Eventide
the next step

SP2016 EFFECTS PROCESSOR/REVERB
OPTIONAL PROGRAM SOFTWARE

The SP2016 CHANNEL VOCODER

Factory Program Number 47

CHANNEL VOCODER is a 20-channel digital vocoder program that can impress the frequency spectrum of the signal in the left input of the SP2016 (called the ANALYSIS signal) onto the signal introduced into the right input (the SYNTHESIS signal). The program incorporates many parameters that allow precise control over the vocoding process, including:

* Adjustable frequency ranges for the analysis and synthesis filter banks
* Analysis/synthesis filter patching control
* Presence and articulation control
* Adjustable amplitude compression on the analysis input
* Variable addition of unvoiced components (eg, sibilance and fricatives)
* Adjustable mix of the synthesis signal for 'fill' (silence bridging)
* Stereo/mono switch, adjustable delay on even/odd harmonics, band delay

The program incorporates a bank of 19 filter channels for the analysis signal and another bank of 19 for the synthesis signal. A 20th channel on the analysis side is used for 'unvoiced energy' detection. The 19 filters in each bank are spaced at harmonic intervals, and both banks may be shifted up and down to cover various portions of the audible frequency spectrum. In addition, the banks may be shifted relative to each other to effect various kinds of spectral dislocation.

Due the multiplicity of functions available, the 13 parameters in CHANNEL VOCODER are organized into 3 distinct 'pages' of options:

ANALYSIS OPTIONS:

PRESENCE	UNVOICED MIXCOMPRESSION
ARTICULATION	FILL SIGNAL

FILTER OPTIONS:

BASE FREQUENCY	SYNTHESIS OFFSET
INVERT/ROTATE PATCH	

OUTPUT OPTIONS:

BANDELAY CHANNELS/BAND	EVEN/ODD HARMONIC DELAY
BANDELAY MAXDELAY	STEREO/MONO OUTPUT

Only one parameter page is active at any one time, and switching between pages is accomplished by pressing the SOFTKEY. The legend appearing after the SOFTKEY is hit will indicate the currently active page - pressing repeatedly allows you to cycle through the 3 pages. When the desired page is selected, you can access the parameters in the page by pressing the PARAMETER key. A fourth SOFTKEY, "ALL OPTIONS", can be enabled and executed if it's desired to defeat the page organization of the parameters and display them in a single long list.

PLEASE NOTE: for ease of use, a selection of 3 parameters, Base Frequency, Presence, and Stereo/Mono (i.e., one parameter from each of the 3 pages) is initially enabled as a 'starter set' when the program or a preset using it is first loaded. This makes it simple for new users to have some control over the program if they're not familiar with the page organization of all the parameters. After the first time the softkey is pressed the parameters will appear only in their respective pages.

ANALYSIS PARAMETERS

COMPRESSION MAX/HIGH/MED/LOW/OFF DEFAULT: MEDIUM

This parameter controls the degree of compression applied to the analysis input signal, from 'maximum' to 'low' and off. Combined with the compression adjustment is a gain boost for higher amounts of compression; the overall effect is that low level signals are boosted when greater degrees of compression are applied.

Compression is most useful for analysis signals that have a wide or uncontrollable dynamic range - in these cases the level signal applied to the analysis input should be set so that the peaks are just lighting the top segment of the input level bargraph. When using analysis signals which are excessively noisy, the compressor should be backed off so that the noise signals don't cause modulation of the synthesis signal, heard as a low-level 'burbling'. (The FILL SIGNAL parameter can often be used to mask such burbling, if necessary).

PRESENCE -5 TO +20 DEFAULT: 0

The PRESENCE parameter controls the amount of boost in the high-mid frequency range of both the analysis and synthesis signals. The affected frequencies will depend on the setting of the BASE FREQUENCY parameter - in general, the higher the base frequency (that is, higher the range of the band-pass filters), the higher the frequencies that will be boosted.

Adding extra presence can greatly increase the intelligibility of synthesized speech sounds by amplifying those portions of the spectrum which are critical for vowel and voiced consonant formation. However, adding too much presence can cause an annoying harshness and cause the SP2016 to clip internally on signals which already have their primary energy in the high-mid band. A good rule of thumb is to use only the minimum amount of presence boost needed to give satisfactory results.

ARTICULATION -5 TO +5 DEFAULT: 0

This parameter controls the degree of apparent 'articulation' of the synthesized voice. More negative values will cause smoother and more gentle modulation of the synthesis signal by the analysis input; higher positive values cause the progressively 'tighter' tracking of the analysis signal. Articulation control is accomplished by changing the time constants of the envelope followers at the outputs of each of the 19 band-pass analysis filters.

High articulation values will cause the synthesis signal to follow the analysis signal more closely, at a setting of +5 producing a quite 'gravelly' or grainy effect. At negative settings, a slow and very smooth modulation can be achieved whose output, while less distinctly intelligible as speech than at higher settings, can be very useful for adding the general formants of vowels to sustained instrumental sounds. Much of the loss of intelligibility at lower articulation values can be made up for by mixing in unvoiced sibillance and fricative components. (See the UNVOICED MIX parameter. Articulation control can also be very effective when used in conjunction with the BANDELAY parameters)

FILL SIGNAL 0% TO 99% DEFAULT: 0%

The FILL SIGNAL parameter controls what percentage of un-vocoded synthesis signal to mix in during those intervals when no analysis signal is present. (This capability is known as 'silence bridging' in some vocoders.) An envelope signal derived from the input compressor (independent of the current COMPRESSION parameter value) is inverted and used to smoothly gate in a low-passed version of the synthesis signal when the analysis signal falls to a low amplitude or is absent.

This parameter is useful any time that it's not desirable for the synthesis signal to just 'disappear' when there is no analysis signal present, or for situations where only a certain portion of a musical passage is to receive a vocoded treatment. The fill signal can also be useful for masking a multitide of irregularities in the analysis signal, for instance, when the analysis signal contains a large unwanted noise component. Adding some fill signal can also help ameliorate the gravel-voiced effect produced when the ARTICULATION parameter is at its maximum.

UNVOICED MIX 0% TO 99% DEFAULT: 0%

This parameter controls the level of synthetic unvoiced speech components to add to the synthesis signal. Unvoiced speech components (i.e., hard consonant sounds like 'k', t', and 'd' and sibilant sounds like 'ssss' 'sh', and 'ch') can be loosely characterized as having a preponderance of their energy in high frequency components. A sharp high-pass filter in the vocoder (the 20th 'channel') splits off likely-looking unvoiced components from the analysis signal and uses their envelope to modulate the amplitude of an internal high-passed noise source. The resulting re-synthesized unvoiced components can be mixed into the vocoded signal in varying degrees using this parameter.

Adding unvoiced components to the output signal can improve the intelligibility of vocoded speech to a surprising degree. While it is not always appropriate for musical effects, very convincing and understandable synthetic voice sounds can be obtained. In particular, when the ARTICULATION parameter is set to a low value, the addition of some unvoiced components can make what were previously indistinct vowel sounds quite intelligible as speech.

FILTER PARAMETERS

BASE FREQUENCY 100 TO 420 HZ DEFAULT: 160 HZ

The basic structure of a channel vocoder consists of 2 banks of filters, one bank analyzing the analysis input signal, and the other filtering the synthesis output signal. The BASE FREQUENCY parameter can be used to specify the frequency of the lowest filter channel in the analysis bank. In the SP2016 CHANNEL VOCODER, the frequencies of the filters within each bank are spaced at harmonic intervals, so the BASE FREQUENCY tunes both the 'fundamental' frequency of the filter bank and the accompanying 18 'harmonics'. Normally, corresponding analysis and synthesis filters are positioned at the same frequency (see the SYNTHESIS OFFSET parameter for exceptions to this) so BASE FREQUENCY effectively specifies the entire range of frequencies over which analysis and synthesis will occur. The following values are the fundamental frequencies of the 35 ranges available within the CHANNEL VOCODER:

100HZ	120HZ	144HZ	184HZ	224HZ	264HZ	344HZ
104HZ	124HZ	152HZ	192HZ	232HZ	280HZ	360HZ
108HZ	128HZ	160HZ	200HZ	240HZ	296HZ	372HZ
112HZ	132HZ	168HZ	208HZ	248HZ	312HZ	388HZ
116HZ	136HZ	176HZ	216HZ	256HZ	328HZ	400HZ

BASE FREQUENCY is useful for tuning the analysis range of the vocoder to best match the spectrum of the analysis signal. Vocoding from a female soprano voice might require a base frequency around 300 HZ, for example, while analyzing a male bass voice might call for a base frequency closer to 100 HZ. Experimentation with the base frequency is to be encouraged, since seemingly slight adjustments in frequency (which often may seem not to be 'correct', from an acoustical point of view) can have a powerful influence on the overall vocoding effect.

Theoretically, an exact match between the BASE FREQUENCY value and the fundamental pitch of the analysis signal will allow each filter in the analysis bank to capture one harmonic of the signal and thus yield the most realistic and intelligible vocoding effect. In practice, however, BASE FREQUENCY will put the filter bank into the right spectral 'ballpark' so that it best matches the (inevitably wandering) fundamental frequency of the analysis signal.

SYNTHESIS OFFSET -34 TO +34 DEFAULT: +0

As discussed previously, corresponding filters in the analysis and synthesis banks are normally positioned at the same frequency. However, the SYNTHESIS OFFSET parameter can be used to shift the base frequency of the synthesis bank (and thus all its 'harmonics') away from the base frequency of the analysis bank. Changing the synthesis range relative to the analysis range in this way results in a 'formant shift': the overall shape of the spectrum remains the same but is shifted up or down in frequency. SYNTHESIS OFFSET is adjusted using a parameter value that specifies the number of frequency ranges over which the synthesis bank should be shifted relative to the analysis bank.

For example (see the table of base frequencies above), if the analysis base

frequency is currently at 160HZ, a SYNTHESIS OFFSET of -1 will set the synthesis base frequency to 152HZ (i.e., the next lower range). Likewise, an offset of +10 would set the synthesis frequency to 240HZ (a shift of 10 ranges up in frequency). It's important to note that you may not shift the synthesis offset past the highest and lowest frequencies shown in the table; the value of the parameter will limit itself to prevent this. For example, if the BASE FREQUENCY is set to 104HZ, the lowest value attainable with SYNTHESIS OFFSET is -1, since the base frequency of the synthesis bank may not be shifted any lower than 100HZ.

The SYNTHESIS OFFSET parameter is useful for shifting the synthesis filters to match a synthesis signal whose fundamental frequency is not close to that of the analysis signal, for instance, when using a low male voice to vocode the higher octaves of a keyboard instrument. The formant shift resulting from a synthesis offset can be used to impart a strange, Harmonizer-type sound to vocoded signals (the 'chipmunk effect'), or it can actually be used to UN-chipmunk a Harmonized signal: put the Harmonized signal into the synthesis input and the un-Harmonized signal into analysis input, then set the SYNTHESIS OFFSET inversely to the pitch ratio (i.e., shift the formant in the opposite direction of the pitch shift). The (albeit rather strange) result more or less preserves the formant of your un-pitch-shifted signal while pitch shifting the actual frequency of it.

INVERT/ROTATE PATCH -9 TO +9 DEFAULT: +0

The PATCH parameter is used to alter the normal mapping of analysis to synthesis filters in the vocoder. Ordinarily, a software 'patchcord' connects analysis filter channel 1 to synthesis filter channel 1, analysis channel 2 to synthesis channel 2, and so on. The PATCH parameter allows the vocoder to be 'repatched' by swapping higher and lower-numbered patchcords (INVERT MODE) or by rotating the patchcord positions left or right (ROTATE mode). In the INVERT mode, the result is a spectral inversion, where the higher-order harmonics of the analysis signal control the lower-ordered harmonics of the synthesis signal and vice-versa. In the ROTATE mode the result is a harmonic transposition, with any shifted-off channels at top wrapped around to the bottom, and vice versa.

In INVERT mode, the PATCH parameter controls the 'degree' of spectrum inversion - positive values cause channels to be swapped starting at the outside positions (i.e., highest and lowest channels), while negative values swap starting at the center (i.e., around channel number 9). For instance, a value of +3 causes channels 19 through 17 to be swapped with channels 1 through 3, respectively. Conversely, a value of -1 causes channel 10 to be swapped with channel 8, and so on. PATCH INVERT values of +9 and -9 are actully redundant, with both values producing a complete inversion of the spectrum, each higher-ordered channel swapped with its lower-ordered counterpart.

In ROTATE mode, the value of the PATCH parameter controls the number of patch positions to rotate the filter connections. Any patches rotated past the highest channel (19) will be rotated back through the lowest channel (1) and vice versa. For instance, a value of +3 causes each channel to be patched 3 channels higher (with the three highest channels patched into the 3 lowest channels); a value of -2 causes each patch to be shifted 2 channels lower (with a corresponding wrap-around at the top). Patch values of +9 and minus 9 again have a redundant effect, swapping the top and bottom halves of the spectrum, while preserving the harmonic ordering within each half.

The PATCH parameter is useful for situations where some dynamic spectrum

modulation is called for but a 'standard' vocoded sound is undesirable. For instance, the use of a judicious degree of spectral inversion can turn standard English into a very convincing alien language! Bizarre accents and strange inflections of human speech can also be created by shifting vowel formants around using the ROTATE mode. Though obviously not a parameter that will come into everyday use, PATCHing may find uses in those situations where radical spectral modification of a sound is necessary.

OUTPUT OPTIONS

BANDELAY [CHANNELS/BAND] 1 TO 9 DEFAULT: 1
BANDELAY [MAXDELAY] 0MS TO +/-740MS DEFAULT: 0 MS

The two parameters grouped under the heading of BANDELAY allow you to apply a variable delay between the envelope output of the analysis filter bank and the control input of the synthesis filter bank. Channel envelopes are grouped into 'bands' and the bands are passed through delays whose lengths increase or decrease with successive bands. (Note that no AUDIO signals are being delayed here - just the envelope information passed between the analysis and synthesis banks.)

The CHANNELS/BAND parameter allows some number of analysis channels from 1 through 9 to be grouped into a single band. When this parameter is selected, brackets on the display will move to enclose the value representing the number of channels per band. A different delay will be applied to the envelope signals corresponding to each band, depending on the MAXDELAY parameter. At the minimum of CHANNELS/BAND = 1 there will be a maximum of nineteen bands; at the maximum of CHANNELS/BAND = 9 a minimum of two bands will result. In general the 19 filter channels are divided among the number of specified bands with any remainder left after the division going into the highest frequency band.

The MAXDELAY parameter specifies the delay that will be applied to the maximally delayed band, from 0 to 740 milliseconds. When the MAXDELAY parameter is selected, the brackets in the display shift to indicate the delay value in milliseconds. Delays for each of the bands of grouped channels are distributed across this maximum length, with each successive band delayed by a fraction of the total. For instance, with MAXDELAY = 600 milliseconds and CHANNELS/BAND = 6, the result is 3 bands: Band 1 (filter channels 1-6) delayed by 0 ms, Band 2 (channels 7-12) delayed by 300 ms, and Band 3 (channels 13-19) delayed by 600 ms. For CHANNELS/BAND = 1 and MAXDELAY = 740, we get 19 bands whose delay lengths increase by about 40 milliseconds for each band, from 0 ms for channel 1 to 740 ms for channel 19. With positive valued delays, increasing channel numbers are delayed by increasing amounts and the highest frequency band is maximally delayed (the delay 'sweeps up'), but with negative valued delays, increasing channel numbers are delayed by DEcreasing amounts and the LOWEST frequency band is maximally delayed (the delay 'sweeps down').

The effect of the BANDELAY parameters on the synthesis signal can be a radical departure from the 'standard' vocoded sound: anything from a doubled, echo-y spectral ghost to a shimmering, sweeping harmonic filter, ambiguously though definitely related to the analysis source signal. At short delays the normal vocoding effect is still evident, though smeared and elongated, but at long delays the rippling effect of delaying each synthesis harmonic's excitation takes over and the identities of both the synthesis and analysis signals. disappear.

EVEN/ODD HARMONIC DELAY 0 TO 380MS DEFAULT: 0MS

Although they may appear functionally similar to the previously discussed parameters, the HARMONIC DELAY parameters actually add delay to the output audio signals of the synthesis filters (rather than adding delay to the envelope signals passed between the analysis and synthesis banks). Different delay lengths may be adjusted for the outputs of the even and odd harmonic filters, from 0 to 380 milliseconds in 60 settings. The delay for the currently selected parameter will be indicated on the display with an arrow pointing at the value. Though not completely silent, delay changes have been designed to be as glitch-free as possible.

The effect of the ODD/EVEN HARMONIC delay parameter is equivalent to a simple post-delay on the signal when the two delays are set equal to each other. When a delay offset is introduced between them the sound first fattens, and then breaks apart into 2 distinct vocoded signals, each with its own character. The doubling effect is especially dramatic in STEREO MODE (see below).

MONO/STEREO OUTPUT DEFAULT: MONO

The MONO/STEREO parameter switches the output of the vocoder between stereo and mono signals. In STEREO mode, the even harmonics are sent to channel 1 (the left output), while the odd harmonics go to channel 2 (right output). In MONO mode the odd and even harmonics are simply mixed together in equal proportions.

The stereo vocoded effect can be quite dramatic, especially with a small amount of delay added (say, less than 80 ms). As various speech sounds come from the analysis signal, the proportion of odd and even harmonics is constantly changing, producing a spatially as well as spectrally dynamic effect as the sound shifts its focus back and forth between the two channels.

CHANNEL VOCODER - QUICK PARAMETER REFERENCE:

ANALYSIS OPTIONS:

PRESENCE - Presence (hi-mid) frequency boost for added intelligibility twenty five stages from -5 TO +20,

UNVOICED MIX - Variable addition of unvoiced sounds (ie, sibilance and fricatives) to synthesis signal, 0% TO 99%

COMPRESSION - Analysis input compressor, 5 settings: MAX/HI/MED/LOW/OFF

ARTICULATION - Analysis articulation tracking control, from 'loose' to 'tight', -5 TO +5

FILL SIGNAL - Variable mix of unvocoded synthesis signal for silence bridging, 0-99%

FILTER OPTIONS:

BASE FREQUENCY- Adjust filter ranges by specifying 'fundamental' frequency of analysis and synthesis filter banks, 100 TO 400HZ

SYNTH OFFSET - Shift synthesis bank frequency range relative to analysis bank; maximum possible shift = -34 TO +34

INVERT/ROTATE PATCH - Alter normal mapping of analysis/synthesis filters by spectral inversion or rotation (harmonic transposition), -9 TO +9

OUTPUT OPTIONS:

BANDELAY CHANNELS/BAND - Filter envelope delay -specify number of filters per delay band, 1 TO 9

BANDELAY MAXDELAY - Specify maximum delay for last band, 0 to 740 MS

EVEN/ODD DELAY - Delay output signals from synthesis filters (separate for even and odd harmonics), 0 TO 380 MS

STEREO/MONO - Switch output between STEREO (even harmonics=left channel, odd harmonics=right channel) or MONO (mixed)

Eventide

ONE ALSAN WAY
LITTLE FERRY, NEW JERSEY 07643
201-641-1200 TWX:710-991-8715