Stills and Frame area to be opened by the operating system's default software for images. Typically this will be a software program such as Photoshop, Imaging or Gimp. This is useful if you would like to do further processing on your image in another software tool. You can then save the image to your hard disk from directly in that program.

To Open an Image in the System's Default Software for Images

1. Select the still image in the Session Stills and Frames area.
2. Right click the mouse to bring up the menu dialog.
3. Select *Open Image in System CTRL-O*.

DRAG & DROP

Stills and frames can also be dragged from the application onto the desktop or into another open software program. Drag and drop currently does not work on movies.

EXPORTING MOVIE FRAMES

You can export individual unprocessed video frames from any movies you have acquired in VideoFOCUS as JPEG, Bitmap or TIFF files.

To Export a Single Frame

1. Select a session in the session list.
2. Select an individual frame in either the thumbnail view or in the video player.
3. Choose *File->Export Image* or Right-Click to access the same menu.

4. Select a file format, BMP, JPEG, or TIFF.
5. Choose a path to send your file to a location of your choice.
6. Enter a name for the file, and hit Export. It will export using the same pixel dimensions and DPI value as the source video.


Note: This is only for unprocessed raw images from your movies. If you have processed a frame it will appear in the Session Stills and Frames area of the application. If that is the case, follow the export instructions listed above

EXPORTING SESSION MOVIES AND AUDIO

You can export the movies you have acquired in VideoFOCUS to many different compressed or uncompressed video file formats. This feature is ideal for converting a library of video files into different formats (QuickTime to AVI, MPEG-4, etc.) for whatever reason, be it for archival or distribution purposes.

To Export a Session Movie

1. Select a session in the session list
2. Choose **File->Export Video**. The Export Video dialog appears.



3. Choose a profile under Files of type. We have formats that support a variety of compressors, or codecs.
4. Alter the Frame Size, if necessary. The actual size is always the default when the dialog opens.
5. Choose to Export Report or not. If the box is checked, a PDF file containing a full report for the exported item will be saved alongside the exported media file.
6. Choose your file location of your choice.
7. Choose Export to save and your new video will be created.

If the exported video is for archive purposes, or if additional work is to be done in

other software, it is suggested that you export to an uncompressed format. If you would like to email the video, choose a compressed format.


EXPORT VIDEO DIALOG BOX OPTIONS

The following table describes the options in the Export Video dialog box.

Export Video dialog options	
Look in	Choose a location to save the exported file.
File name	Enter a name for the file. The default value used is the name of the video session.
Files of type	Choose one of the standard profiles using the dropdown menu, or click Advanced and configure a Custom Profile .
Frame Size	Choose the desired W x H pixel resolution and the session will be resized to match while exporting the video to the file. If the proportions do not match that of the session, the exported video will be Letter Boxed automatically, or you can force letterboxing by selecting it.
Export Report	Check if you want to have a PDF report saved alongside the image file to detail how it was produced. The report includes MD5, SHA1, and SHA-256 hashes of the associated file(s).

CUSTOM PROFILE DIALOG

You can create, delete, view and update your own custom profiles, where you can choose your own format, video codec, and audio codec. Just click on the Advance button to bring up the Custom Profile Dialog.



COMPRESSION FORMATS (CODECS)

The compressor (codec) you choose depends on the intended purpose of the media file. For example, to save files for archival purposes, in the best possible quality, it is best not to compress the files at all, by choosing a profile that has a raw video codec option. Otherwise, the compression profile formats listed in the export dialog are standardized among video applications; some are widely used, while others are less so.

Recommended Compression Formats - QuickTime

These files, appended with the .MOV extension, will open in QuickTime player.

1. **Uncompressed (Raw):** No compression, best image quality, produces a large file. The large file size often prohibits it from being a reasonable format if the file is intended for electronic transfer, via email or the web.
2. **JPEG 2000:** JPEG compression, high quality, smaller file size. The loss of quality is minimal.
3. **MPEG-4:** MPEG compression, moderate quality, small file size. The loss of quality is pronounced, but the small file size makes the format ideal for electronic transfer, via email or the web.
4. **H264:** H264 compression, high quality, small file size. The loss of quality is minimal, and the small file size makes the format ideal for electronic transfer, via email or the web.

Recommended Compression Formats - AVI

These files, appended with the .AVI extension, will open in Windows Media Player or QuickTime player.

1. **Uncompressed (Raw):** No compression, best image quality, produces a large file. The large file size often prohibits it from being a reasonable format if the file is intended for electronic transfer, via email or the web.
2. **JPG2K Compressor:** JPEG compression, high quality, smaller file size. The loss of quality is minimal if the quality slider is set to "best".

EXPORTING AUDIO

If you have files that contain an audio track, you can export them along with the video as a QuickTime (.mov) file, AVI or most other formats listed.

If you wish to export the audio only, you will first need to select the session and click *Processing->Extract Audio* to produce an audio-only session. Then, you can select *File->Export Media* and select a profile with a wave WAV format. This will create a WAV file that contains the audio from the movie.


BATCH EXPORT OF MOVIES AND STILLs

You can export any number of movies or stills at once, or in a succession of concurrent exports.

To Batch Export Session Movies

1. Select the sessions in the session list. Ctrl-click for non-contiguously ordered session movies, Shift-click for contiguously ordered session movies.
2. Choose *File->Export Video*. The Export Video dialog appears.
3. Choose a file selection format (AVI or QuickTime). Both QuickTime and

AVI support a variety of compressors, or codecs.




When batch exporting, the File Name field is not editable.

4. The file names for the selected movies in a batch export will be preserved in the export process and cannot be edited in the export panel. If you wish to change the name of a file for batch export, do so before exporting it.
5. As the session movies export, a progress bar appears. If multiple session movies are exported in a batch, each session movie will be concurrently exported.

To Batch Export Images

1. Select the images in the Session Stills and Frames list. Ctrl-click for non-contiguously ordered session movies, Shift-click for contiguously ordered session movies.
2. Choose *File->Export Image*. The Export Session Movie dialog appears.
3. Choose a file selection format (TIF, BMP or JPG).
4. The file names for the selected images in a batch export will be preserved in the export process and cannot be edited in the export panel. If you wish to change the name of a file for batch export, do so before exporting it.
5. As the images export, a progress bar appears. If multiple session movies are exported in a batch, each image will be concurrently exported.



CHAPTER 9

CREATING REPORTS

For any imported, captured, or processed sessions and still images in VideoFOCUS, a comprehensive report can be created that details what has been done and the exact steps needed to reproduce.

TO CREATE A REPORT FOR A SESSION

After creating a filtered or modified session in the Captured and Imported Sessions area, first ensure that all of the desired fields have been entered under both the Case Properties and the Session Properties, then produce the report by clicking *View->View Report*. A preview PDF report will then open in a browser window; if you are not satisfied with all of the report text, you may close the browser and go back to the relevant Case and/or Session Properties to edit the required fields and then try viewing the report again.

When you are finally satisfied with the report text, you can produce a final PDF report that gets saved along with the exported video file by checking the “Export Report” option in the Export Media window.


Once the PDF has been saved alongside the video file, when the PDF is opened, clicking on any of the pictures will cause the video and/or image file to be opened in the system player or viewer associated with the video or image file type.

TO CREATE A REPORT FOR A STILL IMAGE

After creating a filtered or modified image in the Session Stills and Frames area, first ensure that all of the desired fields have been entered under the Case Properties, the Session Properties for any parent sessions, and finally the Image Properties, then produce the report by clicking *View->View Report*. A preview PDF report will then open in a browser window; if you are not satisfied with all of the report text, you may close the browser and go back to the relevant Case, Session, and/or Image Properties to edit the required fields and then try viewing the report again.

When you are finally satisfied with the report text, you can produce a final PDF report that gets saved along with the exported image file by checking the “Export Report” option in the Export Media window.

Once the PDF has been saved alongside the video file, when the PDF is opened, clicking on any of the pictures will cause the video and/or image file to be opened in the system player or viewer associated with the video or image file type.



Example VideoFOCUS PDF Report



APPENDIX A

TROUBLESHOOTING


This appendix contains information about contacting Salient Sciences Technical Support.

CONTACTING TECHNICAL SUPPORT

For phone support, call (+1) 919.572.6767 (option 2) from 9:00 AM to 5:00 PM, EST.

For online support, email support@salientsciences.com

As a Salient Sciences customer, you have access to our online support at support.salientsciences.com to ask questions, look up answers to frequently asked questions and download documentation.



APPENDIX B

VIDEOFOCUS REGISTRATION

LICENSE INFORMATION

VideoFOCUS can license your software in two different ways:

HARDWARE DONGLE REGISTRATION

VideoFOCUS can be issued with a dongle. If you have a dongle:

1. Install the HASP drivers. If you need a HASP driver, please go to our FAQ webpage at <http://support.salientciences.com/faq.html> for instructions on how to download the latest driver from the manufacturer.
2. Insert the dongle in any empty USB port. Ensure that the red light on the dongle is solid before proceeding with use of VideoFOCUS.
3. Launch VideoFOCUS. The license on the encoded dongle should be accepted.
4. If you are unable to register and are giving a warning that states "No Hardware Key Detected," please double check that the drivers are installed correctly and that the dongle is recognized by your computer.

SOFT REGISTRATION & CONNECTED TO THE INTERNET

If you were not issued a dongle with your copy of the software, the first time the application is launched, VideoFOCUS will see that "No Hardware Key" is detected and it will attempt to license your product over the Internet.

1. Launch VideoFOCUS.
2. A dialog will be displayed requesting an installation ID. This Installation ID will have been provided to you by Salient Sciences.
3. After entering the Installation ID, follow the prompts to fully register your product. Be sure to include end user information including: Company Name, Contact Name, E-mail, and Phone Number
4. VideoFOCUS will now launch.

SOFT REGISTRATION & NOT CONNECTED TO THE INTERNET

If you were given an Installation ID but your machine is not connected to the Internet you can still activate your product.

1. Launch VideoFOCUS.
2. A dialog will be displayed requesting an installation ID. This Installation ID will have been provided to you by Salient Sciences.
3. Follow the instructions as prompted by the VideoFOCUS registration installer.

If you experience any issues registering VideoFOCUS, please email support@salientciences.com or call support at (+1) 919.572.6767 (option2).

UPDATE YOUR LICENSE

If you are not connected to the Internet, you can update your license when you receive a new activation code from Salient Sciences.

1. Choose *Help->Update License*.
2. Enter in your activation code. If you do not have one, contact support@salientosciences.com.