



Voxengo Primary User Guide

(for VST, AudioUnit audio plug-ins made by Voxengo)

<http://www.voxengo.com/>

Document revision date: 22-Apr-2011

Contents

Introduction 5

Common Features of Voxengo Audio Plug-Ins 5

User Interface Layout 6

Title Bar 6

Control Buttons 6

Plug-In Control Interface 6

Hint Line 6

Main Control Buttons 7

Plug-In Instance Name Edit Box 7

Oversampling Selector 7

Bypass Switch 8

A/B Switch Button 8

“A > B” (“B > A”) Button 8

Reset Button 8

Sat Switch 8

Standard Controls 9

Knob 9

Keyboard Value Entry 10

Value List Selector 10

Slider 10

Level Meter 10

Equalizer 11

Equalizer – Group Editing 12

Equalizer – Spectrum 13

Equalizer – Narrow-Band Sweeping 13

Envelope Editor 13

Scroller 14

Window Resizer 14

Spectrum Mode Editor 14

DC Filter Mode Editor 16

Static Spectrums Editor	16
Undo/Redo	18
Preset Manager	19
Main Preset Manager	19
Local Preset Managers	20
Channel Routing Window	21
Layout	21
Input and Output Routing Selection	21
Show All Channel Meters	21
Mid-Side Pairs Selection	21
Channel Group Assignments	22
Key Signal Sources	22
In Channel Labels	23
Group Names	23
Routing Presets	23
Channel Group Selector	24
Channel Group List	24
Plug-In Settings for a Specific Plug-In	25
Color Scheme	25
User Interface Scale	25
Show Hints	25
Oversampling	25
Meter “Density Mode”	26
Meter Integration Time	26
Meter Release Time	26
Meter Peak Level Hold Time	26
Global Voxengo Plug-In Settings	27
Auto Oversampling Level	27
Mouse Wheel Precision	27
Shift Key Precision	27
Drag Precision	27
Show Knob Labels	27

Tablet Mode	27
Show Virtual Keyboard	27
Radial Knob Mode	28
Control Surface Crosshair	28
Anti-Denormal Noise	28
VST Built-In Refresh	28
VST Function Sync	28
No Multi-Channel Operation	28
Remember Window Positions	28
Visual Settings	29
Plug-In Files' Locations	30
Additional Information	31
Selecting Best Audio Block Size	31
64-Bit Audio Processing	31
Plug-In Registration	31
CPU Load Notice	31
Known Issues	31
Knowledge Bit – Correlation	33
Stereo Correlation	33
Knowledge Bit – Aligning Microphones	34
Alignment Steps	34
Microphone Delay Alignment	34
Microphone Phase Alignment	35
Microphone Gain Alignment	35
High-Frequency Drop Issue	35
Why Go Multi-Microphone?	36
Questions and Answers	37

Introduction

Voxengo audio plug-ins feature a set of standard interface elements that are shared among all Voxengo plug-ins.

This guide describes functionality of these elements, focusing on the interaction between the user and the plug-in. This means that this guide does not describe an actual application of user interface elements (some may control Gain, some may control Frequency, etc.), but shows the ways of using these elements efficiently from the user perspective. After reading this guide you will be able to use common features of all Voxengo audio plug-ins. Any non-standard or special features are described in the guides for the plug-ins that contain these non-standard or special features, and are not covered by this guide.

Beside user interface description this guide contains important technical information about Voxengo plug-ins.

Common Features of Voxengo Audio Plug-Ins

- Preset manager
- A/B comparison switch
- Internal channel routing
- Mid/side processing
- Undo/redo history
- Input channel naming
- Channel grouping
- Channel group naming
- Settings (coloration, scale, hints)
- Global settings (control precision, etc.)
- Plug-in instance naming
- Oversampling
- Bypass switch
- External side-chain (where applicable)
- Contextual hint messages

User Interface Layout

The graphical interface of every Voxengo plug-in consists of four distinctive parts: the title bar, the control buttons array, the plug-in control interface and the hint line.

Title Bar



The title bar shows plug-in's version, plug-in's title, and also contains the Plug-In Instance Name text box, the Oversampling selector and the “Bypass” switch.

Control Buttons



This part of the user interface contains several control buttons: the “Presets” button which opens the preset manager popup window; the quick preset selector button which is displayed on top of the “Presets” button; three change history-related buttons; the A/B comparison button; the “A > B” (“B > A”) button; the “Reset” button; the “Routing” button which opens the channel routing popup window; the quick routing preset selector which is displayed on top of the “Routing” button; and the “Settings” button which opens a popup window with the plug-in's copyright information and additional settings.

Plug-In Control Interface

This part of the plug-in varies greatly from plug-in to plug-in, and an in-depth description is available on a per plug-in basis, in a separate user guide.


Hint Line



This interface element displays hint messages and may also display other informational messages. The hint line can be disabled in the “Settings” window.

Main Control Buttons

Every Voxengo plug-in features a set of standard control buttons. They are shown and described below.



Plug-In Instance Name Edit Box

This control element allows you to give a name to the current plug-in instance. This name is used for informative purposes only: it helps you to remember the purpose of the plug-in, or gives you a note of a name of the host audio channel the plug-in is inserted to. For example, the instance name can be set to “Vocals”, “Drum Bus”, “Polysynth”, “Bass”, “Lead Guitar”, etc.

Oversampling Selector

This selector allows you to select a “quality factor” for the plug-in. An oversampling allows a plug-in to run on a higher internal sample rate thus offering a better overall sound quality. Almost all types of audio processes benefit from an oversampling: probably, only gain adjustment, panning and convolution plug-ins have no real use for it. An oversampling helps plug-ins to create more precise filters with minimized warping at highest frequencies, to reduce spectral aliasing artifacts in compressors and saturators, to improve a level detection precision in peak compressors. The higher the oversampling setting is, the more CPU resources plug-in will require since a CPU load is increased proportionally to the specified oversampling setting: at the “8x” oversampling setting plug-in uses exactly 8 times more CPU time (and that is excluding the time necessary to perform the oversampling itself).

The “Auto” oversampling option is the special option that enables the oversampling when a project is being bounced (exported) to a wave file; in this mode, when the plug-in works in the real-time mode the oversampling is set to the “1x” option (some plug-ins may use “2x”). It should be noted that not all audio host applications provide the plug-in with the required plug-in state information: in those cases the “Auto” option works as the “1x” (or “2x”) option all the time (this is the case for some older Mac OS X audio applications that do not report the off-line bouncing state to the plug-in). The “Auto” oversampling mode engages the “4x” oversampling mode if the working sample rate is below 72 kHz; the “2x” mode is engaged if the sample rate is below 144 kHz; otherwise the oversampling is not engaged at all (for example, if the working sample rate is equal to 176.4 kHz or 192 kHz the oversampling is not engaged). The maximum oversampling level used when the “Auto” mode is engaged can be adjusted in the global settings window. Note that when the plug-in enters the bouncing state, the “Off-Line Render” notification will be displayed on its interface. When this notification is displayed, the “Auto” oversampling option will be working correctly.

In the “Settings” window of most Voxengo plug-ins you can choose the oversampling filter type to use. When the “Min-Phase” type is selected Voxengo plug-ins are using poly-phase IIR low-pass filters with at least 106 dB stop-band attenuation and 6% transition band’s width (which starts at $F_s/2$) for the oversampling. Please note that

these poly-phase filters impose a phase coloration which sounds slightly different on various working sample rates.

Voxengo plug-ins using the “Lin-Phase” filter type use linear-phase filters with 150 dB stop-band attenuation and 4% transition band’s width (centered at $F_s/2$). Note that while working at 44.1 kHz sample rate, linear-phase oversampling filters will cut frequencies above 20 kHz. This is a standard side-effect of oversampling and should be considered normal.

Bypass Switch

The Bypass switch was mainly designed for evaluation of changes the plug-in made. The Bypass switch will not reduce the plug-in's CPU load when switched on. The Bypass state is not saved between project sessions, and is not restored to its previously enabled state when the project is reloaded. This was done so to reduce a possible frustration due to the Bypass switch left enabled by a mistake.

A/B Switch Button

By pressing the “A/B” switch button, you can switch between the current and an opposite (“A” or “B”) plug-in states.

You can use the “A/B” switch button to copy programs between Session Bank’s slots. To do so, in the Session Bank you first need to switch to a program you want to copy and then press the “A > B” (or “B > A”) button. Next, switch to a program where you want to put the first program and then finally press the “A/B” switch button again. (The preset manager and its preset banks are discussed later in this guide.)

“A > B” (“B > A”) Button

The “A > B” (“B > A”) button copies the current plug-in state to an opposite state. You may then switch to this opposite state by pressing the “A/B” button.

Reset Button

This is the master reset button. By pressing this button you will reset the plug-in to its default state. The default state can be chosen in the Preset Manager window.

Sat Switch

Some plug-ins feature a built-in soft-knee output limiter. This switch enables this limiter. This limiter controls output signal level and produces “console saturation” effects. Note that neither limiter mode nor soft-clipping threshold is adjustable.

Standard Controls

Knob



The knob control consists of four parts: the name of parameter it controls, the knob position marker, the center of the knob, and the parameter value readout.

The knob position marker is represented by a rounded light indicator. You may drag this marker to adjust the value of the parameter using the circular knob adjustment approach: during dragging, you may move the mouse cursor farther away from the knob to increase parameter value adjustment precision. You may click on the marker to set the desired parameter value immediately. This way of adjusting the knob is available only if the “Radial Knob Mode” global setting was enabled.

The center of the knob has a defined color that corresponds to the parameter or parameter group the knob controls. You may drag the center of the knob to adjust the value of the parameter with up and down mouse movements, linearly. While dragging the center of the knob with the left mouse button, by additionally pressing the right mouse button you can enable a high precision adjustment mode (which can otherwise be enabled by holding down the “Shift” key). Dragging precision can be adjusted in the global settings window.

The parameter value readout displays the current value of the parameter.

When knob is hovered with the mouse cursor, a scale ring is displayed that shows approximate parameter values at different knob positions (“knob labels”). Thousand values are suffixed with an asterisk (e.g. “2*”). This ring can be disabled in the global settings window.

In some plug-ins you may use the right mouse button to enable knob linking. The linking between knobs is available when two given knobs are logically linked (for example, input and output gain knobs, low and high frequency knobs, etc.). You may hold the “Ctrl” (“Command” on Mac OS X) key before dragging a knob with the right mouse button to enable inverse knob linking: in this mode every positive increment of the knob you drag results in a negative increment in the linked knob.

Knob can be also controlled with the mouse wheel. Double-clicking on the knob resets it to the default state.

Front	Cntr	LFE	Side
-91.6	-91.6	-91.6	-91.6

When working in surround, dual-mono or any other multi-group processing mode, you may hold the “Alt” key to enable linked change of the knob in all channel groups with active channels in them. This may not work for all knobs knob since such feature may not be applicable to a parameter a particular knob controls.

Keyboard Value Entry

Most readout values displayed on the user interface (gain, frequency, parameter value readouts) can be clicked for keyboard value entry (some plug-ins whose parameter value readouts can be dragged require double-click for keyboard entry).

Value List Selector



This type of control allows you to choose a value or an option from the list. You may press the selector button with the left mouse button to display the value list. If you later depress the mouse button on the desired value instead of the selector itself, the list will be closed and the value on which you have depressed the mouse button will be selected.

Note that mode preset lists are different to value lists in that mode preset lists do not highlight a currently selected mode when you open the list.

Value lists (but not mode preset lists) offer you additional ways of item selection without opening the list: you may use mouse's "forward" and "backward" buttons to scroll through the values, and you may press the right mouse button to switch between the current and default value.


Slider



Voxengo plug-ins feature horizontal and vertical sliders. A slider can be dragged with the left mouse button to perform adjustments. While dragging the slider with the left mouse button, by additionally pressing the right mouse button you can enable a high precision adjustment mode (which can otherwise be enabled by holding down the "Shift" key).

If plug-in presents you an array of sliders (like the "Overtone GEQ" plug-in does) by holding down the right mouse button on the slider array you will engage the "drawing" mode that allows you to position sliders in the "free-hand" mode.

Level Meter



Many Voxengo plug-ins feature standard level meters. Each level meter contains level scale (in decibel) and several level bars that correspond to channels ("A", "B",

etc.) of the currently selected channel group (level meter displays all available channels if the “Show All Channel Meters” switch was enabled in the “Channel Routing” window). In some cases level meters are shown in the “minimal” variant with only a single level bar that averages level from all channels that belong to the currently selected group. Gain reduction meters are displayed inversely (from top to bottom), and may show positive gain reduction values since they are showing gain reduction changes relative to the average gain reduction over 2-second time frame.

Some level meters may show a small horizontal white bar which represents peak level. In output level meters such peak level may turn red which means that the output level has entered the area above the 0 dBFS signal level and clipping may occur if the plug-in stays last in the audio host application’s signal chain (if the plug-in is staying in the intermediate position before other plug-ins clipping may not necessarily occur). In non-output level meters used in the plug-in a red peak level may be only an informative indication that some predefined threshold level was reached.

Level meter ballistics and peak level hold time can be defined for all instances of the plug-in in the “Settings” window.

Output level meters usually feature the “Out/In” readout that displays the RMS level difference between the input and output signals of the plug-in (this level difference estimation is based on 3-second integration time).

Equalizer

Several Voxengo plug-ins feature an equalizer with the built-in real-time spectrum analyzer. While the visual appearance of the equalizer is similar among Voxengo plug-ins, the equalizing algorithm used by any given plug-in may be different from other plug-ins.



This picture shows the “control surface” with control points which can be dragged with the left mouse button to adjust the filter’s gain and frequency that corresponds to this control point (you may additionally hold the “Shift” key to enable precise adjustments).

You may select the filter type of the control point by pressing the right mouse button on that control point. By default, the filter type is set to “Off” making the control point appear grey. Alternatively, you may use the mouse’s “forward” and “backward” buttons to scroll through the filter types.

The readouts you see show mouse cursor position within the control surface. The readout also displays the musical note (and detune in cents) that corresponds to the

frequency position. The rightmost readout shows mouse cursor's position within the spectrum power range.




Clicking on a control point with the mouse button will cause the point to be encircled, and the corresponding filter's frequency response curve will be shown in orange. The white curve shown in the equalizer window shows the summary frequency response of all currently enabled filters. When the control point is encircled, three entry fields will be also displayed which you can use to specify filter's parameters via keyboard precisely (these entry fields may not be available in a particular plug-in).

While dragging a control point with the left mouse button, you can adjust the filter's bandwidth by additionally holding the right mouse button (or holding down the “Alt” key). Alternatively, you can use the mouse wheel to adjust the filter's bandwidth. You can hold the “Ctrl” (“Command” on Mac OS X) key while dragging a point to enable its gain adjustment only; holding the “Ctrl” (“Command” on Mac OS X) and “Alt” key simultaneously enables control point's frequency adjustment only. Double-clicking on a control point with the “Ctrl” (“Command” on Mac OS X) key held returns this point to the 0 dB gain position.

Double-clicking on a control point without any key being held temporarily disables it. Second double-click will re-enable the control point. A disabled control point will be re-enabled whenever you move it. Some plug-ins may allow you to double-click the control surface anywhere to add a new filter (the filter type will be auto-selected depending on the frequency where the filter was added at).

Equalizer – Group Editing



You can also perform editing operations on a group of control points. To select several control points, drag inside the control surface itself. A box will appear, and all control points inside this area will become selected. Later you can move a group of selected points like you are working with a single point. To add control points to the current selection, you have to hold the “Shift” key before starting to drag inside the control surface. While holding the “Shift” key, instead of dragging you may also click on the control point you want to add or remove to/from selection. If several points were selected the double click on any of the selected points with the “Ctrl” (“Command” on Mac OS X) key held will reset all selected points to 0 dB. To deselect any currently selected points, simply click the control surface anywhere.

You may right-click the control surface to select all control points at once.

The “Up/Down Arrow” button allows you to scale the gains of all filters up and down. The “Inv” button inverts the gains of all filters. The “Reset” button resets filters to their default state. The “Presets” button opens preset manager for the equalizer.

Equalizer – Spectrum

Beside control points and frequency response plot the equalizer window also displays the Fourier spectrum analysis plot. The spectrum analysis and displaying parameters can be selected via the “Spec” mode selector. The “Spectrum Mode Editor” can be used to customize these parameters further. You may also click the control surface with the left mouse button anywhere to reset spectrum analysis display. Note that Voxengo plug-ins use Hanning windowing function to perform FFT analysis.


A red vertical line that can be displayed if the visible frequency range is wide: this line shows maximal frequency of the input signal and depends on the input sample rate. A second red vertical line may be displayed if oversampling is in use – this second line informs you about the internal frequency range used by the plug-in at this moment.

Equalizer – Narrow-Band Sweeping

You may enable the “narrow-band sweeping” function by first holding the “Ctrl” (“Command” on Mac OS X) key and then dragging the control surface with the left mouse button. As a result of this action you will see the curve of the band-pass filter that only passes the selected frequency range. You may adjust the bandwidth of the filter with the mouse wheel. Such mode is useful at highlighting the resonances in the sound.

Note that the band-pass filter’s curve is applied on top of the existing equalizer curve, and so the curve you see when engaging the narrow-band sweeping is composed of the existing equalizer curve and band-pass filter’s own equalizer curve.

Envelope Editor



Some Voxengo plug-ins feature an envelope editor. The function set of this editor is basically the same as that of the equalizer. The main difference lies in the manner the curve is being drawn. Control points in the envelope editor are always connected by a line or a smooth curve, without intersections, and there are no “filters” with their corresponding “filter shapes” attached to each control point. In the envelope editor you may add and remove control points freely by double-clicking in the desired position within the envelope editor’s area.

To delete a control point you should double-click it. The leftmost and rightmost control points cannot be deleted – they will be moved to envelope’s default position when double-clicked. If you are double-clicking a control point with the “Ctrl” (“Command” on Mac OS X) key held, the control point will be set to envelope’s default position.

While any single control point can be moved freely around the control surface, movements of grouped control points are constrained. There are two types of group movements possible: when you drag a control point (when more than a single control point is selected) the vertical movement is constrained and you can move points horizontally; when you drag a line segment the horizontal movement is constrained and you can move points vertically. If you hold the “Ctrl” (“Command” on Mac OS X) while dragging the constraints are removed.

Scroller



Some Voxengo plug-ins feature horizontal and vertical scrolling controls with zooming functionality. Such scrollers are usually attached to the sides of an equalizer “control surface” discussed above. You may move a scroller only after using its zooming function.

A pair of horizontal and vertical scrollers is accompanied by a rhomb which can be used to control position of both scrollers at once in a single X-Y coordinate space.

Scrollers and rhombs can be double-clicked to switch the visual state of the view they are attached to. You may use this functionality to quickly switch between zoomed and non-zoomed visual states of a view.

Window Resizer



This control that is present in some Voxengo plug-ins and is located in the bottom right corner of the user interface window can be dragged to adjust the size of the window. Note that some audio host applications may not allow you to adjust plug-in’s window freely – in that the window will appear cropped after resize.

Spectrum Mode Editor



The “Spectrum Disable” switch completely disables spectrum analysis function of the plug-in.

The “Type” selector specifies which spectrum analysis type should be used. The “RT Avg” mode engages the real-time spectrum averaging analysis. This type of analysis produces RMS-averaged spectrum over the specified period specified by the “Avg Time” parameter. The “Max” type of analysis produces cumulative maximum power spectrum. The “Avg” type of analysis produces cumulative average power spectrum. The “RT Max” produces real-time maximum spectrum with spectrum fall-off (for better spectrum maximum estimate it is suggested to use a higher “Overlap” setting). If you need an “infinite” peak hold, please use the “Max” analysis type.

The “Filled Display” switch enables additional semi-transparent filling of the spectrum area.

The “2nd Spectrum” switch enables the secondary spectrum curve which is displayed in a darker color. The “2nd Type” selector selects analysis type to use for the secondary spectrum. For example, by setting the “2nd Type” to “RT Max” and “Type” to “RT Avg” you may see the average and maximum spectrums simultaneously. Note that the secondary spectrum uses the same “Block Size” and “Avg Time” parameters as the primary spectrum.

The “Block Size” selector specifies the block size of the FFT (fast Fourier transform) spectrum analyzer. Higher block sizes provide more resolution in the lower frequency range, but decrease time coherence (time precision) in the higher frequency range – the higher frequency information becomes over-averaged. Also, at higher “Block Size” settings the spectrum is refreshed less frequently. This can be mostly fixed by increasing the “Overlap” parameter which increases spectrum refresh frequency. Note that you will generally need to increase the “Block Size” when working at increasingly higher sample rates since the specified “Block Size” setting is used over the full spectral bandwidth, and so at higher sample rates analyzer’s resolution in the visible frequency range will be lower for the given “Block Size” setting.

The “Overlap” parameter controls the overlap between the adjacent FFT spectrum analysis windows. Higher overlap values allow spectrum to be updated more frequently.

The “Avg Time” parameter specifies averaging (fall-off) time (in milliseconds) used when the “RT Avg” or “RT Max” analysis is active. This value specifies time it takes spectrum level to fall down by 20 dB.

The “Smoothing” parameter selects spectrum smoothing function’s resolution in octaves. Note that smoothing being solely visual produces 6 dB per octave drop when stationary sine-wave signals are analyzed. E.g. even if the signal consists of 2 sine-waves (1 kHz and 2 kHz) of equal peak amplitude, the 2 kHz sine-wave will look like it is 6 dB quieter. This happens because fast Fourier transform produces a narrower spectrum for high-frequency stationary signals in comparison to low-frequency stationary signals. This drop does not practically appear when non-stationary (musical) signals are being analyzed.

The “Freq Low” and “Freq High” parameters (if they are available in a given plug-in) specify the visible frequency range (in Hertz) of the spectrum view.

The “Range Low” and “Range High” parameters select accessible spectrum power range (in decibel).

The “Slope” selector (which is defined in dB per octave) allows you to adjust spectrum analyzer display’s slope around 1kHz. Skewing the spectrum can be useful because

higher frequencies usually have weaker power in comparison to the lower frequencies. By choosing an appropriate spectrum slope, you can compensate for this fact and make the spectrum plot look more convenient and meaningful.

DC Filter Mode Editor



This window allows you to fine-tune the DC filter, and create user presets if required. DC filter (also known as high-pass filter) is applied after all processing stages, and is used for removing ultra low-frequency content introduced by plug-in's processing modules.

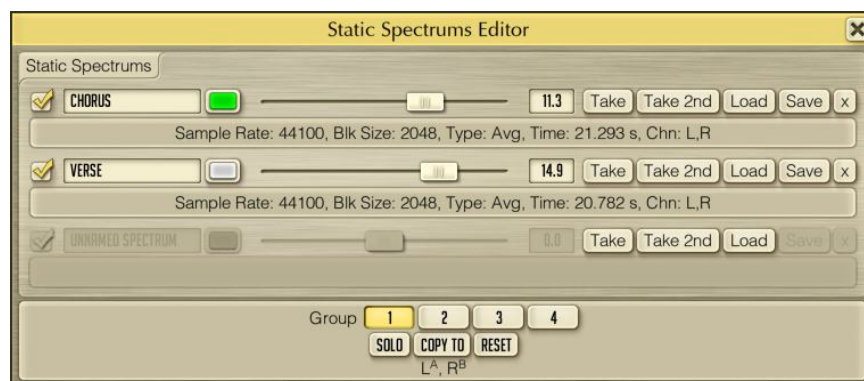
The “DC Filter Enable” switch enables DC filtering stage.

The “Freq” parameter specifies the corner frequency of the DC (high-pass) filter.

The type selector specifies DC filter type. The “Steep” filter types offer a steeper cutoff frequency response, but a higher phase shift. The “Soft” filter types have a smoother cutoff frequency response and a lower overall phase shift. The number before the filter type denotes dB/octave roll-off of the filter (e.g. “-18 Soft” means “-18 dB/octave soft filter”).

Note that DC removal filter induces phase shift and thus may affect sonic character of a sound material. For best results you may specify DC filter frequency as low as possible and use non-steep filter types (preferably, -6 dB/octave one). Due to its phase shifting, the DC filter may be also used creatively to adjust the sonic character of a sound material in the low-frequency region.

Static Spectrums Editor



Some equalizer Voxengo plug-ins feature static spectrums display that can be controlled via the editor window shown above. You may select the display name of

the spectrum slot, its color, the shift in decibel of the static spectrum on the control surface. The static spectrum can be shown or hidden via the visibility check box.

The “Take” and “Take 2nd” buttons take a snapshot of the primary or secondary spectrum, respectively. If after pressing any of these buttons no spectrum is taken, it means that no spectrum is currently available – you either have to configure the spectrum mode or start the audio playback first. Before taking a spectrum, please do not forget to choose an appropriate spectrum analysis type via the “Spectrum Mode Editor” – usually “Avg” or “Max”, and analyze long enough so that spectrum becomes general enough. When analyzing a song, it may be a good idea to store separate spectrums for verse, chorus and bridge parts as they may have distinctively differing spectral balance.

The “Load” and “Save” buttons allow you to load and save the spectrum file into the static spectrum slot. The files have the “.csf” (compressed spectrum file) file extension.

The “x” button resets the spectrum in the selected slot.

When a spectrum is available in the static spectrum slot, the spectrum’s parameters are also displayed. They correspond to parameters that were configured via the “Spectrum Mode Editor” when the spectrum was taken.

Undo/Redo

Voxengo plug-ins offer a convenient way of managing changes you apply to the plug-in's state. When you apply a change to any control, its previous and new values are stored in the special “change log” which you can then use to “unwind” and “rewind” any previously made changes. This way you can be sure that no single unwanted or occasional change you do affects the state of the plug-in in a wrong way.

Undo/redo control consists of three buttons: the “Undo”, the “History” and the “Redo”:



The “Undo” button allows you to un-do changes made to the plug-in's state.

The “History” button opens change log (history) that lists all changes made to the plug-in's state in the order you have made these changes before. You can re-do and un-do changes to any level of depth, instantly. Change log keeps track of last 32 changes you have made. Note that a parameter change is logged with group name placed in parentheses (e.g. “Gain (Ls) change”).

The “Redo” button allows you to re-do changes that were already un-done via the “History” or “Undo” buttons.

Preset Manager

Main Preset Manager

Every Voxengo audio plug-in features the main preset manager which you can use to store plug-in state presets you define.



Presets in the main preset manager are shared among all instances of the same Voxengo plug-in. All presets within the main preset manager are stored in user preset banks. Beside user preset banks two special banks exist: the “Session Bank” and “Factory ROM” bank.

The “Session Bank” contains “programs” rather than presets. Each program in the “Session Bank” contains its own undo/redo change log, meaning that when you switch the program in the “Session Bank” you are also switching the undo/redo context. A currently selected program is specified by the “right arrow” symbol. The “Session Bank” lists programs that “mirror” programs of the VST host audio application (in AudioUnit-compatible applications this bank can be used as a temporary preset palette). When you are activating a program in the “Session Bank” you are actually switching the current program in the VST audio host application (in AudioUnit-compatible applications program switch in the “Session Bank” is similar to a plain preset switching).

The “Factory ROM” bank lists presets that were pre-programmed by the plug-in's producer. These presets cannot be changed, but can be activated. The “Factory ROM” bank is also loaded into the “Session Bank” every time a new instance of the plug-in is created in the audio host application.

The main preset manager contains the following control buttons:

The “+”, “-”, “Load” and “Save” buttons allow you to add and remove the bank or preset, and load and save the bank or preset to and from a disk file. Right-clicking the “+” button inserts preset into the current list position rather than to the end of the list if the left mouse button is clicked. The “U” button allows you to update a currently selected preset by storing the current plug-in state in it.

Please note that Voxengo plug-ins use a proprietary format to store presets and preset banks. We suggest that you add a comprehensive name prefix to bank and preset files so that you would not mix up presets created in different Voxengo plug-ins. During a save operation, the plug-in will offer you a default name prefix automatically. Voxengo plug-in preset files possess the “.cpf” file extension while preset bank files possess the “.cbf” file extension.

The “Set as Default” button allows you to designate a selected preset to become the default preset. The default preset is loaded every time a new plug-in instance is created in the audio host application, or when the master “Reset” button is pressed. If you want to restore the original default preset you should press the “Set as Default” button on the “Default” preset in the “Factory ROM” preset bank.

Double-clicking on a preset name loads the specified preset state. The same does the “Activate” button. A preset load operation can be undone by pressing the “Undo” button.

To change a name of the preset or bank you should at first select the required item and then after a small delay click the item again. Please do not mix this operation with a double-click which activates a selected preset.

Local Preset Managers

Voxengo plug-ins also feature so called “local” preset managers: for example, the one available in the “Channel Routing” window. These preset managers do not possess preset banks, and are not related to the main preset manager: they work independently from the main preset manager.



Local preset managers usually manage presets for a specific module of a specific Voxengo plug-in. The “Update” button can be used to quickly update selected preset to the current state.

Presets available in the “Routing Presets” local preset manager are globally accessible to Voxengo plug-ins of all kinds. These presets greatly minimize routine adjustments you usually need to perform after creating a new plug-in instance in your audio host application. Routing presets that are not compatible with the current plug-in are shown grayed out with the “(n/a)” prefix.

Channel Routing Window

Layout

The “Channel Routing” window consists of a set of selection buttons assembled in the following groups: the “Input Routing”, the “Mid-Side Pairs”, the “Group Assignments”, the “Key Signal Sources” and the “Output Routing”:



Input and Output Routing Selection

These groups of buttons allow you to route external plug-in inputs to internal plug-in channels and vice versa: to route internal plug-in channels to external plug-in outputs. A plug-in has a pre-defined number of internal channels while a number of input and output channels may vary depending on the host audio application track or bus where the plug-in resides. Note that if input routing selector appears red, it means that this selector refers to a non-existent input channel: you may correct this warning by selecting an existing channel. External side-chain inputs are denoted by parenthesized labels: for example, “(IN3)”, “(IN4)”.

Show All Channel Meters

This option enables displaying of all channel meters and statistics counters regardless of the currently selected channel group. When this option is not enabled, only meters belonging to the currently selected channel group will be shown. Enabling this option is useful when you are using dual mono or mid-side processing: this option will allow you to see channel meters for left and right, or mid and side channels together.

Mid-Side Pairs Selection

These buttons (if available in a particular plug-in) allow you to assign internal channels to mid-side encoding/decoding pairs. The mid-side encoding is a wide-

spread technique that allows you to process stereo signal's middle (center) and side (spatial) information independently of each other thus offering a great deal of control over that signal's stereophony. Mid-side encoding works with paired channels only and thus requires two channels to be assigned to the same mid-side pair. An input signal is mid-side encoded before it is processed by the plug-in, and decoded afterwards before it is routed to an output of the plug-in.

Channel Group Assignments

The plug-in allows you to assign its internal audio channels to logical groups of channels. Each group of channels is affected by its own set of parameter values (it can be an EQ shape, a gain factor, an overdrive setting, etc.)

Individual audio channels can be attached to channel groups. For example, this allows you to have individual EQ settings for channel 1 and for channel 2, separately. To achieve this you may simply assign channel 1 to group 1, and channel 2 to group 2.

In a surround setup you may assign left and right channels to group 1, and surround channels to group 2, and apply an EQ shape to them independently from each other.

Each plug-in audio channel can be assigned to a single channel group only. Channel grouping also affects channel-linking in case of dynamics processing and other processes that estimate signal loudness envelope: channels assigned to the same group will be linked during processing and signal loudness estimation. The mode of such linking (max of all, average of all, etc.) may be specified with an additional plug-in control, if available.

Key Signal Sources

This set of buttons (if available in a particular plug-in) allows you to assign key signal sources to all internal channels. For example, instead of using channel's "A" self signal to trigger processing performed on this channel you may assign channel "B" to trigger processing performed on the channel "A". Key signal sources are usually used in dynamics processing plug-ins. The technique when any particular channel is used as a "key" for another channel is called ducking: "key" channel "ducks" another channel. E.g. bass guitar track may be "ducked" by a drum track if necessary.

If some internal channel is only used as key signal source (e.g. this channel is an external side-chain signal), this channel can be "muted": in such case this channel will not be processed and sent to the output together with channels that were not muted. Such "muted" channel will be available to plug-in's internal algorithms only.

Please note that external side-chain sources are always represented by individual input channels in the plug-in. While audio applications that support AudioUnits usually offer a convenient one-click "side-chain" input selection for the plug-in, these side-chain inputs are shown as additional input channels in the "Channel Routing" window of the plug-in, with their names parenthesized: e.g. "(IN3)". These channels should be routed in the same way as ordinary channels are routed, including setting up required key signal sources assignments. Even though such way of setting up external side-chain does look somehow complicated, it offers a very deep control over side-chaining: for example, it allows you to perform mid-side encoding of the side-chain input before it is used for keying; another useful feature is that you can swap channels of the side-chain input effortlessly.

In Channel Labels

By pressing this button you can open the channel label assignment popup window. You can give individual names to input channels in order to simplify a further channel routing setup understanding. Channel names you provide are also displayed on the user interface to inform you about channels that are currently being adjusted by the plug-in's controls.

You can also import channel labels from the host audio application by pressing the "Import labels from host" button. However, not all audio host applications provide distinctive input channel names: in that case entering your own names manually is the only existing option (Mac OS X AudioUnits do not provide input channel names).

Group Names

This button opens the group name assignment popup window. You may assign names to channel groups for easiness of channel group selection.

Routing Presets

The "Routing Presets" button opens a local preset list that covers settings present in the "Channel Routing" window, including channel labels. These presets are shared among Voxengo plug-ins of all types: note that due to this fact a particular preset (e.g. the one that uses mid-side pairs or external side-chain assignment) may not be applicable to a specific plug-in where the used feature is unavailable (not all plug-ins possess mid-side processing and external side-chain) – in this case such preset will be shown grayed out, with the "(n/a)" prefix.

Channel Group Selector

Almost every Voxengo plug-in features the channel group selector which allows you to select the channel group for which parameters are currently being edited or monitored.



This selector features a set of buttons that represent channel group names (indices, by default). Only groups actually assigned to the internal channels in the Channel Routing window are shown. You may use these buttons to switch between channel groups.

The “Solo” switch allows you to solo channels that belong to the currently selected group. This can be useful if you want to focus listening on group’s channels only. The state of the “Solo” switch is not saved between project sessions, and is not restored to its previously enabled state when the project is reloaded.

The “Copy to” button allows you to copy parameter settings defined for the current channel group, to another channel group. The “Reset” button resets parameters defined for the current group to the default values.

Channel Group List

$(L+R)^A, IN3^C, IN4^D$

A Voxengo plug-in also presents a list of input channels being routed to a currently selected channel group. This list is interconnected with the “Channel Routing” window, and displays routing settings defined by it. This list contains names of input channels, comma delimited. Internal channel name (“A”, “B”, “C”, etc.) that accepts the corresponding input channel is displayed in a superscript style (these internal channel names are also displayed on the level meters). If more than one input channel is routed to the same internal channel, the sum is displayed in the form “(IN1+IN2)”. Additionally, when the internal channel is assigned to a mid/side group, its input channels are taken into parentheses with the “m” (mid) or “s” (side) prefix: “s(IN1 & IN2)” which means “side part of the mid/side pair consisting of IN1 and IN2 input channels”.

When a plug-in supports side-chain, any side-chain assignments are displayed with a special postfix: “:sc()”. For example, “L:sc(scL)” means that the input channel “L” is affected by the side-chain signal from the input channel “scL”.

Plug-In Settings for a Specific Plug-In

Voxengo plug-ins feature a set of settings that affect instances of the same plug-in only: they do not affect the same settings of other Voxengo plug-in types. This is convenient since this way you can select different color schemes for different plug-in types (for example, a Voxengo compressor may be colored in reds while a Voxengo equalizer may be colored in grays). The same applies to the user interface scale and the hints display switch: for example, if you have not still mastered a specific Voxengo plug-in you may leave the hints for this plug-in visible while you may disable the hints for a Voxengo plug-in you know well.

Color Scheme

The color scheme selector consists of several small preview icons that show possible color scheme selections. You may press a desired icon to change the color scheme of all instances of the plug-in. Color scheme switches immediately in all open plug-in instances, and instances that will be opened later.

User Interface Scale

By using this setting you may adjust the visual scale of the plug-in. For example, if your screen resolution is very high (1920 x 1200 or more) you may enable the largest scale (125%). On the other hand, if your screen resolution is low (1024 x 768) or you want to have many plug-in instances visible on the screen, having a smaller user interface size can be beneficial: in that case you may select the 80% scale. Note that the visual scale switch requires a full audio application restart.

Show Hints

This switch allows you to enable or disable hint messages that appear on the bottom line. The hint state changes immediately in all open plug-in instances, and instances that will be opened later.

Oversampling

This selector, if available in a particular plug-in, allows you to choose the type of oversampling filter to use. Oversampling is discussed in the topic “Main Control Buttons – Oversampling Selector” above. The “Min-Phase” option selects minimum-phase oversampling filter which is easy on CPU and does not add a considerable processing latency. The “Lin-Phase” option selects linear-phase oversampling filter which is CPU-intensive and adds around 11 milliseconds of processing latency. At the same time, the linear-phase filter adds no phase coloration to the sound material while the minimum-phase filter adds such coloration (which is, however, mostly unnoticeable).

Note that in VST audio host applications this switch should be adjusted with care. While with most hosts a simple playback restart is enough for the change to take place, in some hosts it may require full application restart and plug-in rescan to change the processing latency. Not doing so may make tracks where the plug-in is used go out of sync. The plug-in window will show the red “Latency Changed” label if

playback restart (or application restart, or current project reload) is required due to the processing latency change that occurred.

Meter “Density Mode”

This checkbox enables “density” metering mode. In this mode you can see levels a signal stays at often. By examining the range of levels a signal stays at you may draw conclusions about effective dynamic range of the sound material. Note that signal level estimation in the “density” mode is affected by the configured meter’s integration and release times. Configured meter’s peak level hold time affects time duration a specific signal level registers on the meter in this mode.

Meter Integration Time

This parameter affects signal level integration time (in milliseconds) of all level meters (this value reflects time it takes for a signal level to fall down by 20 dB, or raise up from any steady level to another steady level). Note that this setting does not affect peak level registered on the level meters, but directly affects the visible difference between the peak and RMS signal levels when a musical signal is metered.

Meter Release Time

This parameter changes level meter’s release time (in milliseconds), and is also known as level meter’s “ballistics”. This is the time it takes a signal to fall down by 20 dB.

Meter Peak Level Hold Time

This parameter adjust time (in milliseconds) a registered 1-sample wide peak level stays on the level meter unchanged.

Global Voxengo Plug-In Settings

Voxengo plug-ins also feature global settings that are shared among all Voxengo plug-in types. These settings mainly customize your interaction with various controls: mouse wheel precision, knob drag precision, etc.

Auto Oversampling Level

This option selects oversampling level used by Voxengo plug-ins when the “Auto” oversampling option is selected. This option selects maximal oversampling level used. The actual used level may be lower depending on the working sample rate (the higher the working sample rate the lower the auto-selected oversampling is). Note that change takes effect only after the host audio application is fully reloaded.

Mouse Wheel Precision

This setting affects a precision of the mouse wheel – the higher the precision the finer the control change that can be done with the mouse wheel.

Shift Key Precision

This setting affects a control change precision when the Shift key modifier is pressed during dragging the control with the mouse.

Drag Precision

This setting affects precision of mouse drag operations. This applies to knobs and readouts.

Show Knob Labels

This switch allows you to enable or disable numeric labels that appear around the knob when you are hovering your mouse over it.

Tablet Mode

This switch enables “tablet” mode allowing you to control Voxengo audio plug-ins with a pen tablet. This mode can be also used if you are operating plug-ins via remote desktop software.

Show Virtual Keyboard

When this switch is enabled, value entry will be accompanied with a popup window showing the “virtual keyboard” you can use instead of your real keyboard to enter various values (numeric and textual). Virtual keyboard is especially useful when the audio host application “blocks” certain keys from reaching plug-in’s user interface.

Radial Knob Mode

This option enables radial operation mode of the knob. You may disable this mode if you usually use linear knob mode only.

Control Surface Crosshair

This option enables crosshair display in the plug-ins on their “control surface” areas – e.g. equalizer, spectrum analyzer.

Anti-Denormal Noise

This switch enables insertion of anti-denormalization noise on plug-ins’ inputs. This noise has RMS of -220 decibel – well below the audible dynamic range. If you are using Voxengo plug-ins in a host audio application that applies such noise automatically you may disable this switch to preserve some CPU resource. Without this anti-denormal noise the filters and oversampling of the plug-in may overload CPU when silence is processed.

VST Built-In Refresh

This switch enables VST specification’s built-in user interface refreshing timer. This switch is available for VST plug-in format only. Enabling this switch will make all user interface refresh operations dependable on the audio host application. This may or may not be preferable, and depends on the overall behavior of the plug-in in a particular audio host application you are using. For example, if you have noticed some strange behavior of the plug-in’s user interface during audio project bouncing you may try enabling this option.

VST Function Sync

This switch enables internal VST function synchronization that solves crash problems in some hosts at the expense of minor added CPU load. For example, Tracktion 3.0.4.8 may call plug-in’s “suspend” function together with the audio processing function: this, of course, leads to a crash. You may enable this switch to ensure no interfering functions will get called simultaneously.

No Multi-Channel Operation

This switch forces stereo and mono (no multi-channel) operation in Voxengo plug-ins. Enable this switch if you are using VST audio host applications that do not support multi-channel VST plug-ins correctly. To take effect, this switch requires audio host application restart.

Remember Window Positions

This switch allows you to enable or disable saving of window’s position when it is reopened after closing. Window’s position is saved relative to the parent window.

Visual Settings

Latest Voxengo plug-ins feature a set of visual settings that allow you to subtly customize the look of the plug-ins. These settings are:

1. Flat panels – when enabled, all buttons and panels of the plug-ins will look flat, without gradient fill.
2. Spotlight – enables a wide light area that looks like a trace of spotlight.
3. Textures – enables texturing of the panels.
4. Shadows – enables shadows on visual elements.

Plug-In Files' Locations

Any Voxengo plug-in is bundled in a separate archive file for every target platform. These archives contain document files, a license agreement, a change history log and a plug-in itself in the form of the dynamic link library (DLL) or the component file, for any given target computer platform. In order to be able to instantiate the plug-in in the host audio application you should put the DLL (component) file into the correct file folder that corresponds to the “plug-ins path” specified in the audio application. It is suggested that you put the DLL (component) file into the “Voxengo” subfolder within the plug-in folder, for easier searching.

While it is mandatory to read and agree to the supplied license agreement before using the plug-in, you may for your convenience put the DLL (component) file separately from its accompanying files, including the license agreement file. However, as license agreement explicitly mentions, the DLL (component) file cannot be re-distributed to other persons or companies/entities separately. You may store the DLL (component) file separately from its accompanying files for your own archival and storage convenience purposes only.

During its operation, the plug-in also creates settings files where it stores global and per plug-in settings, including presets. On Windows systems the data files created by Voxengo audio plug-ins reside in the “\Documents and Settings\<UserName>\Application Data\Voxengo\Audio Plug-Ins\” folder; on Mac OS X systems the data files reside in the “/Users/<UserName>/Library/Preferences/Voxengo/Audio Plug-Ins/” folder. You may safely remove, copy and replace these files, including the “Voxengo\Audio Plug-Ins\” folder itself. All Voxengo plug-in settings and presets are available to the specific user of the computer only.

Additional Information

Selecting Best Audio Block Size

The most efficient audio card block size (latency) to work with Voxengo plug-ins is 512 samples (plug-in's internal processing overhead may become considerable at low latencies such as 64 samples – the plug-in may require 45% more CPU time at such low latency).

64-Bit Audio Processing

Plug-in conforms to the VST specification version 2.4, and so it is able to interface with audio host applications in full 64-bit floating point resolution (if supported by the host audio application).

AudioUnit version of plug-in is technically able to process 64-bit floating point audio data.

Plug-In Registration

Commercial Voxengo plug-ins should be registered before you can use them in non-demo mode. Each commercial plug-in in its Settings window contains a special area where user registration details can be entered. Please use the standard copy&paste functions to ensure a correct product key entry (you may right-click the product key entry field to paste the key). User name is not required for entry. When details were entered correctly, you will see the “Registered to” text in the Settings window: in this case plug-in is working in a non-demo mode.

CPU Load Notice

Before judging the plug-in's CPU load performance please make sure you understand that on multi-processor (multi-core) computers CPU load metering offered by the audio host application you are using can only be approximate. The reason for this is the fact that depending on the track configuration of your project, you may get a varying maximum number of plug-ins that can work at the same time. For example, if your project consists of only a single track with a long effects sequence, you may quickly run out of CPU resources while if you are putting these effects on parallel tracks, the maximum number of these effects you can use will double at the least.

Also pay attention to the “Oversampling” setting of the plug-in, because running a plug-in at 4x oversampling setting is the same thing as running 4 separate plug-ins at 1x oversampling.

Known Issues

These issues are known and may be resolved in the future. In most cases, these issues do not make Voxengo plug-ins unusable:

1. Plug-in's preset manager on Mac OS X may appear behind plug-in's interface window. This problem can't be resolved from the plug-in side since it is related to

window hierarchy and how it was setup by the audio host application: almost no Mac OS X audio host applications honor audio plug-ins to open their own utility windows.

2. Routing preset selector shows presets that make use of mid/side processing or external side-chaining even if the plug-in does not support these functions. Loading such presets does not have any negative impact on the plug-in other than causing confusion about specific function availability in the plug-in.
3. Bias Peak Pro 6 VBox on Mac OS X. First issue happens after opening the preset/bank load/save dialog from within plug-in's own preset manager: the plug-in's interface stops responding afterwards. Another issue with VBox is more serious and may lead to an application crash. We believe this is caused by VBox's own problem that is caused by VBox calling plug-in's "Reset" and "Render" functions simultaneously from different threads.

Knowledge Bit – Correlation

As you probably already know, correlation between two independent signals is defined in the range between -1.0 and 1.0, inclusive. If correlation is close to 1.0 between any two signals, it can be said that signals are “in-phase” (0 degree phase difference). If correlation is close to -1.0 the signals are “out of phase” (180 degree phase difference). Correlation values close to 0.0 usually tell that any two signals have no correlation and so they may either be 90 degree apart or significantly delayed relative to each other, or both.

Correlation between two sinusoidal signals is equal to mathematical cosine of the phase difference between these signals.

Note that while simple sinusoidal stationary (constant) signals require only phase rotation to change correlation value, musical signals and signals which are complex in their nature may require time alignment (delaying) as well. Correlation of any signal with any independent (uncorrelated) noise signal is 0.0 and cannot be changed with phase nor time alignment.

Stereo Correlation

Speaking about stereo (two-channel) signals, “acceptable” correlation value range lies between 0.0 and 1.0. Negative correlation values mean that channels are out of phase and this usually works problematically – stereo field becomes “unreal” (“surround”) and causes disorientation. Beside that, out of phase stereo information is not mono-compatible and such signal sounds poorly if the listener is not placed on the central listening axis.

In order to create any useful spacious stereo-image the correlation values close to 0.0 should be used. Also note that an uncorrelated stereo-signal in comparison to a correlated stereo-signal at an equal peak level usually sounds louder to the listener by around 1.25 decibel, because uncorrelated channel sound coming from a speaker does not cancel channel sound coming from another speaker spaced 60-degree apart from it as much as correlated channel sound cancels it.

Knowledge Bit – Aligning Microphones

This topic discusses possible ways of achieving a good alignment between several microphones providing a maximally focused and feature-rich sound. While this topic is easier perceived from the perspective of Voxengo PHA-979 plug-in user, information contained herein can be used universally.

When you are placing a couple of microphones to capture a single instrument, this means you are creating a microphone array consisting of two microphones.

When doing microphone alignment you should always consider only two microphones and a single sound source at a time. One of the microphones should be considered as the “master” microphone and delay, phase and gain should not be changed. When working with three or more microphones in an array, each additional microphone should be aligned to the master microphone, one by one, so that only two microphones are considered at a time.

Each microphone in this array has these basic parameters: distance from the sound source, delay relative to the master microphone and phase relative to the master microphone. Relative delay can be negative meaning the sound reaches the microphone earlier than the master microphone.

All directed microphones should generally capture a single sound source or otherwise you may face problems getting all sound sources to sound good when trying to align microphones, because in such case you will have to make complex decisions to find a good balance of per-microphone phases and delays.

Alignment Steps

Aligning any given microphone in the array relative to the master microphone basically consists of three steps: 1) finding an optimal delay; 2) finding an optimal phase; 3) finding an optimal gain.

First of all, it should be said that microphone delay, phase and gain alignment is required only if you are capturing the direct sound of an instrument with a microphone array. Microphone gain alignment may become unimportant or even problematic if you are capturing “space” where the instrument is located (i.e. surround microphone array, room microphones, etc), because at longer distances from the sound source the microphones may capture room’s reflections non-proportionally to instrument’s direct sound.

Microphone Delay Alignment

This method is best used with percussive instruments that produce a sharp transient. After doing a multi-microphone recording you should inspect the recorded waveform at some high zoom setting, and write down the sample position of the transient in both the master microphone track and the microphone track you are aligning (you can also use PHA-979’s “Calculator” for this). The time difference will be the delay required to be applied for the microphone to be aligned to the master microphone track.

Microphone Phase Alignment

In order to find the best phase alignment (rotation) value of a recorded microphone track relative to the master microphone track you have to put that microphone, master microphone and a sound source on a single imaginary plane. Let's call these imaginary points "Z", "M" and "S", respectively. Let's "R" be the angle between the lines "S"-"Z" and "S"-"M" on that plane. Now, rotate that plane in your imagination around "S"-"M" line, preserving distances between the points, so that the line "S"-"Z" becomes parallel to the floor. Imagine that the point "Z" being at your chest. Then, if the point "M" imaginary appears left-hand of your chest you should set the phase rotation to "R", while if the point "M" appears right-hand of your chest you should set the phase rotation to minus "R".

The phase alignment should basically be done with directionality axes (vectors) of all microphones pointing directly at the sound source. If this is not the case, an additional phase rotation angle should be added to make each microphone on-axis.

Any stereo microphone pair used in configurations like "ORTF", "X/Y" or "A/B" should be treated as a single microphone, with the directionality of this pair being equal to vector sum of directionality vectors of each microphone in the pair.

Note that a microphone by design may have its polarity inverted by 180 degrees relative to the master microphone. When phase-aligning such microphone you may need to additionally flip its phase by 180 degrees.

Microphone Gain Alignment

When doing microphone gain alignment you may use the following approach. This approach requires some simple signal or sine-wave (collectively, a "test tone") to be played from sound source's position: this can be a handclap or a sweeping sine-wave tone. Power of signal reaching each microphone can then be analyzed easily at all required frequencies to find the required gain adjustment value for each microphone. For more consistent results you may need to analyze powers of lower frequencies below 500 Hz only.

Alternatively, you can use the theoretical rule that states that the power of signal when it propagates through the air drops by 6 dB with each doubling of the distance from the sound source. E.g. if microphone A was placed 10 inches from the sound source while microphone B was placed 20 inches from the sound source, microphone B should be made louder by 6 dB to match the loudness of a signal in microphone A. Of course, for this approach to work, the microphone pre-amp gain and "sound pressure to voltage" characteristic of the microphones should be matched (or microphones of the same model should be used) which makes the first "test tone" gain alignment an easier approach in practical situations. Also note that this "theoretical" approach does not account for any additional sound reflections that are reaching the microphones and which may cause interference, effectively increasing or decreasing the signal power of the recording.

High-Frequency Drop Issue

When a signal from two given microphones that were aligned and gain-compensated is being mixed, a problem of high-frequency drop may arise. This usually happens when you mix two spaced microphones placed at different angles and distances to the

sound source. The main reason of such behavior lies in the difference of propagation of frequencies in space.

While lower sound frequencies usually penetrate the matter quite easily and propagate through the space without much damping and reflections (due to a longer wavelength), higher frequencies are easily absorbed and reflected even by small obstacles. Thus, in most cases, higher frequencies from the same sound source reaching two spaced microphones usually become uncorrelated, even after the microphone delay and phase alignment has been done. Due to their small wavelength higher frequencies also require only a very short time delay to become uncorrelated even if no additional absorption or reflection is taking place. The lower frequencies reaching two aligned microphones are usually always correlated.

When you are mixing two correlated signals in equal proportion you get 6 dB of overall loudness gain. But when you are mixing two uncorrelated signals the gain can vary between 3 and 6 dB (depending on the level of correlation) causing high-frequency drop relative to the overall gain in the lower frequencies. This drop becomes bigger with each additional spaced microphone added to the array, and, for example, may reach 6 dB when you mix four spaced microphone signals.

Why Go Multi-Microphone?

If you generally wonder why one should use multi-microphone arrays instead of a single microphone when recording a sound source, you should understand that various acoustical instruments have quite a complex sound emitting profile: the spectral content of side, front, bottom, etc. surfaces of the instrument may be quite different. For example, snare drum's drumhead side does not have much "crack" sound in it while snare drum's side provides a lot of "crack"; on the bottom of the snare drum the sound of resonating snares is prevailing. All these aspects cannot be captured with a single microphone if one wants to achieve defined and focused recorded sound: getting such sound requires close microphone placement with a further time and phase alignment between the recorded microphone tracks.

Questions and Answers

Q. Why some Voxengo plug-in controls appear dimmed after a new Voxengo plug-in instantiation or after loading a previously saved project?

A. This is a normal behavior since Voxengo plug-ins use dynamic internal processes allocation. When some controls appear dimmed it means that internal processes that correspond to these controls were not created still. Such behavior should not be seen as an erroneous one: these processes will be created when an audio playback is first started or when audio bouncing is initiated, and so the visual controls will also become active then.

Q. I may hear that when I switch from the 44.1 kHz working (project) sample rate to the 96 kHz sample rate the overall plug-in's sound changes a bit. Is this normal?

A. This is an expected plug-in behavior since the working sample rate may affect several internal signal processing aspects, especially digital filter shapes. Thus we are making an expressed disclaimer about the plug-in's sonic qualities when working at different sample rates: the plug-in may produce sonically different results working at different sample rates. However, these differences are usually minor any should not be construed as an invalid plug-in operation. It should be also noted that Voxengo plug-ins are designed to work at any sample rates equal or above 44.1 kHz: lower sample rates are not supported.

Q. On multi processor systems are Voxengo plug-ins optimized to make use of more than one processor?

A. Voxengo plug-ins do not have special multi-processor optimizations, but they work fine if audio host application puts two or more instances on several cores/processors. This means that a single instance of a plug-in won't be spread over two cores/processors, but if you load two instances placed on different tracks/buses, they will be spread. This was tested in several hosts that support multi-processing, and the boost is around 70% in dual-core mode in comparison to single-core mode.

Q. Could you please tell me an easy way to setup mid/side processing over a stereo input?

A. To setup mid/side processing you will need to do the following: a) open the "Channel Routing" window; b) assign the mid/side pair "MS1" to channels "A" and "B"; c) assign channel "A" to group 1, channel "B" to group 2. After this has been done, the mid channel will be affected by settings defined in the group 1, and the side channel will be affected by settings defined in the group 2.

Q. How do I setup the mid/side processing?

A. This is done in the "Channel Routing" window. First of all, you should choose a pair of internal channels which you would like to process in mid/side mode. You should assign these channels to the same mid/side pair (e.g. "MS1"). Then you will need to assign these channels to separate channel groups as otherwise you won't be

able to separately control the processing applied to both. For example, you may choose to assign the “mid” channel to group 1, and the “side” channel to group 2. After this has been done, you will be able to set the separate processing parameters for the “mid” channel by selecting group 1 for editing. The “side” parameters are accessible by selecting group 2 for editing. Alternatively, you may simply load the “Mid-Side Stereo” routing preset that will setup the necessary routing for you.

Q. I cannot get the GUI scale function to work. Tried closing and reopening the GUI but this didn't have any effect.

A. For the change to take effect when changing the user interface scale, you'll need to reload your audio application completely.

Q. I do not seem to find the “register” icon for the plug-in.

A. Please open plug-in's “Settings” window – you'll see the “Enter Key” button there.

Q. If I save a preset using host application's facilities – the name of the instance (which I can set in the text field of the plug-in) gets saved as well. So if I change presets the instance name gets overwritten.

A. Unfortunately, when you are using host application's own preset management this overwrites instance name when you load any previously saved preset. This cannot be resolved as otherwise instance name will not be saved with the project at all. You should use plug-in's internal preset manager if you want to keep instance name unchanged when loading a preset.

Q. If I click on the frequency, say “2K” and enter a “3” then it will jump to “20.00” and never come back.

A. Please make sure you are entering “3K” instead of just “3”, because entering “3” means “3 Hertz” – not “3 kilo Hertz”.

Q. It seems that the plug-in is taking some CPU resources even if I disable all of its stages (modules).

A. This is to be expected since even if plug-in processes nothing it still performs some routine tasks that are considered to be plug-in's processing overhead. Plug-in may also take additional CPU resources if its user interface is open since level meters (and spectrums) have to be updated even in the absence of the signal. Also the oversampling (quality) setting higher than “1x” alone makes plug-in take more CPU resources.

Q. Why the output signal metering of the plug-in looks inconsistent with the levels I see on my host audio application's level meters?

A. This can be caused by two reasons: 1) If you are using mid-side processing the output level you see is actually mid or side channel level, it is not a final output level. This could not be implemented in a different manner because if plug-in showed only the final output signal level you would not be able to see the signal level of the mid or

side channel separately. 2) If you are using the oversampling setting different from the “1x”, the signal level you see is the output signal level before the final downsampling is performed. Since after downsampling new peak overshoots may appear you may see these peaks on the host’s level meter while not on the plug-in’s one. However, these overshoots usually have a low magnitude and should not cause any problems. If in doubt, please use a peak limiter on the master bus.

Q. I can only see one stereo out on the plug-in, that's not right, is it?

A. This means that your audio host application does not support VST 2.3 multi-channel operation, or you are inserting this plug-in to a mono or stereo track/bus. Many audio hosts still limit plug-in support to 2 channels input/output only, however this is not plug-in’s limitation as it was successfully tested in hosts that support multi-channel operation.

Q. It seems that plug-in’s Routing window may display up to 16 input channels while there are only 8 internal channels inputs can be routed to. This looks like a limitation, is it so?

A. Indeed, plug-in may accept up to 16 input channels (8 normal inputs + 8 side-chain inputs) while it only possesses 8 internal channels. This means one cannot use full multi-channel external side-chain signal to process multi-channel input signal. Since side-chaining is usually used for mono and stereo signals only, this limitation should not cause much inconvenience in real situations, not to mention that dealing with multi-channel external side-chain signal can be mentally stressful.

Q. It would be nice if there was a quicker way to shuffle through the popup list options (in the Routing window, various mode and preset selectors).

A. You may hold the left mouse button on the selector, and then release the mouse button on the required popup list option. It is a one-click way of selection.

Q. It seems that gain reduction meter’s scale is small – down to -9 dB only. Does it mean a particular Voxengo plug-in is not designed for higher compression values?

A. Since Voxengo plug-in’s gain reduction meter shows gain reduction value relative to an average gain reduction in 2 second time frame, 3 to -9 dB scale is enough for most situations ranging from soft to heavy compression. In this sense, gain reduction meter shows instant compression ratio changes; that is why it can also display positive values showing moments when compression is relaxed after a period of a relatively high compressive action.

Q. Why gain reduction meter may show positive values?

A. The reason gain reduction meter may show positive values is due to the fact it shows gain reduction changes relative to 2-second average gain reduction, that is why it may fluctuate to positive values from time to time: this moments show when gain reduction becomes lower than the average gain reduction.

Q. I'm trying to set up side-chaining in Cubase 4 but I'm not sure if I doing it correctly. I have 1 track sending to the left input of a plug-in and another track sending to the right of a plug-in. I have setup this plug-in on a stereo group. I then go into plug-in's routing and Mute the B side-chain signal as I do not want to hear it but it does not seem to be effecting the compression.

A. In this scenario beside muting the "B side-chain signal" you also need to set key signal source for channel "A": it should be the "B" you have muted.

Q. I can't figure out how to setup side-chain in a Voxengo plug-in.

A. If you are sending stereo side-chain signal to a plug-in, in this plug-in you have to assign Input channels 3&4 to internal channels C&D. Then you need to assign Key Signal Sources for internal channels A&B: you should use C&D (C to A, D to B). Then you'll need to mute channels C&D, and you are done: channels A&B will be affected by signal in channels C&D (which represent side-chain inputs).

Q. When loading a Voxengo plug-in into my audio host application, and opening plug-in's "Channel Routing" window, I see a single input and a single output channel only while the plug-in is inserted into a multi-channel track. What should I do?

A. Usually you simply need to press the "Play" in your host audio application once so that this application reports an actual number of input and output channels to the plug-in.

Q. Hello, when using the equalizer I am able to drag the control points around however there is no effect on the audio.

A. You'll have to right-click the control point and select its type which by default is "Off" (filtering is disabled).

Q. Will upgrading to the latest plug-in version affect the settings for songs already in progress where previous version is being used?

A. If you are applying a minor version update (e.g. 2.0 to 2.1, 3.0 to 3.2, etc.) this will not affect settings, and the update can be considered "transparent". Major version upgrades will not interfere with your existing song settings since major version upgrades come with a different module name and they can be used side-by-side with a previous major version of the plug-in. However, next major version will not automatically "inherit" settings and presets you created in the previous plug-in version as, technically speaking, such next major version should be perceived as a brand new plug-in. Next major version plug-in may also feature reworked sound algorithms that may not sound exactly like a previous version at similar settings.

Q. What is oversampling, and should I use "2x", "4x" or "8x" setting?

A. Oversampling allows a plug-in to run at a higher internal sample rate: this usually produces a finer sound. For 44.1kHz/48kHz audio using "4x" setting is suggested. For 88.2kHz/96kHz audio "2x" is usually enough. You do not need oversampling at

higher project sample rates. The “8x” setting is an extreme setting which is available when you need an extra fine sound, especially when plug-in settings are specifically tuned to produce strong harmonic distortion.

Q. When I enable oversampling in Voxengo plug-in and route its output to the same, but unprocessed track in order to sum two signals (processed and unprocessed) in a desired proportion I get weird phasing effect.

A. This is to be expected since oversampling in most Voxengo plug-ins is not linear-phase. Thus, it produces phase shift making it impossible to sum processed and unprocessed tracks. This can be fixed by inserting the same plug-in to unprocessed track and enabling oversampling there, but keeping other aspects of the plug-in at minimal effect setting.

Q. Is there anyway to change the default setting, the way the plug-in loads? I would like to set some specific parameters every time I first launch the plug-in.

A. Please use the “Set as Default” function in plug-in's own Preset Manager to set a new default plug-in state.

Q. I have tried to use a particular Voxengo plug-in on a track in parallel to an unprocessed copy of the same track. I have got some phase cancellation effects when summing these tracks. What is going wrong?

A. Voxengo plug-ins, depending on the algorithm, may produce some amount of phase shift across the frequency spectrum. This is especially true when oversampling was set to any value above “1x” (if a plug-in does not implement linear-phase oversampling). While this should be considered a normal behavior, it may stop you from using the “parallel compression” technique and the likes. To avoid this problem, many Voxengo plug-ins feature the “Dry Mix” parameter which implements a built-in dry/wet signal mix control. If you still need to implement a dry/wet mix control by means of two tracks, you have to insert an instance of the same plug-in to the “dry” (unprocessed) track and use neutral settings in that instance together with the same oversampling setting.

Q. I’m testing Voxengo plug-in for “bit transparency”, and it seems that even in the “Bypass” mode most plug-ins produce output with 32-bit floating point resolution from 16-bit input signal. I was expecting them to produce original 16-bit signal.

A. This is to be expected since the “Bypass” mode in Voxengo plug-ins should be considered temporary-use mode. Even when in the “Bypass” mode Voxengo plug-ins still add noise that prevents “denormal number” problem from happening. Since this noise is very quiet, it increases technical bit depth of the signal. You may disable the “Anti-Denormal Noise” option in the global settings to make plug-ins behave “bit-transparently” in the “Bypass” mode.

Q. How can I get -48 (-96) dB per octave low- and high-pass filters?

A. All you need to do is add as many -24 dB per octave filters as necessary, all tuned to the same center frequency. For additional convenience, you may put each filter to its own Gain value so that you can have a single-click access to every filter. Do not forget that you can also select all filters to move them together.

Q. I'm not sure whether it is a bug or not but should the Underlay spectrum follow the same Block size as the Main spectrum? Currently, no matter which Block size I set, the Underlay spectrum block always stays at 2048.

A. You have to copy spectrum mode settings from the current group to the underlay group – this can be done by opening spectrum mode editor and using the “Copy” function there.

Q. I'm now evaluating one of the Voxengo equalizers, and I was hitting upon some strange behavior when using it. I used two “Peaking” filters with a Gain setting of about -3dB and a rather wide bandwidth. What I've got is that with the equalizer “on”, the track's peak level was higher than with the equalizer “off”. I do not get this.

A. All filters usually produce additional oscillation (known as Gibbs phenomenon), thus increasing momentary peak levels while the average power of the signal may become lower. This fact should not be considered as flaw, it is how mathematical side of signal processing “works”.

Happy sound processing!