

MIT OpenCourseWare
<http://ocw.mit.edu>

21M.361 Composing with Computers I (Electronic Music Composition)
Spring 2008

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.

21M.361: Composing with Computers I (Electronic Music Composition)

Peter Whincop

Spring 2008 OCW

[All lab notes are being significantly revised for next term to reflect upgrades in software, different organization, and the incorporation of my own notes on DSP.]

Lab 3.2: Recording Using Peak; Peak Plug-Ins: SuperFreq-10, SoundSoap; Noise Reduction, etc.

Peak Recording:

1. As usual, check the Sound Out under the Audio Menu.
2. Check the Audio>Record Settings. (i) The Record Disk must have some disk name. If mta-studio1-hd isn't listed, click on the drop-down menu and see if it is not grayed-out. If it is, then there are permissions problems and you should contact me so I can fix them. Otherwise make sure the disk is checked and is 'primary.' (ii) File Format should be set to AIFF. (iii) The other checkable fields are self-explanatory, though I would have Use Record Timer... checked, and Open after Saving checked. (iii) Clicking on Device and Sample Format will give you the chance to make sure 96 in 1&2 are being used, and set the sample rate, bit depth, and stereo/mono.
3. To record, go to Audio>Record.... First check that a signal is getting into Peak by looking at its level meter in the transport window. Recording begins when you click on the circle (record) button in this window. Once you do that, a waveform will start scrolling jaggedly at the top of the window, usually with a latency of around 1s. Make sure you see some activity. Recording stops when you press the square (stop) button in this window. You will be asked to save the new Audio Recording, and if you checked the right box under Record Settings, the newly recorded and saved soundfile will open in Peak.
4. Two other things I can't think of where to put: (i) Under DSP you can split stereo files, or convert a single mono file into a stereo file with a fixed but settable panning. (ii) If when you try to copy to the clipboard Peak has a fit at you about not having enough disk space, there are either permissions problems which I will have to fix, or the Scratch Disk has been changed. Go to Peak>Preferences... and correct it.

Peak Plug-Ins:

1. Plug-ins can be found under their own menu. You can use five plug-ins, and they are serial just as in Pro Tools. The BIAS (the manufacturer of Peak) plug-ins are most useful for now. Once you have selected a plug-in, a check mark will be placed beside it in the sub-menu. Even if you click off the plug-in window, it will still be active. You have to select 'none' in the appropriate sub-menu to disable it. You can see which plug-ins you have active under the Windows menu, just above the open soundfiles list.
2. You play your sound and change the plug-in settings to your satisfaction. But you can't (yet) record dynamic changes to your soundfile using the plug-in. All you can do is bounce them, found under the plug-ins menu. This will alter the file (in memory), so remember to Save As... with a different name. After it has bounced, Peak will ask you if you want to disable the plug-ins. You should, otherwise it will apply (but not bounce unless you tell it to) the same plug-ins to the sounds you have just applied the plug-ins to.
3. If you are daring, look up V-Box. It is under the plug-ins menu, and there is documentation on how to use it. Basically it is a way of using multiple plug-ins in series and parallel, according to a flow diagram you create yourself. We might cover it one day if we have time.

SuperFreq-10:

1. This is a stupidly-named filterbank consisting of 10 filters. There are other SuperFreqs, 4, 6, and 8, referring to the number of

individual filters.

2. You can change the visual order of the individual filters, but the control order (all the knobs) remains the same.
3. Each filter can be bypassed by clicking at the bottom of the appropriate column. The entire plug-in can be bypassed by clicking in the Bypass button at the bottom of the window.
4. You can set each filter type to peak (bandpass or band-reject/notch), highpass, highshelf, lowpass, or lowshelf. The most useful one is the peak filter. It functions rather like the SND filterbank, or the parametric EQ on the rack. High or low filters of either type are more appropriate at the ends. A shelf filter boosts high or low, leaving everything at 0dB change, i.e., no change. A high or lowpass filter keeps their end of the spectrum at 0dB and reduces everything else.
5. You can change frequency and dB of each filter visually/using the mouse. Q (quality, resonance, or slope) must be changed using the knob. Q basically says how wide a filter will be, how focused it will be on the frequency you select. Low Q is wide, high Q is narrow. More technically, Q is the ratio of center frequency to bandwidth, which is defined as the difference in frequencies which correspond to a 3dB reduction in amplitude.
6. Changing one filter appears to change other filters; the dots don't move but the curve does. The curve is the aggregate frequency response of the filter; areas shaded light blue are where changes occur: above the 0dB line, the sound is amplified, below, attenuated.
7. If you are bumping up a whole of frequencies, which is always tempting, you might make the sound clip. You can reduce the overall output with the fader on the right.
8. Most plug-ins have Presets, Save and Load buttons at the bottom of the window. Save, well, saves the settings. Load, idem. Presets, again, obvious.

Noise Reduction Using SoundSoap:

1. There are two SoundSoaps, SoundSoap Pro and SoundSoap (2.2). The latter is far simpler, so I'll start with that. It is just a simple version of part of SoundSoap Pro, with an easier interface, but fewer options.
2. SoundSoap calls what it does 'Broadband' noise reduction. Here is how it works. When you record anything, you should always record a few seconds of the ambient sound, e.g. the classroom without anyone talking. This is the noise you wish to remove from the entire recording. SoundSoap 'learns' that noise by analyzing the frequencies of the noise over a short period of time, averaging it so that time is no longer a variable, and subtracts the amount of each frequency it analyzed from the sound you are cleaning up. It's a pretty amazing way of doing it. And it doesn't do too much damage to what you want to keep: if, say, there is a certain amount of 700Hz in the noise, and some in my voice, it will take out the right amount, since it knows how much there was in the noise. Of course this works best when the noise is fairly consistent, like from a bad recording device, a/c sounds, even fairly homogeneous traffic sounds. It is imperfect though. Noise has a time-varying character; it's fairly random, for all intents and purposes. One chunk of noise might be to our ear exactly the same as the one that follows it, but it most likely contains slightly different amounts of every frequency. A result of this is the melodic tinkling sound I played from Cat Power's Baby Doll (from You Are Free)—now on the MIT Server. There are ways around this, discussed below, but generally out of our realm for now.
3. Set the cursor on the waveform to where there is noise. Click on Learn Noise, and press play (spacebar). Peak will learn the amount of each frequency in this ~1s sample. After that, when you play your sound with Broadband On, that level of noise will be removed from the sound as it is playing. Since it is a plug-in, you would have to bounce it to retain the noise-reduction effect. If you

do not have the consecutive ~1s of noise that SoundSoap needs, you can duplicate the smaller fragment a number of times, just so it can have its ~1s.

4. You can switch Broadband (noise reduction) on and off, and you even have the option of keeping only the noise. This can be an interesting source of compositional material: if you bounce the plug-in with noise only selected, that's what you will get.

5. You can 'tune' the noise reduction. According to the manual, it "Helps you 'tune in' the hiss and background noise you want to remove. The Learn Noise button sets this control automatically—or you can adjust it manually." You can also alter how much noise reduction is applied: when the knob is vertical, it takes out the amount of noise it learned—"it is set automatically by Learn Noise (or you can adjust it manually)." If you increase the knob's value, it will take out more noise, in other words, more of each frequency that it detected as noise. This can be compositionally interesting in two ways—you can alter your sounds, or the amount of noise reduced is greater, therefore the noise only option gives you more to play with. Decreasing the knob's value will take out less noise, in other words, it will leave some noise in your sound. Playing with the Tune Noise knob can yield compositionally unexpected results too.

6. You can Preserve Voice, "which is a special pre-filtering mode that removes frequencies outside of the range of the human voice." Remove Rumble "takes out low-frequency noise of 40Hz and below—often found as background noise in many DV clips." Sound Enhancer does what it says, in that it "adds 'sparkle' back into the audio signal. It can help restore frequencies that may have been lost in old or degraded media." It must be psychic. I imagine it is a little like iTunes' feature with the same name. "Hum Removal is used for eliminating hum that was introduced into digital audio and video files from electrical interference. Use 60Hz for North and South America," and 50Hz for most other places. I have used Hum Removal to good advantage in the past (due to my sloppy wiring of toys). Remove Click and Crackle does exactly that, though not very well. You set a threshold level—moving it upwards remove more clicks and crackle, but possible does (undesirable?) things to your sound.


6. You don't just have to select noisy parts of your sound for the Learn Noise process; you can choose to play any ~1s chunk. Experiment with this for interesting compositional results. But in general, it is still advisable to keep a few seconds of noise lying around in case you want to NR a recording. I sometimes save them as separate files, for future use.

7. SoundSoap Pro has four modules. You activate each module by checking its box at the bottom. You bring up its sub-window by clicking above the box. Take care with this. It is easy to set a module up, but not check its box, while another module isn't visible, but checked.

8. Both SoundSoap and SoundSoap Pro can be used as Audiosuite effects, and as RTAS mono or stereo inserts, in Pro Tools. This is pretty fantastic; as inserts you can automate many of its parameters, thus change noise reduction amount, noise reduction on or noise only, etc. And with the non-Broadband modules of SoundSoap Pro, all kinds of fantastic things (like variable threshold noise gating).


9. Here is the lazy way of explaining the four modules.

SoundSoap Pro User Interface: Hum & Rumble Tab



Hum & Rumble Tab

On/Bypass Radio Buttons/Checkbox – These controls allow the Hum & Rumble tool to be turned on or off, while other tools remain active. These controls are useful as they allow this tool to simply be turned off when not required – and are also useful for comparing the difference between having the Hum & Rumble tool on or off.



Hum & Rumble Tools may be turned on or bypassed using the radio buttons in the tool area, or with a checkbox in the tab

Hum Meter – SoundSoap Pro’s Hum Meter is another visual feedback tool, which along with your ears, is a useful tool for helping to locate the fundamental frequency causing a particular hum. Typically, most hums occur at either 50Hz or 60Hz – however, there may be cases in which they occur at other frequencies. The Hum Meter works in conjunction with the Hum Frequency Slider. To find the fundamental frequency causing a hum, simply sweep the Hum Frequency Slider while listening to a media file containing hum. As you get closer and closer to the fundamental frequency, the Hum Meter’s reading intensifies. When you see the strongest reading in the Hum Meter, you’ve found the fundamental frequency responsible for causing the hum, and can then use other Hum & Rumble controls to adjust how much the hum is reduced.



The Hum Meter provides visual feedback that helps identify the frequency causing a hum

Hum Frequency Slider – The Hum Frequency Slider allows you to reduce a specific frequency, and has a range of 20Hz – 200Hz. Moving this slider determines the fundamental frequency that will be targeted (a specific value may also be entered by typing a value in the numerical entry field). If the frequency causing a hum is unknown, the Hum Frequency Slider may be swept as the noisy media file is played, and the Hum Meter, located directly above the Hum Frequency Slider, will register activity as you approach the offending fundamental frequency. The setting that gives the strongest reading in the Hum Meter is the frequency (or very close to it) causing the hum problem.



Parameters can be adjusted with a slider, by entering a numerical value, or using the Up/Down Arrow keys on your keyboard

! *When typing in values or using the Up/Down Arrow keys to adjust values in a numerical field, be sure to press the Return or Enter key on your keyboard.*

Q (Bandwidth) Slider – The Q Slider determines the bandwidth of the notch filter being applied to a particular frequency. The range of the Q slider is from 10 to 50 – A larger Q value affects only a very narrow frequency range, and a smaller value affects a wider frequency range. Hum will vary depending on the source – a very narrow Q (higher value) will eliminate only the frequency set with the Hum Frequency Slider, and will do very little to frequencies that are very close, either higher or lower. On the other hand, a wider Q (smaller value), will still affect the same fundamental frequency, but will also affect some of the surrounding frequencies. Ideally, the narrowest Q (largest value) possible should be used, so as to affect just the offending frequency. If the Hum Meter shows a strong reading, but some hum is still audi-

ble, it's best to experiment with the Q value setting, as well as the Depth and Harmonics Sliders.



Q (bandwidth) values can be adjusted with a slider, or by entering a numerical value

Depth Slider – The depth slider controls the amount of audio being cut out by the notch filter responsible for eliminating hum, and is directly related to both the Q Slider and Hum Frequency Slider. The Depth slider has a range of 0dB (no audio being cut out) to 50 dB of reduction. The Depth setting cuts the most amount of audio at the fundamental frequency, as determined by the Hum Frequency Slider, and affects the surrounding frequencies proportionally, depending on the Q value that is used.



Notch filter depth can be adjusted with a slider, or by entering a numerical value

Harmonics Slider – The Harmonics Slider designates how many harmonics of the fundamental frequency should be removed in either direction from the fundamental frequency. The range of harmonics that can be removed is 1 – 9, always as integers, never as fractions. When dealing with hum, there is always a fundamental frequency, and commonly there are harmonics at other frequencies, due to resonance. Harmonics typically will have similar characteristics to the hum found at the center frequency, but will be proportionally lower in amplitude. For example, a 60Hz hum may have harmonics at higher frequencies, such as 120Hz, 180Hz, 240Hz, and so on. Assuming the fundamental frequency is 60Hz, a setting of “3” on the Harmonics Slider would also remove the hum that occurs at 120Hz and at 180Hz.



SoundSoap Pro can eliminate up to 9 harmonics of the fundamental frequency

Harmonics Tilt Knob – The Harmonics Tilt Knob allows you to fine-tune the relationship between how much of the fundamental frequency is cut, and how much of the signal is cut at each harmonic.

It is common for a hum that occurs at a certain harmonic of the fundamental frequency to have a lower amplitude than the hum at the fundamental frequency. The Harmonics Tilt Knob controls the relationship between the depth of each successive notch filter created by the Harmonics slider. With a value of “0.00”, each additional notch created by raising the value of the Harmonics slider cuts at the same value as the notch placed at the fundamental frequency.



The Harmonics Tilt Knob allows fine-tuning of harmonic reduction

For example, if the Hum Depth slider is set to cut by 20dB, and the Harmonics Tilt Knob is set to “0.00”, each additional notch created by increasing the value of the Harmonics slider will also cut each harmonic by 20dB. However, if the Hum Depth slider is set to cut by 20dB, and the Harmonics Tilt Knob is set to “0.50”, then each successive notch created by increasing the value of the Harmonics slider will cut at half the depth of the previous notch.

Rumble Reduction Slider – The Rumble Reduction Slider focuses on very low frequency noise, commonly known as “rumble”. Rumbles typically occur around 40Hz or below, and SoundSoap Pro’s Rumble reduction tool has a range of 20Hz – 100Hz. The Rumble reduction tool is a high-pass filter (it lets high frequencies pass through, while blocking low frequencies), with a slope of 12dB and a Q value of .5.




Separate Rumble reduction control is available for very low frequency noise

Global Tools - As with all the other tool tabs, SoundSoap Pro's global tools may be accessed from within the Hum & Rumble tab. These global tools include Preset Buttons, A/B/C/D Compare Buttons, Help Button, Noise Only Mode Button, Spectrogram, and the ability to monitor other tool settings with the Tool Tabs/Tool State Thumbnails. These tools are described in more detail at the beginning of this section.

That concludes the overview of the Hum & Rumble Tool – the next section covers Click & Crackle reduction with SoundSoap Pro.

3

SoundSoap Pro User Interface: Click & Crackle Tab



Click & Crackle Tab

On/Bypass Radio Buttons/Checkbox – These controls allow the Click & Crackle tool to be disabled, or bypassed, while other tools remain active. This control allows the tool to simply be turned off when not required, to save CPU power for other tasks. It is also useful for comparing the difference between having the Click & Crackle tools on or off.



Click & Crackle Tools may be turned on or bypassed using the radio buttons in the tool area, or with a checkbox in the tab

Click Threshold Slider – The Click Threshold slider dictates how bad clicks must be to be detected and removed. The range of the Click Threshold slider is 1.00dB – 25.00dB. When clicks are present in the media file being restored, if they fall below the threshold value, they remain unaltered – if they're above the threshold, they are removed. The general rule with the Click Threshold slider is, moving it to the left will remove more clicks, and moving it to the right removes fewer clicks.



The Click Threshold slider allows fine-tuning click reduction

Crackle Threshold Slider – The Crackle Threshold slider dictates how bad crackle must be to be detected and removed. The range of the Crackle Threshold slider is 12.0dB – 25.0dB. When crackle is present in the media file being restored, if it falls below the threshold value, it remains unaltered – if it's above the threshold, it's removed. The general rule with the Crackle Threshold slider is, moving it to the left will remove more crackle, and moving it to the right removes less crackle.



The Crackle Threshold slider allows fine-tuning crackle reduction

Click Meter – The Click meter is another of SoundSoap Pro's visual feedback tools. It illuminates in red to indicate that a click is being repaired.




The Click Meter – When settings detect click & crackle, the meter lights up, as in the picture on the right

3

Global Tools – As with all the other tool tabs, SoundSoap Pro's global tools may be accessed from within the Click & Crackle tab. These global tools include Preset Buttons, A/B/C/D Compare Buttons, Help Button, Noise Only Mode Button, Spectrogram, and ability to monitor other tool settings with the Tool Tabs/Tool State Thumbnails. These tools are described in more detail at the beginning of this section.


Congratulations! You've now learned about SoundSoap Pro's Click & Crackle Tools – the next section discusses the Broadband Tool..

SoundSoap Pro User Interface: Broadband Tab



Broadband Tab

On/Bypass Radio Buttons/Checkbox – These controls allow the Broadband tool to be disabled, or bypassed, while other tools remain active. This control allows the tool to simply be turned off when not required, to save CPU power for other tasks. It is also useful for comparing the difference between having the Broadband tool on or off.




Broadband Tools may be turned on or bypassed using the radio buttons in the tool area, or with a checkbox in the tab

Learn Noise Button – The Learn Noise button is used to automatically “learn” a profile of frequencies responsible for causing a particular type of Broadband noise. There are two modes of operation available with the Learn Noise button. The first mode, called “Snapshot Mode”, (identified by a camera icon) works by depressing the button and initiating audio playback within the host application. SoundSoap Pro analyzes the minimum amount of audio (from the current playhead position) needed to accurately create a noise profile, and then automatically snaps the 12 Threshold sliders to the appropriate positions to reduce/remove the unwanted broadband noise. The second mode is a timed Learn Noise mode. To use this mode, one would click the portion of the Learn Noise button with the clock icon, and then initiate playback. SoundSoap Pro continues to analyze audio (beginning from the original playhead position) that is played until the clock icon is clicked a second time. At this time, the 12 Threshold sliders are snapped into the appropriate position to reduce/remove the Broadband noise present, based on the average noise profile learned over the duration indicated by the user.



Meters – The 12 meters in the Broadband Tool indicate the level of audio contained within each frequency range. The meters on the left-most side of the interface indicate low frequency ranges, while the meters on the right-most side of the interface indicate high frequency ranges – with meters between these two extremes indicating the midrange frequencies between the two. The actual frequency range of each band is not fixed – it depends on the learned noise profile.



The meters indicate audio level in each frequency band, and help in fine-tuning settings for individual bands

Threshold/Reduction View Slider – This horizontal slider allows viewing either just the Threshold sliders – just the Reduction sliders – or can be positioned so that both Threshold and Reduction sliders can be viewed and adjusted at the same time. By default, the Threshold/Reduction Control slider is set to the center position, so that both Reduction and Threshold sliders are visible and adjustable.



The Threshold/Reduction Control slider allows viewing just Threshold sliders, just Reduction sliders, or a percentage of both simultaneously

Reduction Sliders – The Reduction sliders control the amount of noise reduction being applied to a particular frequency range. Reduction sliders can be used in a number of ways – they may all be locked together and moved simultaneously by moving a single slider – or they may be temporarily unlocked, moved to different relative positions, and then locked back together and moved uniformly at their respective positions. In addition, they can all be unlocked and moved independently of each other – or, some sliders may be locked together, while others

remain unlocked and can be moved independently. By default, the Reduction sliders are set to a value of -15dB, and have an overall range of 0dB to -40dB.



Reduction sliders may be moved in groups or independently

Threshold Sliders – The Threshold sliders set the level that the signal must exceed for processing to be enabled in each frequency range. The sliders default to being locked together, so moving one slider will move the other sliders of the same type (reduction or threshold) and retain their relative positions. Individual sliders may be fine-tuned by unlocking them and moving them independently. You can optionally choose to lock any number of sliders together so they can be adjusted in unison. By default, the Threshold sliders are set to a value of 0dB, and have an overall range of 0dB to -70 dB.



Threshold sliders may be moved in groups or independently

Lock/Unlock Reduction Sliders Button – This control either locks or unlocks all the Reduction sliders so that they may be moved independently of one another – or moved as a group, each in its relative position. In addition to the master Lock/Unlock button, each individual slider has its own Lock/Unlock control positioned at its base. These controls allow the locking or unlocking of each frequency band's Reduction slider, without having to unlock all the sliders simultaneously. By default, all the Reduction sliders are locked together.



Above: All Reduction sliders locked together – they move together
Below: All Reduction sliders unlocked – they now move independently

Lock/Unlock Threshold Sliders Button – This control either locks or unlocks all the Threshold sliders so that they may be moved independently of one another – or moved as a group, each in its relative position. In addition to the master Lock/Unlock button, each individual slider has its own Lock/Unlock control positioned at its base. These controls allow the locking or unlocking of each frequency band's Threshold slider, without having to unlock all the sliders simultaneously. By default, all the Threshold sliders are locked together.



Above: All Threshold sliders locked together – they move together
Below: All Threshold sliders unlocked – they now move independently

Attack Knob – The Attack knob sets the attack time in milliseconds for each frequency band. By default, the Attack value is set to 75 milliseconds.



The Attack knob controls how fast broadband noise reduction is turned on

Release Knob – The Release knob sets the release time in milliseconds for each frequency band. By default, the Release value is set to 100 milliseconds.



The Release knob controls how fast broadband noise reduction is turned off

Attack Tilt Knob – The Attack Tilt knob adjusts the ratio of Attack between the low frequency bands and the high frequency bands. By default, the Attack Tilt knob's value is set to 1.0.



The Attack Tilt knob varies the ratio of Attack values between low and high frequency bands

For example, if the Attack Knob's value is set to 50 ms, and the Attack Tilt knob is set to "1.0", then the Attack time for the lowest frequency band will be 50 ms, as will be the Attack time for all other frequency bands. If however, the Attack Tilt knob is set to "2.0", then the Attack time for the lowest frequency band remains at 50 ms, but the Attack time for the highest frequency band will be 25 ms - or, a 2:1 ratio between the lowest and highest frequency bands. In this case, moving to the right from the lowest frequency band to the highest, each successive band will have a proportionally lower attack value. Essentially, the Attack Tilt knob simply varies the ratio of Attack time between the lowest and highest frequency bands, and sets the Attack time of any midrange frequencies to a proportional value in between.



quency band remains at 50 ms, but the Release time for the highest frequency band will be 25 ms - or, a 2:1 ratio between the lowest and highest frequency bands. In this case, moving to the right from the lowest frequency band to the highest, each successive band will have a proportionally lower release value. Essentially, the Release Tilt knob simply varies the ratio of Release time between the lowest and highest frequency bands, and sets the Release time of any midrange frequencies to a proportional value in between.

3

Global Tools - As with all the other tool tabs, SoundSoap Pro's global tools may be accessed from within the Broadband tab. These global tools include Preset Buttons, A/B/C/D Compare Buttons, Help Button, Noise Only Mode Button, Spectrogram, and the ability to monitor other tool settings with the Tool Tabs/Tool State Thumbnails. These tools are described in more detail at the beginning of this section.

Release Tilt Knob - The Release Tilt knob adjusts the ratio of Release between the low frequency bands and the high frequency bands. By default, the Release Tilt knob's value is set to 1.0.

You've now learned about SoundSoap Pro's Broadband Tool - the next section goes over the Noise Gate Tool, which can be used to eliminate any residual noise leftover after using the previously mentioned tools.




The Release Tilt knob varies the ratio of Release values between low and high frequency bands

For example, if the Release Knob's value is set to 50 ms, and the Release Tilt knob is set to "1.0", then the Release time for the lowest frequency band will be 50 ms, as will be the Release time for all other frequency bands. If however, the Release Tilt knob is set to "2.0", then the Release time for the lowest fre-



SoundSoap Pro User Interface: Noise Gate Tab




Noise Gate Tab

On/Bypass Radio Buttons/Checkbox – These controls allow the Noise Gate tool to be disabled, or bypassed, while other tools remain active. This control is useful as it allows the tool to simply be turned off when not required, to save CPU power for other tasks. It is also useful for comparing the difference between having the Noise Gate tool on or off.



Noise Gate Tools may be turned on or bypassed using the radio buttons in the tool area, or with a checkbox in the tab

Visual Display – The Noise Gate's visual display shows the audio waveform being processed, as well as the noise gate's threshold indicators. Options for viewing the waveform include displaying just the left or right channels, or adding the left and right channels together – displaying the overall audio signal between the two channels as a single audio waveform. In addition, there is a Freeze button that appears when the mouse cursor is placed over the visual display. Clicking the button temporarily pauses waveform scrolling, and clicking the button when the waveform is paused starts it scrolling again.



Various elements of the Noise Gate's visual display

Threshold Slider – The Threshold slider determines the audio level required for the Noise Gate to open. The threshold level is also represented by a pair of white, par-

allel horizontal lines, which appear over the audio waveform in the visual display. By default, the Noise Gate Threshold slider is set to a value of 20dB, and has a range of 0dB to -60dB.



The Noise Gate's Threshold can be adjusted with a slider, or by entering a numerical value

Reduction Slider – This slider provides a proportional amount of noise reduction, based on the settings of the Noise Gate's Threshold Slider. A lower setting in the Depth slider provides a lower ratio between the threshold level to the amount of noise reduction, whereas a higher setting increases this ratio. By default, the Noise Gate's Depth slider is set to a value of 2.00, and has a range of 1.00 to 5.00.



The Noise Gate's Depth can be adjusted with a slider, or by entering a numerical value

Threshold Attack Knob – The Threshold Attack knob determines how quickly the Noise Gate opens. By default, the Attack value is set to 30 milliseconds, and has a range of 10 milliseconds to 500 milliseconds.



The Noise Gate's Attack can be adjusted with a knob, or by entering a number

Threshold Release Knob – The Threshold Release knob determines how quickly the Noise Gate closes. By default, the Release value is set to 100 milliseconds, and has a range of 50 milliseconds to 1000 milliseconds.



The Noise Gate's Release can be adjusted with a knob, or by entering a number

3

10. As with everything, misuse is a great thing. Keeping just the clicks and crackle can be extremely fruitful, especially when we use Max/MSP (and especially in the advanced course); the clicks can be used to trigger events. And clicks are good to collect for making beats. The clicks resulting from a click and crackle removal operation are quite varied.

11. There are a few other plug-ins that are part of the Master Perfection Suite from BIAS. Briefly, GateEx is a gate/expander, which can be used as a noise gate as in SoundSoap Pro, and as a dynamic expander—we'll be looking at dynamic things next lab. All these plug-ins have tons of controls for fine-tuning production, less for our messing around. PitchCraft is a pitch transposer that can retain vocal quality, unlike with normal pitchshifting, which e.g. makes us sound like the chipmunks when our voices are transposed up an octave. The manual pages that follow include a section (p. 38) on how to change your voice into that of a robot, child, person of another gender, etc. Repli-Q learns the average spectrum of a recording, say of a Radiohead song, and you can then slap it onto one of your duller produced pieces. It doesn't do anything magical and subtle like impart reverb, spatialization, individual track EQ-ing, but it is a cheap way of making something sound half-good. Reveal doesn't change sound, but shows you all kinds of things about it, many of which won't be useful to you. For instance, the phase diagram is very useful in production, but we don't (yet) have much use for it. The Squeez plug-ins are dynamic compressing, limiting, and upward expanding tools that are supposed to be pretty decent.

Chapter 3: GateEx

Introduction

GateEx is a full-featured noise gate and downward expander plug-in. When used as a noise gate, GateEx allows an audio signal to be reduced in amplitude when it falls below a user-defined threshold level. When used as a downward expander, GateEx allows the apparent dynamic range of an audio signal to be increased.

Some common examples of how a gate would typically be used are:

- In dialogue recordings, a gate might be used to remove (or minimize) the presence of background noise or room tone that is audible in the amplitude dips between the speaker's words or phrases. When the subject speaks, the rising amplitude of the voice triggers the gate to open, allowing the voice to pass through, and as words trail off and the amplitude drops, the gate closes and prevents any low level background noise from being audible.
- In music recordings (we'll use a drum kit as an example), a gate might be used to remove or isolate one type of drum from another. When used properly on a drum recording, a gate can accomplish such tasks as isolating the snare drum, and removing any bleed picked up through the snare mic from other nearby drums, such as the kick or hi-hat.

A common example of how a downward expander would typically be used:

- A downward expander can be applied to increase the apparent dynamic range. This increase in

dynamic range is accomplished by decreasing the gain amount in the relatively quiet portions of the recording, so that the apparent noise floor is lowered.

3

GateEx User Interface & Controls

The GateEx plug-in is designed to be intuitive and easy to use. This section discusses the graphical user interface, and will familiarize you with each of the various controls.


Global Controls

GateEx features a variety of parameters designed for specific tasks, and also features some "global" controls, that affect the state of the entire plug-in. Below is a brief description of these "global" tools, and how they operate.

A/B/C/D Compare Buttons

The A/B comparison feature is common in audio programs, and gives the ability to preview different settings before committing to them. Essentially, various tool settings are made and stored as snapshot "A" –

GateEx Graphical User Interface



another group of settings can then be assigned as snapshot "B" – and then clicking the A or B buttons in the interface instantly toggles between the two. As GateEx offers many fine-tuning options, it features not just two comparison modes – but instead includes A/B/C/D comparisons, for four unique snapshots to compare between, before committing to any file processing. These snapshots are similar to saving and loading presets, but are instantaneous, and do not require selecting presets from a menu to load them. To permanently save one of these snapshots, simply select the desired snapshot, and save it as a custom User Preset.

Factory Preset Pop-Up Menu

This menu displays the current factory preset, if one has been selected. Clicking this menu shows a list of all available factory presets. A factory preset may be selected by choosing it from this pop-up menu.

Factory Preset Selection Arrows

These arrows are used to navigate through factory presets. Clicking the down arrow moves to the next available factory preset, and clicking the up arrow moves to the previous factory preset.

User Preset Save Button

Saves a user preset, which includes the state of all GateEx parameters.

User Preset Load Button

Recalls a user preset, includes the state of all GateEx parameters.



User presets do not appear in the Factory Preset pop-up menu, so be sure you are aware of the directory where you are saving them into, so they may be recalled easily later.

Bypass Button

The Bypass button is used to temporarily disable the plug-in. When the plug-in is bypassed, audio playback is not affected by the plug-in.

Gate/Expander Controls

The Gate/Expander controls each affect a unique characteristic of the overall gate/expander behavior. These are described briefly below.

3

Input Level Meters

The Input Level Meters show the level of the original audio signal being fed into the plug-in, before it is affected by any of the various controls.

Output Level Meters

The Output Level Meters show the level of the processed audio signal, after being affected by the active parameter settings.

Threshold Slider

The Threshold slider determines the amplitude level at which the gate opens or closes. If the signal level exceeds the Threshold setting, the gate opens, and remains open for as long as the signal remains above the Threshold. If the signal falls below the Threshold setting, the gate closes, and remains closed until the signal exceeds the Threshold setting.

By default, the Threshold slider is set to a value of -60dB, and has a range of 0 to -60dB.

Attack Knob

The Attack Knob determines how quickly the gate opens when the audio signal being fed into it exceeds the threshold level.

By default, the Attack Knob is set to a value of 30ms, and has a range of 1-1000ms.

Release Knob

The Release Knob determines how quickly the gate closes when the audio signal being fed into it falls below the threshold level.

By default, the Attack Knob is set to a value of 100ms, and has a range of 50-3000ms.

Depth Knob

The Depth Knob determines how far the gate closes (i.e., how deeply the signal is attenuated) when triggered by a signal that exceeds the threshold level. By default, the Depth control is set to its maximum level, -90dB, which causes audio signals to be completely silenced when the gate is in the closed position. The Depth control is useful as it allows a variable level of signal reduction when the gate is closed, and can help produce a more subtle gating effect. Another useful effect of having a Depth setting lower than the maximum amount is that it allows the gate to open faster than if it were in the fully closed position, and can be useful if a faster attack time is needed, and the Attack knob is already set to operate at its lowest value (1ms).

By default, the Attack Knob is set to a value of -90dB, and has a range of 0 to -90dB.

Hold Knob

The Hold Knob allows the gate to remain in the open position for a user-specified amount of time, even after the signal level has fallen below the threshold level. This control is especially useful when the audio signal is very

close to the threshold level, and might cause the gate to open and close repeatedly, very quickly, as the signal exceeds and then falls below the threshold over a very short period of time. This rapid opening and closing of the gate can create an unwanted artifact called "chatter" or "jitter", and the Hold control helps to avoid it.

By default, the Hold Knob is set to 0ms, and has a range of 0 – 3000ms.

Hysteresis Knob

The Hysteresis control is similar to the Hold control in its practical application of helping to avoid/eliminate "chatter", but achieves this in a different way. Rather than simply maintaining the gate's open position for a user-specified duration, the Hysteresis control allows setting independent threshold levels for opening or closing the gate, by specifying the difference between them in decibels. With this approach, the gate opens when a primary threshold level is exceeded by a suitable signal level, and closes again only when the signal level drops below a secondary threshold level.

By default, the Hysteresis Knob is set to 0dB, and has a range of 0 – 10dB.

Lookahead Knob

As its name implies, the Lookahead Knob reads data from a file a bit earlier than it is actually played back. This allows the gate to react earlier than it would, based on its threshold setting alone. This feature is useful as it helps to preserve transients and attacks that might be eliminated (or minimized) due to the gate opening after they had already occurred.

By default, the Lookahead Knob is set to 0ms, and has a range of 0 – 100ms.

Ratio Knob

The Ratio Knob is used to set the rate at which audio signal reduction is applied, when the signal exceeds the threshold level.

By default, the Ratio Knob is set at 1:1, and has a range of 1:1 – 50:1.

Gain Slider

The Gain slider is used to set the output level. By default, it is set to 0dB, and has a range of 0 to -90dB.

particular areas of the waveform display as you scroll through its timeline.

Waveform View Radio Buttons

This group of radio buttons allows choosing between having the Waveform Display calculated by the content of a stereo signal's left channel (L), right channel (R), or both channels (L+R).



When using L+R mode, a single waveform displays the sum of both channels.

Graphic Displays

GateEx features two graphic display areas that are helpful in making the appropriate settings quickly. These are described below.

Waveform Display

The Waveform Display shows a graphical representation of the audio signal being fed into GateEx in real-time. This display is particularly helpful in quickly setting the threshold level, as it creates a graphical relationship between the signal level and the threshold level.

Gate Threshold Indicators

These indicators appear as horizontal white lines superimposed over the Waveform Display, and give a graphical representation of the current threshold value, in relation to the input signal level.

Freeze Waveform Button

The Freeze Waveform Button allows you to stop and start the waveform from scrolling. This is useful for examining

Zoom Button

This control is located where the X & Y-axes meet. It can be used to zoom in/out on either axis independently, or both at the same time. This allows for higher precision viewing of the graph display.

Horizontal Scrollbars/Arrows

When zoomed in on the horizontal axis, the scrollbars and/or scroll arrows may be used to move the Waveform Display earlier or later in time, allowing examination of waveform events anywhere along the timeline.

By default, the zoom level is set to “zoomed out all the way”, and the horizontal scrollbars/arrows have no effect.

Vertical Scrollbars/Arrows

When zoomed in on the vertical axis, the scrollbars and/or scroll arrows may be used to move the Waveform Display up or down, allowing examination of events that occur in the waveform's positive or negative phase.

By default, the zoom level is set to “zoomed out all the way”, and the vertical scrollbars/arrows have no effect.

3

Graph Display

The Graph Display shows a number of key parameter settings in an easy-to-read graph format, with each axis labeled in decibels. The X-axis represents audio input level, while the Y-axis represents output level. The graph makes it fast and easy to make settings, and allows easy plotting of input and output level values.



When Hysteresis is used, the Attack threshold setting in the graph display is represented in blue, and the Release threshold setting appears in green.

Using GateEx

This section contains some basic “hands-on” getting started tips.

1. Use the Ratio Knob to set the rate at which reduction increases when the audio signal exceeds the current threshold level. A good starting point is to set it to the highest setting (i.e., 50:1).



2. Adjust the Threshold slider until the gate opens to allow the desired audio content to pass through it, and closes to prevent lower level, unwanted audio from being passed.



3. Fine-tune your settings by making adjustments to the following controls, and try these tips:

- **Attack and Release** – Use faster times for percussive material (such as drums), or slower times for less-percussive material (such as voice).
- **Depth** – Set the depth so the proper amount of reduction is happening when the gate is closed. This allows for subtle gating where not all the signal is removed when the gate is closed, or more drastic gating where all the signal is removed when the gate is closed.
- **Hold** – Use the hold control to ensure the gate stays open for the desired length of time (i.e., so that it does not begin to close prematurely and reduce part of the desired audio signal).
- **Hysteresis** – Use the hysteresis control to set independent threshold levels for opening/closing the gate if the audio signal hovers near the threshold setting.
- **Lookahead** – Use the Lookahead control if you need the gate to open a little before the audio signal goes above the threshold level in order to preserve any transients or attacks.

If you are unfamiliar with using gates, try using some of the included factory presets, which are representative of common gate settings. These should provide a good starting point, from which you can fine-tune the controls to best suit the audio material you are working with.



Conclusion

Now that you are familiar with the various GateEx controls, and its basic workflow, please continue along to the next chapter, where you will learn about the PitchCraft plug-in.

Chapter 4: PitchCraft

Introduction

PitchCraft is a pitch correction and transposition plug-in that can be used to correct and/or change pitch from a variety of sources. The high quality algorithm includes formant correction to retain the natural sound of a voice – and may also be used to intentionally modify the characteristics of a voice.

This plug-in provides high-quality professional pitch correction/transposition, with an easy to use, intuitive interface.

Some common examples of how PitchCraft might be used include:

- To correct pitch problems in vocal or instrumental recordings
- Change pitch of a performance to match a different key than was used in the original performance
- Pitch and formant modification, as a sound design tool
- Virtually altering the size of the vocal cavity and vocal cords, for effects such as “gender transformation”, voice disguising, changing an adult’s voice into a child’s voice and vice-versa

PitchCraft User Interface & Controls

The PitchCraft plug-in is designed to be intuitive and easy to use. This section discusses the graphical user interface, and will familiarize you with each of the various

controls. The user interface is divided into a number of sections, each with dedicated controls or displays used for pitch transposition and correction.

Global Controls

PitchCraft features a variety of parameters designed for specific tasks, and also features some “global” controls, that affect the state of the entire plug-in. Below is a brief description of these “global” tools, and how they operate.

Reference Source Pop-Up Menu

This menu allows you to choose whether the reference source (i.e., the signal being fed into PitchCraft) is a single (left or right) channel of a stereo file, or both channels simultaneously.




If you are working with a mono file, be sure to set the Reference Source Pop-Up menu to “Left”.

Pitch Range Pop-Up Menu

This menu allows you to choose the desired pitch range to use. Using the appropriate pitch range optimizes PitchCraft for certain source types. The plug-in will not detect pitch values outside the bounds specified for a given source. If you hear artifacts in the audio output,

PitchCraft Graphical User Interface



you might want to try a different setting, as artifacts are generally produced when the source pitch falls outside the bounds specified by the value selected in the Pitch Range pop-up menu.

Making a pitch range selection in the Pitch Range pop-up menu assigns a corresponding pitch range scale along the Y-axis of the Tuning History Graph. Pitch range options include:

- Normal – D2 (73 Hz) to E6 (1318 Hz)
- Female – E3 (165 Hz) to E6 (1318 Hz)
- Male – D2 (73 Hz) to B4 (494 Hz)
- Lower Register – E0 (20 Hz) to E4 (330 Hz)
- Higher Register – E4 (330 Hz) to D#7 (2500 Hz)

A/B/C/D Compare Buttons

The A/B comparison feature is common in audio programs, and gives the ability to preview different settings before committing to them. Essentially, various tool settings are made and stored as snapshot "A" – another group of settings can then be assigned as snapshot "B" – and then clicking the A or B buttons in the interface instantly toggles between the two. As PitchCraft offers many fine-tuning options, it features not just two comparison modes – but instead includes A/B/C/D comparisons, for four unique snapshots to compare between, before committing to any file processing. These snapshots are similar to saving and loading presets, but are instantaneous, and do not require selecting presets from a menu to load them. To permanently save one of these snapshots, simply select the desired snapshot, and save it as a custom User Preset.

Factory Preset Pop-Up Menu

This menu displays the current factory preset, if one has been selected. Clicking this menu shows a list of all available factory presets. A factory preset may be selected by choosing it from this pop-up menu.

Factory Preset Selection Arrows

These arrows are used to navigate through factory presets. Clicking the down arrow moves to the next available factory preset, and clicking the up arrow moves to the previous factory preset.

User Preset Save Button

Saves a user preset, which includes the state of all PitchCraft parameters.

User Preset Load Button

Recalls a user preset, includes the state of all PitchCraft parameters.



User presets do not appear in the Factory Preset pop-up menu, so be sure you are aware of the directory where you are saving them into, so they may be recalled easily later.

4

Bypass Button

The Bypass button is used to temporarily disable the plug-in. When the plug-in is bypassed, audio playback is not affected by the plug-in.

Tuning History Graph Display and Tuning Meter

These controls are used to view an audio signal's tuning history over time, and also to customize and navigate the Graph display.

View Notes/Hz Radio Buttons

These controls allow you to choose whether the Graph display's Y-axis is represented as Notes or Hz (frequencies).

Graph Display

This display plots the pitch of both the input (raw) and output (processed) signals, and is useful for comparing the difference between the two. The yellow line represents the input pitch, and the orange line represents the output, or corrected pitch. The X-axis represents time, and the Y-axis represents the currently selected scale, and may be viewed as notes or Hz (frequencies).

Zoom Button

Zooming in or out is accomplished using the Zoom Button, located at the point where the X and Y-axes meet. This control allows a very detailed zoomed-in view, or can be zoomed-out for a "big picture" view. Option-clicking on the Zoom Button resets the zoom level.

- Zoom In Horizontally – Click and drag to the right
- Zoom Out Horizontally – Click and drag to the left
- Zoom In Vertically – Click and drag up
- Zoom Out Vertically – Click and drag down

Tuning Meter

The Tuning meter appears vertically along the right side of the interface, and displays the current amount of tuning. The divisions along the Tuning meter's scale is marked at 200-cent intervals, with the maximum tuning amount of 200 cents at the top of the scale, zero cents in the middle, and -200 cents at the bottom.

Transposition Controls

This group of controls is used for altering the pitch and formant of the reference source being fed into PitchCraft.

Transpose On/Off Checkbox

This control allows enabling/disabling all Transposition controls.

Pitch Slider/Numerical Field

This control allows raising or lowering the pitch value (in cents), while maintaining the original formant structure (i.e., the "timbre" or "tonal color" of the voice).



When working with vocal sources this control is the equivalent to modifying the length of the vocal cords.

Each division along the Pitch Slider's scale represents the associated musical interval (major second, major third, etc).



Sliders/pop-up sliders can be reset to their default values by Option-clicking on them.



When numerical fields are in an editable state (i.e., after clicking in them with the mouse), their values may be changed by typing in new values, by using the Up/Down Arrow keys on your keyboard, or by using the scroll wheel on a Mac OS X-compatible scroll wheel mouse.

Formant Slider/Numerical Field

This control allows raising or lowering the format value

(in cents) through a tilt of the spectral envelope, stretching the spectrum at the high end, and compressing it at the low end. The formant slider acts as a particular type of filter, while the pitch control performs the usual pitch transposition on the instrument.



When working with vocal sources this control is the equivalent to modifying the volume of the vocal cavity.

Each division along the Formant Slider's scale represents the associated musical interval (major second, major third, etc). When the Formant slider is set to zero, formants are preserved independently of pitch transposition (known as "formant-corrected pitch transposition").

Pitch Correction Controls

This group of controls is used for correcting problems with the pitch of the reference source being fed into PitchCraft.

Correct Pitch On/Off checkbox

This control allows enabling/disabling all Pitch Correction controls.

Tuning Pop-Up Menu

This menu allows you to choose the desired type of tuning to be used in the pitch correction process. Tuning is expressed in number of cents with respect to the equal-tempered note, for each of the twelve notes, and includes the following options:

- Equal – 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0

- Classic Just – 0, -29, 4, 16, -14, -2, -10, 2, 14, -16, 18, -12
- Harmonic – 0, 5, 4, -2, -14, -29, -49, 2, 41, 6, -31, -12
- Indian Raga – 0, -10, 4, -6, -16, -2, -12, 2, -8, -18, -4, -14
- Arabic – 0, 30, -20, -50, -45, 2, 23, 6, -14, -43, -70, 10
- Bali/Java – 0, 27, -58, -42, -123, -130, -126, 211, 24, -61, -57, -142
- Tibetan – 0, -42, 32, 10, -22, 22, 18, 25, -27, -4, 19, -14
- Carols Alpha – 0, -12, -44, 66, 88, -110, -132, -164, -176, -198, 220, 242
- Pythagorean – 0, -10, 4, -6, 8, -2, 12, 2, -8, 6, -4, 10
- Mean – 0, -24, -7, 10, -14, 3, -21, -3, -27, -10, 7, -17

Scale Pop-Up Menu

This menu allows you to choose the desired scale to be used in the pitch correction process. Scales are expressed in number of semitones with respect to the root note

- Chromatic – All Notes
- Major – 1, 2, 3, 4, 5, 6, 7
- Dorian – 2, 3, 4, 5, 6, 7, 1
- Phrygian – 3, 4, 5, 6, 7, 1, 2
- Lydian – 4, 5, 6, 7, 1, 2, 3
- Mixolydian – 5, 6, 7, 1, 2, 3, 4
- Minor – 6, 7, 1, 2, 3, 4, 5
- Locrian – 7, 1, 2, 3, 4, 5, 6
- Whole Tone – 2, 4, 6, 8, 10
- Blues – 3, 5, 6, 7, 10
- Harmonic Minor – 2, 3, 5, 7, 8, 11
- Melodic Minor – 2, 3, 5, 7, 9, 11
- Major Pentatonic – 2, 4, 7, 9
- Minor Pentatonic – 3, 5, 7, 10

4

- Major Bebop – 2, 4, 5, 7, 9, 10, 11
- Dominant Bebop – 2, 4, 5, 7, 9, 10, 11
- Super Locrian – 1, 3, 4, 6, 8, 10
- Diminished – 2, 3, 5, 6, 8, 9, 11
- Hungarian Minor – 2, 3, 6, 7, 8, 11
- Arabian – 2, 4, 5, 6, 8, 10

Key Pop-Up Menu

This menu allows you to choose the root note of the scale selected in the Scale pop-up menu. Options include:

- A
- A#
- B
- C
- C#
- D
- D#
- E
- F
- F#
- G
- G#

Pitch Correction Time Slider

This control allows you to adjust the speed at which pitch correction is applied.

Detune Slider

This control allows you to adjust the reference frequency used for tuning (i.e., from 440Hz to 445Hz). This control should be adjusted only if attempting to match the pitch with another source that is not in "concert pitch" (i.e., not tuned to A 440).

Keyboard Controls

The Keyboard controls are used for defining custom notes and scales to be used in the pitch correction process. These controls are described below.

Ignore Note Checkboxes

These checkboxes allow certain notes to be ignored in the pitch correction process, even when they are present in the scale being used. The presence of a check in an Ignore checkbox means that the note represented by the key that appears directly below it in the interface will be ignored. The absence of a check means the note will be used in the pitch correction process.

Keyboard Display

The virtual keyboard is used to display the active notes being used in the pitch correction process, and can also be used to define a custom scale for pitch correction.

Use Note in Scale Indicators

The Use Note in Scale indicators appear as orange dots superimposed over keys in the Keyboard display. Scales can be created or modified using any MIDI device, such as Peak's MIDI keyboard window, or a MIDI controller/keyboard. Any MIDI "Note On" message from a controller or sequence is added to the current scale, if the note is already part of the current scale, the "Note

On” message will remove it from the current scale. When a Use Note in Scale indicator appears over a key, that key’s note is detected and its pitch is corrected. When an indicator does not appear over a key, that note is not part of the current scale.

Custom Scale Controls

These controls allow the creation of a custom scale to be used for pitch correction. Custom scales are automatically saved into the Scale pop-up menu, and may be recalled for use later.

Current Scale Display

Displays the name of the current scale, as well as appropriately placed Use Note in Scale indicators over the keys/notes in the active scale.

Using PitchCraft

This section includes hands-on directions for accomplishing common pitch modification tasks.

To Define a Scale:

- Choose the desired Scale Preset from the Scale pop-up menu.

To Define a Custom Scale:

- 1a. Click the desired keys in PitchCraft’s virtual keyboard – clicking a key toggles the Use Note in Scale indicators into the On or Off position.

Or:

- 1b. Click the keys of a MIDI input device (such as

Peak’s virtual MIDI keyboard) – the MIDI Note On messages toggle the Use Note in Scale indicators into the On or Off position.

2. When you have finished defining your custom scale, click the Done button – the custom scale is added to the Scale menu as a preset.

To Correct Pitch:

1. Select the appropriate type of reference source from the Reference Source pop-up menu.
2. Select the appropriate pitch range for the reference source, using the Pitch Range pop-up menu.
3. Enable the Correct Pitch section, by checking the Correct Pitch checkbox.
4. Define a scale or set of notes using the Scale, Key, and Keyboard controls.
5. Play the audio source. As audio enters the plug-in and is analyzed, the current pitch is displayed in the Tuning History Graph. The plug-in determines which note in the current scale is “correct” and applies the appropriate amount of pitch correction to the input source, to match the note in the scale.

The Pitch Correction Time control can be used to smooth out the correction effect if the pitch becomes “warby”. It can also be used to create interesting “robotic” voice effects by setting the pitch correction time to zero.



The Transpose and Correct Pitch sections may be enabled simultaneously, with transposition and formant modification being performed in parallel with pitch correction. Advanced features include the ability to detune the reference frequency (for example from A = 440Hz to 443Hz), the ability to define the scale via a MIDI controller, and the ability to save up to 4 temporary presets.



PitchCraft is designed for monophonic sources and it will not work on polyphonic material. Please keep in mind that material with excessive reverb or background noise is not considered monophonic. It is advised that you clean up the audio signal as much as possible before processing it with this algorithm. PitchCraft may produce interesting creative effects when used with polyphonic audio material, but keep in mind this is an algorithm carefully designed for transparent processing of monophonic sounds.

To Transpose Pitch and/or Formant:

1. Select the appropriate type of reference source from the Reference Source pop-up menu.
2. Select the appropriate pitch range for the reference source, using the Pitch Range pop-up menu.
3. Enable the Transpose section of the interface, by checking the Transpose checkbox.
4. Adjust the Pitch and/or Formant sliders until the desired effect is achieved.

Voice Modification Using the Pitch & Formant Controls

The Pitch and Formant controls can produce very interesting special effects, especially when the source is a recording of a human voice. When using these controls for modifying a human voice, the following guidelines may be useful:

- Transposing pitch upward is the equivalent of shortening vocal cords
- Transposing pitch downward is the equivalent of lengthening vocal cords
- Shifting formants upward is the equivalent of reducing the size of the vocal cavity
- Shifting formants downward is the equivalent of enlarging the size of the vocal cavity

Some examples of voice modification possibilities are listed below:

To Transform a Female Voice into a Male Voice:

1. Using the Pitch slider, transpose pitch downward substantially.
2. Move Formant slider down until a natural sound is achieved, (which usually occurs when the Formant slider's value is slightly smaller than the Pitch slider's value).

To Transform a Male Voice into a Female Voice:

1. Using the Pitch slider, transpose pitch upward.
2. Move the Formant slider up until a natural sound is achieved.

To Transform an Adult's Voice into a Child's Voice:

- Raise the Formant slider's value – notice how "age" decreases as the formant value increases.

To Transform a Child's Voice into an Adult's Voice:

- Lower the Formant slider's value – notice how "age" increases as the formant value decreases.

To Transform a Human Voice into an "Alien" Voice:

1. Raise the Pitch value all the way (i.e., slider all the way to the right).
2. Lower the Formant value all the way (i.e., slider all the way to the left).

To Disguise a Human Voice:

1. Lower the Formant value slightly.
 2. Adjust Pitch value to suit level of voice disguise.
-

Chapter 5: Repli-Q

Introduction

Repli-Q is a powerful EQ matching/spectral balancing plug-in, which is designed to analyze, edit, and compare the spectral content of audio recordings. Once analysis has taken place, spectral characteristics of one recording may be applied to another. By viewing the frequency spectrum plotted as graphical data (and by listening of course), a user can discern where there are peaks, “holes”, or other problems in the spectral content of the recording. Various steps in the recording and mastering process can result in problems such as studio resonance caused by the natural comb filtering effect of a poorly shaped recording/listening environment – or equalization used to compensate for a studio’s equipment and environmental deficiencies. Repli-Q allows these sorts of spectral problems to be located and repaired using its high quality linear-phase filters.

Some common examples of how Repli-Q might be used are:

- Mimicking and applying a favorite artist’s spectral (EQ) “style” to your own recordings.
- In the mastering process, it can give a cohesive sound to a collection of songs that may have been recorded and mixed in different locations/environments.
- It can be used to create a collection of profiles – microphones, for example. These profiles may later be applied to audio recordings to give them the characteristic sound of a particular type of microphone.

Repli-Q matching profiles may be created for other types of equipment, such as consoles, compressors, recording environments, etc.

Repli-Q User Interface & Controls

Repli-Q is designed to be intuitive and easy to use. This section discusses the graphical user interface, and will familiarize you with each of the various controls.

5


Global Controls

Repli-Q features a variety of controls designed for specific tasks, and also features a number of “global” controls, that affect the state of the entire plug-in. Below is a brief description of these “global” tools, and how they operate.

A/B/C/D Compare Buttons

The A/B comparison feature is common in audio programs, and gives the ability to preview different settings before committing to them. Essentially, various tool settings are made and stored as snapshot “A” –

Repli-Q Graphical User Interface



another group of settings can then be assigned as snapshot "B" – and then clicking the A or B buttons in the interface instantly toggles between the two. As Repli-Q offers many fine-tuning options, it features not just two comparison modes – but instead includes A/B/C/D comparisons, for four unique snapshots to compare between, before committing to any file processing. These snapshots are similar to saving and loading presets, but are instantaneous, and do not require selecting presets from a menu to load them. To permanently save one of these snapshots, simply select the desired snapshot, and save it as a custom User Preset.

Factory Preset Pop-Up Menu

This menu displays the current factory preset, if one has been selected. Clicking this menu shows a list of all available factory presets. A factory preset may be selected by choosing it from this pop-up menu.

Factory Preset Selection Arrows

These arrows are used to navigate through factory presets. Clicking the down arrow moves to the next available factory preset, and clicking the up arrow moves to the previous factory preset.

User Preset Save Button

Saves a user preset, which includes the state of all Repli-Q parameters.

User Preset Load Button

Recalls a user preset, which includes the state of all Repli-Q parameters.



User presets do not appear in the Factory Preset pop-up menu, so be sure you are aware of the directory where you are saving them into, so they may be recalled easily later.

Bypass Button

The Bypass button is used to temporarily disable the plug-in. When the plug-in is bypassed, audio is not affected.

Profile Controls

The Profile controls are used for learning, saving, and loading spectral and EQ profiles, as well as for matching and smoothing between different profiles. This section discusses each control's function.

Learn Spectrum Button

This control analyzes the spectral content of the audio source being fed into Repli-Q. When the Learn Spectrum button is clicked, Repli-Q immediately begins analyzing, or "learning" the input signal.

When audio playback is stopped, Peak is still sending digital silence (a zero amplitude signal for all samples) to the plug-in. If the Learn Spectrum button has been clicked, Repli-Q is learning the spectrum of digital silence – which is a flat spectrum of very low amplitude (-96 dB). If playback starts, Repli-Q continues averaging the actual signal spectrum with the already averaged flat -96 dB spectrum. The resulting final spectral average will be false.



The recommended way to operate Repli-Q is:

- Start playback first, and then click the Learn Spectrum button, in that order. Then, click the Learn Spectrum button again to stop learning.
- In other words, once the Learn Spectrum button

5



is clicked, start playback immediately, and when playback is stopped (or ends on its own), click the Learn Spectrum button again to stop the learning process.



When numerical fields are in an editable state (i.e., after clicking in them with the mouse), their values may be changed by typing in new values, by using the Up/Down Arrow keys on your keyboard, or by using the scroll wheel on a Mac OS X-compatible scroll wheel mouse.

Save Spectrum Button

This control saves the current spectrum as a Repli-Q Spectral Profile, and automatically places it in the Reference pop-up menu as the current/active profile.

Reference Pop-Up Menu

This pop-up menu contains all Repli-Q Spectral Profiles created during the current session (i.e., since the plug-in was instantiated) – and also allows loading Spectral Profiles saved in earlier sessions.

Save EQ Button

This button saves the current EQ settings as a Repli-Q EQ Profile, and automatically loads it in the EQ pop-up menu as the current/active profile.

EQ Pop-Up Menu

This pop-up menu contains all Repli-Q EQ Profiles created during the current session (i.e., since the plug-in was instantiated) – and also allows loading EQ Profiles saved in earlier sessions.

Matching Slider/Numerical Field

This control allows variable spectral matching between an input source's profile and a reference file's profile.

By default, the Matching slider is set to 50%.



Sliders can be reset to their default values by Option-clicking on them.

Smoothing Slider/Numerical Field

This control allows variable EQ smoothing, and helps to prevent resonance artifacts.

By default, the Smoothing slider is set to 0%.

Gain Slider/Numerical Field

This control increases or decreases the output level.

By default, the Gain slider is set to 0dB, and has a range of -20 to 20 dB.

Gain Compensation Checkbox

When active, this control compensates for gain differences introduced by an EQ setting, and makes the output level identical to the input level. The presence of a check in this box means that Gain Compensation is active.

Soft Clip Checkbox

When active, this control acts as a soft limiter to ensure that the output level does not clip due to gain differences introduced by an EQ setting. The presence of a check in this box means that Soft Clip is active.

Spectrum/EQ View Radio Buttons


These radio buttons allow choosing between viewing of spectral profiles and EQ profiles.

Graph Display

The Graph Display operates in two modes, Spectral Profile Mode and EQ Profile Mode. The mode depends on which of the Spectrum/EQ View radio buttons is selected. Various elements of the Graph Display Area change depending on the mode being used.

Spectrum Profile View Mode Display

In Spectrum View Mode, spectral data is shown in the graph. The X-axis is labeled in Hz (Frequencies), and the Y-axis is labeled in dB (decibels). Three different colored lines are plotted in the Spectrum graph, each represents a unique spectral profile.



Yellow – Represents the spectral profile of the Reference file, which has the desirable characteristics you are attempting to match, and eventually apply to another file (i.e., which is referred to here as the "Input" file).


Green – Represents the spectral profile of the Input file, which is the file onto which you are attempting to apply the Reference's desirable spectral characteristics.

Blue – Represents the average spectrum of the Output profile (i.e., the actual degree of matching between the Input and the Reference). The position of the Blue line within the graph is linked to the position of the Matching slider, and indicates the current output spectrum. When the Matching slider is set to 0%, it matches the Input's (Green line) spectral characteristics perfectly. When set to 100%, it matches the Reference source's (Yellow line) spectral characteristics perfectly. When the Matching slider is set somewhere between 0 –

100%, it represents a blend of both the Reference and Input's spectral characteristics.

EQ View Mode Display

In EQ View Mode, frequency data is shown in the graph. The X-axis is labeled in Hz (Frequencies), and the Y-axis is labeled in dB (decibels). The frequency data shown in the graph depends on the position of the Matching slider. When the Matching slider is set to 0%, the frequency data displayed belongs to the spectral "Input". When it is set to 100%, the frequency data shown belongs to the spectral Reference. When the Matching slider is set to some value between 0 – 100%, a blend of the frequency content of both the Input and Reference is displayed.



Editing EQ Filtering Properties

In EQ View Mode, clicking on the graph display allows editing of the EQ's filtering properties.

- Click in EQ Graph – Displays the current EQ profile's filtering properties, such as frequency, gain, and Q (bandwidth).
- Click & Drag in EQ Graph – This allows editing the EQ profile's frequency and gain filtering properties.
- Shift-Click & Drag in EQ Graph – This allows editing the EQ profile's frequency and Q (bandwidth) filtering properties.

Zoom Button

Zooming in or out is accomplished using the Zoom Button, located at the point where the X and Y-axes meet. This control allows a very detailed zoomed-in view, or can

5

be zoomed-out for a “big picture” view. Option-clicking on the Zoom Button resets the zoom level.

- Zoom In Horizontally – Click and drag to the right
- Zoom Out Horizontally – Click and drag to the left
- Zoom In Vertically – Click and drag up
- Zoom Out Vertically – Click and drag down

Using Repli-Q

This section contains some basic “hands-on” getting started tips, and introduces some typical Repli-Q workflows.

To Match the Spectrum of an Audio File:

1. Open a Reference file in Peak and select the portion to be used as a Reference – the “Reference” file being the one with the desirable spectral characteristics that you wish to apply to another file.



Note that different sections of a Reference file may have differing spectral characteristics. You may wish to denote specific portions of the reference file using Peak's Markers or Regions.

2. Instantiate Repli-Q, and start playback of the reference file in Peak.
3. Click the Learn Spectrum button – the spectral profile is plotted as Repli-Q learns it.
4. Click the Learn Spectrum button again to end the spectral “learning” process.
5. Click the Save Spectrum button to save the learned spectral profile – it will automatically be loaded into the Reference pop-up menu as the

current/active spectral reference profile. This profile is plotted as a yellow line in the Graph display.

6. Open another audio file in Peak – this will be the material onto which we apply the Reference audio file's desirable spectral characteristics.
7. Click the Learn Spectrum button and start playback, the Input file's spectral profile is plotted as a green line in the Graph display.
8. Click the Learn Spectrum button again to end the learning process. As soon as the second learning process is stopped, the Output spectrum's profile is immediately calculated (appears as a blue line in the graph display). At this point, the Graph display will show three differently colored plotted lines – yellow is the Reference profile, green is the Input profile, and blue is the output profile (whose corresponding EQ characteristics will be applied to the Input file when the plug-in is active). Each of the spectral profiles' corresponding EQ profiles may be viewed by switching to the EQ Graph Display (by using the EQ View radio button above the Graph Display).
9. In the Spectral Profile Graph, use the Matching slider to adjust the degree of matching between the Reference and Input profiles – when the Matching slider is at 0%, the Output profile (blue line) coincides with the Input profile (green line), which means that the Input will not be modified. When the slider is at 100%, the Output profile is identical to the Reference profile (yellow line), applying maximum equalization matching to the Input.
10. Use the Smoothing slider to smooth the Output profile (blue line) if it presents too many sharp peaks and dips, which can cause undesirable resonance, or inconsistent performance over different sets of reference monitors.

For additional control, the EQ profile's properties may be manually edited – described in additional steps below.

- In the EQ Graph Display, click and drag to modify EQ gain properties (and/or shift-click and drag to modify the EQ's bandwidth filtering properties).
- Repli-Q is now ready for processing – play back the input and the new EQ settings will be applied.



At any point after learning the Input and matching/smoothing/editing the EQ, that EQ may be saved with the Save EQ button. This is useful when the same EQ is needed in the future. Saving the EQ is not required for processing the current input, however, if not saved, the current EQ will be lost when learning a different spectral profile, or de-instantiating the plug-in.

To Process Using a Saved EQ Profile:

- Instantiate Repli-Q and load a saved EQ profile using the EQ pop-up menu.
- Match, smooth and edit the EQ as desired – matching will simply scale the EQ between a flat 0dB curve and the loaded EQ profile.
- Open an input file in Peak and play back – the loaded EQ profile is applied.

To Use Repli-Q as an Editable Linear Phase EQ:

- Instantiate Repli-Q and select the EQ Graph Display – a flat 0dB curve is displayed.
- Play back audio material that needs EQ changes, and perform as many manual edits to the EQ curve as desired until the desired effect is achieved.
- Bounce the plug-in – the custom EQ settings are applied.

To Automatically Compensate for Gain Changes:

- Check the Gain Comp checkbox – gain changes introduced by EQ settings are compensated for automatically, keeping the output level

approximately equal to the input level. This is useful in evaluating actual differences in EQ characteristics independently of differences in loudness – as increases in loudness sometimes translate to “better sounding”.

To Prevent Clipping the Output Signal:

- Check the Soft Clip checkbox – potential clipping caused by EQ settings that boost output amplitude is now avoided.

To Manually Control Output Level:

- Uncheck the Gain Comp checkbox.
- Use the Gain slider to adjust the output to the desired level.

Conclusion

Now that you are familiar with the various Repli-Q controls, and its basic workflow, please continue along to the next chapter, where you will learn about the Reveal plug-in.



Chapter 6: Reveal

Introduction

Reveal offers an integrated solution to spectral, power, and phase analysis. Its components can be broken down into three categories:

- The tools may also be used for analyzing real-time live input, in order to analyze room acoustics, and potential phase issues due to speaker placement and speaker tuning.

Frequency Analysis Tools

- Oscilloscope
- Spectral Analyzer (Graph & Meter Displays)
- Spectrogram

It takes a combination of audio analysis tools to get a comprehensive understanding of audio content, and Reveal provides them all, in a single easy-to-use plug-in.

Level Analysis Tools

- Peak and RMS Level History Display
- Peak and RMS Level Meters

Reveal User Interface & Controls

Reveal is designed to be intuitive and easy-to-use. This section discusses the graphical user interface, and will familiarize you with each of the various controls.

Phase & Pan Analysis Tools

- Lissajous Phase Scope
- Pan Power History Display

This combination of tools provides audio engineers with all the elements needed to analyze recordings (or live input), and is indispensable during the audio production process. Audio engineers typically use these types of analysis tools in the following scenarios:


- In the studio, they are used to determine problem areas in a mix, or to observe how processing or mixing is affecting the audio content.

Global Controls

Reveal includes many controls designed for specific tasks, and also features several "global" controls, that affect the state of the entire plug-in. Below is a brief description of these "global" tools, and how they operate.

6

Reveal Graphical User Interface



Resolution Pop-Up Menu

This control sets the graphical resolution used in Reveal's various displays. The resolution may be adjusted to suit the CPU of the computer on which Reveal is being used. Higher resolution settings are intended for faster CPUs, and lower settings are intended for slower CPUs – you may want to experiment with this setting to get the best results for your particular computer system.

Global Freeze Button

Clicking the Global Freeze button temporarily stops movement in all displays, allowing for closer examination of their data. Clicking this button again resumes activity in all displays.

Input Source Pop-Up Menu

This menu allows you to choose whether the input source (i.e., the signal being fed into Reveal for analysis) is a single (left or right) channel of a stereo file, or both channels simultaneously.



If you are working with a mono file, be sure to set the Input Source pop-up menu to "Left".

Zoom Buttons

Technically, these are not "global" controls, but they appear in several of Reveal's tools, so they will be covered here.



The Zoom Button appears in several of Reveal's tools, and allows zooming along the X & Y-axes.

Zooming and scrolling is possible within tools that feature this control, which allows a very detailed zoomed-in view, or can be zoomed-out for a "big picture"

view. Zooming is accomplished using the Zoom Button, located at the point where the X and Y-axes meet. Option-clicking on the Zoom Button resets the zoom level.

- Zoom In Horizontally – Click and drag to the right
- Zoom Out Horizontally – Click and drag to the left
- Zoom In Vertically – Click and drag up
- Zoom Out Vertically – Click and drag down

A/B/C/D Compare Buttons

The A/B comparison feature is common in audio programs, and gives the ability to preview different settings before committing to them. Essentially, various tool settings are made and stored as snapshot "A" – another group of settings can then be assigned as snapshot "B" – and then clicking the A or B buttons in the interface instantly toggles between the two. As Reveal offers many fine-tuning options, it features not just two comparison modes – but instead includes A/B/C/D comparisons, for four unique snapshots to compare between, before committing to any file processing. These snapshots are similar to saving and loading presets, but are instantaneous, and do not require selecting presets from a menu to load them. To permanently save one of these snapshots, simply select the desired snapshot, and save it as a custom User Preset.

Factory Preset Pop-Up Menu

This menu displays the current factory preset, if one has been selected. Clicking this menu shows a list of all available factory presets. A factory preset may be selected by choosing it from this pop-up menu.

Factory Preset Selection Arrows

These arrows are used to navigate through factory presets. Clicking the down arrow moves to the next

6

available factory preset, and clicking the up arrow moves to the previous factory preset.

User Preset Save Button

Saves a user preset, which includes the state of all Reveal parameters.

User Preset Load Button

Recalls a user preset, which includes the state of all Reveal parameters.



User presets do not appear in the Factory Preset pop-up menu, so be sure you are aware of the directory where you are saving them into, so they may be recalled easily later.

Bypass Button

The Bypass button is used to temporarily disable the plug-in. When the plug-in is bypassed, audio playback is not analyzed by the plug-in.

Tool-Specific Controls & Displays


While all of Reveal's analysis tools may be viewed simultaneously in the View All Displays Tab, you may also view each tool independently. Viewing a single tool offers a much larger display area, which can translate into more accurate data interpretation. Having this option allows you to view many types of data simultaneously, or a single type of data in great detail.

Navigating between the various tools can be done in two ways – clicking a tool's mini-display in the View All Displays Tab will take you to that tool's full size display –

or, clicking the desired tool's tab in the lower part of the Reveal interface will do the same. In this section, we'll take a look at each of Reveal's tools.

Oscilloscope Tool

The Oscilloscope displays a trace-style scope waveform. Left channel data is displayed in green, and right channel data is displayed in red. The graph display's X-axis measures amplitude, and is labeled in dB (decibels), while the Y-axis measures time, and is labeled in samples.



Reveal's Oscilloscope Display

Peak and RMS Power History Tool

This tool analyzes an audio file's Peak and RMS ("Root Mean Square" or average) level history. By plotting the Peak and RMS levels in a graph it's possible to monitor the perceived volume of an audio signal over a period of time. The data plotted in the graph display area includes:

- Left Channel Peak Level – Orange
- Right Channel Peak Level – Red
- Left Channel RMS Level – Yellow
- Right Channel RMS Level – Pink




Reveal's Peak and RMS Power History Tool Display

The graph display's X-axis measures amplitude, and is labeled in dB (decibels), while the Y-axis measures time, and is labeled in minutes and seconds. The precision of the intervals along each axis varies, and depends on the zoom level being used.

Spectrogram Tool

The Spectrogram analyzes an audio file's spectral content, and displays it in a graph. A color gradient is used to represent the varying amount of power present in different parts of the frequency spectrum. Red represents the highest amplitude for a given frequency, while green represents the lowest amplitude.



Reveal's Spectrogram Display

Analysis Window Slider/Numerical Field

The Analysis Window control allows varying frequency resolution versus time (attack/decay) resolution. The higher the Analysis Window setting, the more frequency resolution and the less time resolution is displayed, and vice-versa.

A higher Analysis Window setting means that partials/harmonics (i.e., sine waves) with very close frequencies will be seen as separate peaks in the spectrum (each peak is a lobe about 7 spectral points wide). If the Analysis Window setting is steadily decreased, at some point the peaks will merge into one large lobe, and it will be impossible to distinguish whether separate partials exist.

A lower Analysis Window setting means that two attacks which occur very close to each other in time will be seen as two separate events (i.e., the spectrum will jump up,

decay, and jump up again). If the Analysis Window setting is steadily increased, at some point the two attacks will be within one Analysis Window, and the spectrum will only go up once.

Therefore, if partials/harmonics are important for the analysis you are performing, the Analysis Window setting should be increased until separate spectral peaks are visible. Too big of an Analysis Window averages too much of the signal, so fast changes (attacks or fast pitch variations) merge into one blurry spectrum.

For monophonic signals, the Analysis Window setting should be around 4 periods of signal (on average, since the pitch is changing) for adequate frequency resolution (i.e. so that the partials are seen as separate peaks in the spectrum).

By default, the Analysis Window slider is set to 50ms, and has a range of 10 – 100ms.

The graph display's X-axis measures frequency, and is labeled in Hz (Hertz), while the Y-axis measures time, and is labeled in minutes and seconds. The precision of the intervals along each axis varies, and depends on the zoom level being used.



Sliders can be reset to their default values by Option-clicking on them.




When numerical fields are in an editable state (i.e., after clicking in them with the mouse), their values may be changed by typing in new values, by using the Up/Down Arrow keys on your keyboard, or by using the scroll wheel on a Mac OS X-compatible scroll wheel mouse.

Pan Power/Phase Correlation Tool

This tool analyzes phase correlation and balance between a stereo signal's left and right channels. Phase correlation data is shown in the Phase Correlation Meter.

6

Maximum correlation (i.e., a strong "In Phase" reading) means the left and right signals are identical. Zero correlation (i.e., a strong "Out of Phase" reading) means the signals are completely unrelated (i.e., panned hard left or right). Pan power data is shown in the Pan Power Graph. If both channels contain equal power, the meter reading is centered. If one channel contains more power than the other, the meter reading appears stronger for that channel.



The Spectral Analysis Tool's Spectral Graph Display




Reveal's Pan Power and Phase Correlation Display

The Phase Correlation meter displays correlation as a percentage, either in phase or out of phase.

The Pan Power graph display's X-axis measures pan power as a percentage, while the Y-axis measures time, and is labeled in minutes and seconds. The precision of the intervals along the Y-axis varies, and depends on the zoom level being used.

Spectral Analysis Tool

The Spectral Analysis tool analyzes frequency content contained in an audio file, and displays it in the form of a Spectral Meter or Spectral Graph. Left channel data is displayed in green, and right channel data is displayed in red.



Reveal's Spectral Analysis Tool - showing the Spectral Meter Display

Display Pop-Up Menu

This control allows configuring the Spectral Meter or Graph to display Instantaneous, Peak, or Average levels.

Response Time Slider/Numerical Field

The Response Time control is used to set how quickly the display responds, as it receives and analyzes audio signals. A low setting causes it to react relatively quickly, and a high setting causes a slower reaction.

By default, this slider is set to 50ms, and has a range of 10 - 100ms.

Analysis Window Slider/Numerical Field

The Analysis Window control sets the amount of data being analyzed, and allows varying frequency resolution versus time (attack/decay) resolution. The higher the Analysis Window setting, the more frequency resolution and the less time resolution is displayed, and vice-versa.

A larger Analysis Window will better separate partials/harmonics, but will also average spectral variations within the duration of the window.

A shorter Analysis Window will better separate spectral variations over time (i.e., will more clearly show an attack of the piano as a sudden rise of the middle portion of the spectrum), but will not separate partials as well, and eventually, with a short enough setting, discrete spectral peaks for each partial will not be visible.

A good rule of thumb is to make this window three to


four times the average pitch period in the signal. Some common examples include:

- Female vocals – (around 300 Hz pitch), the window should be around 4/300, or 0.01333 seconds, or 13.33 milliseconds long.
- Bass – (around 50 Hz), the window should be around 4/50 or 0.08, or 80 milliseconds long.
- Complete Mix – Use the lowest-pitch instrument (usually the bass) as the guideline for setting the Analysis Window size.

The graph display's X-axis measures frequency, and is labeled in Hz (Hertz), while the Y-axis measures amplitude, and is labeled in dB (decibels).

Lissajous Phase Scope Tool

The Phase Scope tool displays the amplitude and phase of two signals (i.e., the left and right channels of a stereo signal) as a Lissajous pattern, in which the left channel's frequency data is plotted along the graph's Y-axis, and the right channel's is plotted along the X-axis. Lissajous patterns are useful for determining phase and amplitude relationships between the left and right signals of a stereo pair.



Reveal's Lissajous Phase Scope Display

Peak/RMS Level Meters Tool

This tool contains both Peak and RMS ("Root Mean Square" or average) Level metering. Other features include Clip Indicators, Clip Statistics, and variable Peak

Hold, RMS Smoothing, and Meter Range. Each of these features is described below.



Reveal's Peak/RMS Level Meters Display

Clip Indicators

These indicators are located just above the Peak Level Meters. They light up in red when the signal exceeds 0dBFS. The Clip Indicators can be reset by clicking on them.

Peak Level Indicators

These indicators appear as horizontal yellow lines, at the top of the Peak Level Meter display. Various options for these indicators may be set using the Peak Hold pop-up menu.

Overall Peak Level Indicator

This indicator displays the overall Peak Level, with an infinite hold time (similar to setting the Peak Hold time for the Peak Level indicators to Infinity). This indicator can be reset by clicking the Reset All button.

Sample Count in Last Clip Display/Button

This control is both a numeric display, and a button – it shows how many samples were clipped in the last clipping incident. Clicking this button resets the clipped sample count to zero.

Total Samples Clipped Display/Button

This control is both a numeric display, and a button – it

6

shows the total number of samples in all clipping incidents. Clicking this button resets the clipped sample count to zero.

Total Clips Display/Button

This control is both a numeric display, and a button – it shows the total number of clips. Clicking this button resets the clip count to zero.

Peak Hold Pop-Up Menu

This control sets the duration that the Peak Level indicators stay at the Peak Level.

By default, this control is set to 3 Seconds, and options include:

- Off
- 1 Second
- 3 Seconds
- 5 Seconds
- Infinity

RMS Smoothing Pop-Up Menu

This control sets the duration over which RMS level values are calculated. By default, this control is set to 3 Seconds, and options include settings between 1 – 10 Seconds.

Meter Range Pop-Up Menu

This control sets the range of both the Peak and RMS meters. By default, this control is set to -96dB, and options include:

- -48dB
- -96dB
- -144dB

Conclusion

Now that you are familiar with the various analysis tools in Reveal, continue along to the next chapter, which covers SuperFreq.

Chapter 8: Squeeze-3 & -5

Introduction

Squeeze-3 & -5 represent the very latest in professional multi-band compression, limiting, and upward expansion technology. Its linear phase filters give precise control over compression settings for three frequency bands in the case of Squeeze-3, and five bands with Squeeze-5. The advanced visual display provides crucial analysis feedback, and allows parameter adjustment with the highest degree of control.

Squeeze-3 & -5 allow independent compression settings to be applied for low, mid, and high frequencies, and each band can be configured to affect the exact frequency range desired.

This chapter will introduce the graphical user interfaces, how the controls work, and gives some basic hands-on getting started tips for working with Squeeze-3 & -5.

Squeeze-3 & -5 User Interface & Controls

Squeeze-3 & -5 are designed to be intuitive and easy to use. This section discusses the graphical user interface, and will familiarize you with each of the various controls.

Global Controls

Squeeze-3 & -5 features a variety of controls designed for specific tasks, and also features a number of “global”

controls, that affect the state of the entire plug-in. Below is a brief description of these “global” tools, and how they operate.

A/B/C/D Compare Buttons

The A/B comparison feature is common in audio programs, and gives the ability to preview different settings before committing to them. Essentially, various tool settings are made and stored as snapshot “A” – another group of settings can then be assigned as snapshot “B” – and then clicking the A or B buttons in the interface instantly toggles between the two. As Squeeze-3 & -5 offer many fine-tuning options, it features not just two comparison modes – but instead includes A/B/C/D comparisons, for four unique snapshots to compare between, before committing to any file processing. These snapshots are similar to saving and loading presets, but are instantaneous, and do not require selecting presets from a menu to load them. To permanently save one of these snapshots, simply select the desired snapshot, and save it as a custom User Preset.


Factory Preset Pop-Up Menu

This menu displays the current factory preset, if one has been selected. Clicking this menu shows a list of all available factory presets. A factory preset may be selected by choosing it from this pop-up menu.

Factory Preset Selection Arrows

These arrows are used to navigate through factory

Sqweez-3 Graphical User Interface



presets. Clicking the down arrow moves to the next available factory preset, and clicking the up arrow moves to the previous factory preset.

User Preset Save Button

Saves a user preset, which includes the state of all parameters.

User Preset Load Button

Recalls a user preset, includes the state of all parameters.



User presets do not appear in the Factory Preset pop-up menu, so be sure you are aware of the directory where you are saving them into, so they may be recalled easily later.

Bypass Button

The Bypass button is used to temporarily disable the plug-in. When the plug-in is bypassed, audio playback is not affected.

Master Compression Controls

A set of master compression controls appears just above the EQ/Threshold Graph Display area. The controls affect all compression bands. Each control is described in further detail below.

Gain Compensation Button

When clicked, this control toggles automatic gain compensation on or off. When Gain Compensation is active, the output signal approximates the level of the original input signal.

Spectrograph View Pop-Up Menu

This control allows choosing whether left channel, right channel, or stereo channel information is shown when the Spectrograph is active. The spectrograph display appears only in the Threshold View (or in the Threshold portion of the "Both" View).

Spectrograph Mode Pop-Up Menu

This control allows the Spectrograph display to be turned on or off. Two display options are available when the Spectrograph is active – Instantaneous and History. Instantaneous mode displays the current spectral qualities of the audio signal. History mode displays an RMS Average of the file's spectral qualities over the past few seconds.

Spectrograph Pre/Post Pop-Up Menu

This control allows the display of spectral data in a pre-processing or post-processing state. It is useful for visually comparing the "before" and "after" states of compression settings before the plug-in is applied.

Soft Clip Button

When active, this control acts as a soft limiter to ensure that the output level does not clip due to gain differences introduced by a compression setting. The presence of a check in this box means that Soft Clip is active.

Digital Ceiling Slider/Numerical Field

This control sets the maximum allowable level in decibels (used by the Soft Clip control's algorithm).

By default the Digital Ceiling slider is set to 0dB, and has a range of -20 to 0dB.



Sliders/pop-up sliders can be reset to their default values by Option-clicking on them.



When numerical fields are in an editable state (i.e., after clicking in them with the mouse), their values may be changed by typing in new values by using the Up/Down Arrow keys on your keyboard, or by using the scroll wheel on a Mac OS X-compatible scroll wheel mouse.

Input/Output Meters

These meters show levels for both the original signal being input, and the output level, after being affected by the current compression parameters.

Master Gain Slider

This control affects output level globally.

By default, the Gain Slider is set to a value of 0dB, and has a range of -12 to 12 dB.

Graph Display

This display shows a variety of compression data for each band in graphical format. The Graph display may be viewed in three modes – EQ, Threshold, or Both (which displays EQ & Threshold data simultaneously). The graph's X-axis represents frequencies and the Y-axis represents amplitude. In Threshold or Both mode, a spectrograph may be optionally displayed (see description of Spectrograph, earlier in this chapter).

The Graph display presents many useful pieces of information in a way that is easy to interpret with just a

quick glance. The display's elements are described below.

Display Mode Radio Buttons

This set of radio buttons allows choosing between the three available display modes – EQ, Threshold, or Both.

Filter Crossover Indicator/Numerical Entry Field

The Filter crossovers appear in the display as vertical lines, and are used to adjust the width of each compression band (and edit the crossover frequency between it and its neighboring bands). Each crossover indicator features a control tab at its top and bottom, and a numerical entry field so that specific custom values may be quickly entered. Numerical entry fields are staggered so as to fit the interface smoothly – the field for the crossover between bands 1 and 2 appears at the bottom of the indicator, the field for the crossover between bands 2 and 3 appears at the top of the indicator, and so on.

Threshold Indicator

The Threshold indicators appear only in Threshold or Both view modes, and are used to adjust the compression threshold for each band. Each band's Threshold indicator corresponds (and is linked) to the Threshold slider in the band-specific control area below the graph display area – moving the indicator in the graph also moves the slider, and vice-versa.

Gain Indicator

The Gain level for each band is indicated by a thin, light blue line (see Graphical User Interface overview on page 72) in the EQ graph display area. This indicator appears in a horizontal position by default (with all bands set to 0dB, but its shape and orientation change with adjustments made to the Gain slider and Maximum Compression Indicator). The Gain Indicator appears in the EQ and Both view modes.

Maximum Compression Indicator

This indicator displays the maximum amount of compression that can be applied. It appears as a thin white line (see Graphical User Interface overview on page 72) in the EQ graph display area. The Minimum Compression Indicator appears in the EQ and Both view modes.

Current Compression Indicator

This indicator displays the compressor's current frequency response at any given instant. It appears as a thin yellow line (see Graphical User Interface overview on page 72) whose position will vary depending on the current compression settings – but will never exceed the amount of compression set with the Maximum Reduction pop-up slider. The Current Compression Indicator appears in the EQ and Both view modes.

Gain Edit Control

The Gain Edit controls appear as small white three-dimensional balls within each compression band in the graph display. These controls are used to adjust a compression band's Gain level, and they correspond with (and are linked to) the Gain slider in the band-specific control area below the graph display area – moving the Gain Edit control in the graph display also moves that band's Gain slider, and vice-versa. The Gain Edit controls appear in EQ or Both View mode.



Each band contains an upper and lower control ball. The upper one controls Gain, and the lower one controls the Maximum Reduction setting (see next control description).

Maximum Reduction Edit Control Balls

The Maximum Reduction Edit Controls appear as small white three-dimensional balls within each compression band in the graph display. These controls are used to adjust a compression band's Maximum Reduction, and

they correspond with (and are linked to) the Maximum Reduction pop-up slider in the band-specific control area. Moving the Maximum Reduction Edit control in the graph display also moves that band's Maximum Reduction pop-up slider, and vice-versa. The Maximum Reduction Edit controls appear in EQ or Both View mode.

Band-Specific Controls

In addition to the global control discussed in the previous section, band-specific controls are also included. These allow fine-tuning of particular compression characteristics that need to be made only within a certain frequency range.

Threshold Slider

This control allows setting a threshold level – audio signals that exceed this threshold are processed with the current compression settings. Signals that fall below the threshold are not affected by the band's compression settings.

By default, Threshold Sliders are set to 0dB, and have a range of 0 to -90dB.

Threshold Input Meter

Each band features an input level meter, which is specific to the frequencies that fall within the range of that band. These meters are useful for setting the Threshold Sliders within each band.

Gain Slider

This control can be used to manually make up gain that may be lost as a result of compression settings within a given band.

By default, the Gain Sliders are set to 0dB, and have a range of -12 to 12dB.

Ratio Pop-Up Slider

This control determines how much a signal is reduced when it exceeds the threshold level for that band. For example, using a 3:1 ratio means that every 3dB over the threshold level results in a signal that is 1dB over the threshold, which is a reduction of 2dB.

By default, the Ratio Pop-Up Sliders are set to a value of 2:1, and have a range of 2:1 to 50:1.

Maximum Reduction Pop-Up Slider

This control allows setting a maximum reduction value for a given band. Some ratio and threshold settings may result in unrealistic reduction values. This slider allows setting a user-defined maximum amount of reduction.

By default, the Maximum Reduction sliders are set to 0dB, and have a range of 0 to 60dB.

Attack Pop-Up Slider

This control determines how quickly compression is applied, when a signal exceeds the threshold for that band.

By default, the Attack Pop-Up Sliders are set to 30ms, and have a range of 0 to 200ms.

Release Pop-Up Slider

This control determines how quickly compression effect is released, once a signal falls below the threshold level for that band.

By default, the Release Pop-Up Sliders are set to 200ms, and have a range of 0 to 200ms.

Knee Pop-Up Slider

This control determines how hard (dramatic) or soft (transparent) the knee of the compression curve is at the threshold point.

By default, the Knee Pop-Up Sliders are set to .50dB, and have a range of .50 to 3.00dB. (This value corresponds to the interval on which the compression curve is smoothed around the threshold.)

Bypass Checkbox

The Bypass checkbox is used to temporarily disable a specific compression band. When a band is bypassed, audio within its frequency range is not affected by compression settings. The presence of a check means that band is bypassed.

Solo Checkbox

The Solo checkbox is used to temporarily disable all other compression bands. When a band is soloed, only audio within its frequency range is affected by compression settings. The presence of a check means that band is soloed.

Group Controls

This set of controls is used to apply common Ratio, Maximum Reduction, Attack, Release, and Knee settings to all bands simultaneously.

Using Sqweez-3 & -5

This section contains some basic “hands-on” getting started tips for using Sqweez-3 & -5.

To Apply Compression (Basic Steps):

1. Instantiate Sqweez-3 or -5 in Peak, open an audio document, and start playback.
2. In each compression band, use the input meter to help gauge the appropriate Threshold level – and set this using the Threshold slider. The amount of compression applied depends on how far the Threshold slider is set below the meter’s highest reading.
3. Click the Gain Compensation button (or use the Master Gain Slider) to make up for overall output level reduction introduced by the compression settings.
4. In each band:
 - Adjust the Gain slider to a level that suits the content in that band.
 - Adjust the Ratio, Attack and Release to suit your audio material.

To Fine-Tune Compression Settings:

1. Click and drag the Filter Crossover Indicators to create compression bands that span frequency ranges best suited for the audio material you are working with.



The Spectrograph Display can help to visualize the audio material’s frequency spectrum, and is useful in determining the ideal position for each filter crossover.

2. Check the Soft Clip Checkbox to prevent unwanted clipping.
3. Use each band’s Bypass and Solo controls to audition changes to each band’s settings.

4. For a more dramatic compression effect, try the following:

- Decrease the Attack or Release values.
- Set the Knee to a higher value.
- Set the Ratio to a higher value.
- Increase the Maximum Reduction level.

Conclusion

You should now be familiar with the various plug-ins in the Master Perfection Suite, and their basic operation. For more information, or for additional help, please visit the BIAS website, where you can find tutorials, access technical support, and search through a database of frequently asked technical questions.

We hope you enjoy the Master Perfection Suite!

Other Noise Reduction Techniques:

1. Peak has a fairly successful click removal DSP function.
2. You can carefully filter out unwanted frequencies, say, caused by a feedback loop. This can be done in a live situation (preferably) as well as when composing.
3. You can use band-delimited noise gates. We can’t do that yet, but we do in the advanced class.