Introduction

This program implements a filter bank channel vocoder with 18 channels. The program was designed for use with a full duplex Creative Labs SoundBlaster card (e.g. the SB32 or SB64) and a Pentium computer running Windows95. The program should function, however, with any sound card that can operate in full duplex mode and has a Windows multimedia driver. The program runs fine on a 120MHz Pentium but may produce choppy sound with slower computers.

What is a "Channel Vocoder", you ask? A channel vocoder is a device for compressing, or encoding, the data needed to represent a speech waveform, while still retaining the intelligibility of the original waveform. The first channel vocoder was developed by Homer Dudley in 1936. It passed the speech signal through a bank of band-pass filters. These filters each covered a portion of the audio spectrum. The energy of each filter's output was then measured, or sampled, at regular time intervals and stored. This collection of filter energy samples then comprised the "coding" of the speech signal. This code could then be transmitted over a communication channel of lower bandwidth than would be neccessary for the raw speech signal. At the receiving end, the speech signal is reconstructed from this code by using the time sequence of filter energy samples to modulate the amplitude of a pulse signal being fed into a bank of filters similar to the ones used for the encoding. The result, while clearly not the original speech signal, is nonetheless intelligible, and one can, in most cases, understand what is being said.



A block diagram of the overall process is shown in the figure below.



SYNTHESIZER SECTION

The idea for the channel vocoder technique arises from the manner in which speech is generated in the human vocal tract. Simply put, the vocal chords produce a periodic, pulse-like, stream of air, which is then acoustically filtered by the elements of the vocal tract: the esophagus, tongue, lips, teeth, and the oral and nasal cavities. As one speaks different sounds, the shape and elastic properties of these elements are being constantly changed in response to neural signals arising from speech centres in the brain. This causes a time-varying filtering, or spectral variation, of the excitation arising from the vocal chords.

The channel vocoder, then, first *analyzes* the speech signal to estimate this time-varying spectral variation. To do this it uses the filters in the filter banks to determine how a particular frequency component of the speech signal is changing with time. On the output end, this analysis of the spectral variations are used to *synthesize* the speech signal by using another filter bank to apply these same spectral variations to an artificial periodic pulse like signal. The output filter bank acts as an artificial vocal track and the pulse signal acts as a set of artificial vocal chords.

Some sounds produced during speech do not arise from the vocal chords, but are produced by turbulent air flow near constrictions in the vocal tract such as may occur between the tongue and the teeth. For example, such sounds as "SSS", "K", "SSHH", "P", and so forth arise in this manner. These sounds would be poorly reconstructed using a pulse excitation source, and so most channel vocoders also have a noise signal that can be used as an excitation source as well. A "Voiced/Unvoiced" detector circuit is used to detect whether the speech signal is arising from vocal chord excitation (Voiced speech) or is arising from noise excitation (Unvoiced speech), and the appropriate excitation source is then selected at the output end.

Channel vocoders were originally developed for signal coding purposes, with an eye (ear?) towards reducing the amount of data that would be needed to be transmitted over communication channels. In fact, speech coding system development continues to this day to be a vigorous area of research and development. These systems have far outstripped the basic channel vocoder idea in complexity, coding efficiency, and intelligibility, however. So why have we produced this program? The reason is that channel vocoders (and the functionally equivalent, but computationally quite different, phase vocoder) have found application to music production. In the 1960's Siemens in Germany produced a vocoder which was used in some recording and in radio and television. The vocoders used in these early musical efforts were very large and unsuited to general use. In the mid-70's a breakthrough of sorts came about when a number of companies, notably EMS (Electronic Music Studios) in England, produced relatively small and easy to use vocoders designed for use in musical applications. After that, the vocoder sound became a staple of the music and entertainment industry. Many extremely popular records (Kraftwerk!), TV shows (the Cylons of Battlestar-Galactica), and movies (Darth Vader in Star Wars) are identified with vocoders.

Although the introduction of these relatively small vocoding systems made it possible for the wide application of vocoders to music and film, they were still quite expensive for your average musician in the street. The EMS vocoder cost upwards of 6500 UK pounds! There are now less expensive vocoders on the market now, but these attain their low price at the expense of limited functionality. This is where the Cylonix vocoder comes in. For a price of 1/200 that of the EMS vocoders you can get a similar level of functionality! The Cylonix vocoder has the following sets of features, some of which are to be found only on the most expensive of the hardware vocoders.

FEATURES OF THE CYLONIX VOCODER

Number of Channels: 18 Channel Filter Cutoff Rate: 48 dB/octave Choice of 3 different filter bandwidths Sibilance Feedthrough Simple Pitch Tracking Spectrum Shifting Spectrum Sampling and Hold Adjustable Analyzer/Synthesizer Spectrum Mapping Matrix Graphical Spectrum Energy Display Stereo Output with Independent Control of 8 Different Parameters on all 18 Synthesizer Channels Level Pan Distortion Echo level and Echo Period Delay Analyzer Slew Rate Compression/Expansion Voice/Unvoiced Detector Internal Variable Frequency and Pulse Width Oscillator and Noise Generator Recording of Vocoder Output Loading of WAV Files for use as Carrier and Modulator Sources Randomization of Spectrum Mapping Spectrum Enhancement and Smoothing MIDI Control of All Sliders and Buttons Separate Carrier and Modulator Gain Controls

Some Other Commercially Available Vocoders

Here are specifications of some commercially available hardware vocoders. Compare the capabilities of these vocoders to those of the Cylonix vocoder. You will get a feeling for the power of the Cylonix program. Please keep in mind that the great majority of these commercial machines cost hundreds or even thousands of dollars! Compare this to the low price of the Cylonix vocoder.

EMS 5000:

Analog Price: around 6500 UK pounds Number of Channels: 22 Filter Order: 8th order Channel levels can be adjusted with external Control Voltage (CV) inputs Pitch extractor 2 built in oscillators (VCOs) with square wave and ramp outputs One noise source Voiced/unvoiced detector Slew/Freeze control can be manually operated with a variable control or freeze switch. Frequency shifter stage with a range of .05Hz to 1 kHz Patching of analyzer and synthesizer channels with a matrix of 22 x 22 patchpoints.

EMS-2000

Analog Price: 1000 UK pounds Number of channels: 16 Filter Order: 6th order Internal pulse oscillator and white noise source Voiced/Unvoiced detector Spectrum Slew/Freeze (Hold)

SYNTON 221

Analog Price: around \$7500 US Number of channels: 20 Filter cutoff: 54 db/octave Channel patching Matrix-panel Internal VCO noise generator Voiced/Unvoiced detection LEDs for spectrum monitoring External control inputs/outputs for each channel

SYNTON 222

Analog Price: \$625 US Number of channels: 12.

SYNTON 202

Analog Price: \$500 US Number of channels: 10 No noise generator.

SYNTON SPX-216

Analog Price: \$1000 US Number of channels: 14 VCO and noise generator Carrier compression & distortion Rear-connector for CV in/out.

ELECTRO-HARMONIX VOCODER

Analog Price 400 UK pounds 14 channels

ROLAND SE70

10 and 21channels in the same unit Digital Effects box

ROLAND SVC-350

Analog 11 channels Compressor on Mic input Spectrum Hold 24db/octave filter cutoff

KORG VC-10

Analog comes with keyboard Compressor on mic input 1 master oscillator and 1 oscillator per key 3 low frequency oscillators (LFO's) for vibrato 20 channels 24 dB/octave filter cutoff

BODE 7702

Analog 16 channels Analyzer-Synthesize channel patching Pulse (buzz) and Noise (hiss) generators Spectrum Sample and Hold

SENNHEISER VSM201

Analog Price: 10,000 UK pounds 20 channels 36 dB/octave filter cutoff Fixed 150Hz pulse generator Noise source Voiced/Unvoiced detector with filter

Döpfer A-129

Analog Price: about \$400 US for basic module Made in conjunction with Kraftwerk. 15 channels 24 dB/decade filter cutoff Modular, allowing complex processing by adding other Döpfer mdoules

ELEKTOR

Analog kit 10 channels 24 db/decade filter cutof Noise generatorf Voiced/Unvoiced detector Sibilance Engineered in conjunction with Synton.

PAIA

Analog kit Price: \$100 US 8 channels 12 dB/decade filter cutoff Sibilance Carrier distortion

Carrier/Modulator Source Selection

The carrier input to the analyzer section can come from a variety of sources. These sources are selected with the following five buttons:



If the Extrn button is on, the carrier input is taken from the sound card Analog-to-Digital converter (ADC). See the <u>Input Mode</u> section for details on how to determine the external input source.

If the Pulse button is on, an internally generated pulse wave is used as the carrier signal. This pulse signal has a duty cycle (ratio of on time to off time) of 1/4. The frequency of the pulse waveform is set by the FREQ slider. The duty cycle, or width, of the pulse waveform can be set with the WIDTH slider. A high value corresponds to a nearly 50% duty cycle and produces a hollow sound. A low value corresponds to a low duty cycle and produces a raspy buzzing sound.

If the Chor1 button is on, three internally generated pulse waves are mixed together to form the carrier signal. These three waves form a diminished chord (e.g. C-Eb-Gb). The frequency of the pulse waveforms are set by the FREQ slider, and their pulse width with the WIDTH slider.

If the Chor2 button is on, three internally generated pulse waves are mixed together to form the carrier signal. These three waves form a major chord (e.g. C-E-G). The frequency of the pulse waveforms are set by the FREQ slider, and their pulse width with the WIDTH slider.

If the Noise button is on an internally generated uniformly distributed white noise signal is used as the carrier signal.

Only one these 5 buttons can be on at one time.

It is also possible to read in a wav file and use it as either the carrier or modulator signal.



Press the Load button to select the wav files which will be read in and be available for use as the carrier and modulator signals. A file dialog box will popup first for the carrier wav. This will ask for the file which will be used to obtain the carrier wav. If "Cancel" is pressed then no carrier wave will be loaded. Once the carrier wav file has been determined another dialog box, shown in the figure below, will popup, *if the wav file was a stereo file*. This dialog box will be used to determine which channel, or combination of channels, will be used as the carrier source. There are 4 different choices here: Left, Right, Left+Right, Left-Right and None. If None is selected then no carrier wav will be loaded. For the other choices, the source of the carrier wav is determined as follows: If Left is selected, the carrier is taken from the Left channel of the input wav file. If Right is selected, the carrier is taken from the Right channels is used as the carrier. After the carrier source (divided by 2) of the left and right channels is used as the carrier. After the carrier source is selected the process is repeated for the modulator source.

Select the Modulator Source		×
Left	Right	
Left + Right	Left - Right	
Nor	ne	

Once these files have been read in, you can press PCar to play the carrier signal wav file or PMod to play the modulator signal wav file. If the RptC button is turned off, pressing the PCar button will play the wav file through once. If the RptC button is on, pressing the PCar button will play the wav file repeatedly. Turning the RptC button off again will terminate the carrier wav playback at the end of the current cycle. Similarly, if the RptM button is turned off, pressing the PMod button will play the wav file through once, while if the RptM button is on, pressing the PMod button will play the wav file repeatedly.

Double-clicking on the PMod or PCar buttons will play both the Carrier and Modulator wave files at the same time. This allows synchronization of the modulator and carrier.

Please note that the internal format used for modulator or carrier sources is 16 bit, monophonic, 11.025 Khz format. If the wav files do not have this format the wav data will converted to this internal format after being read in by the program.

One benefit of using wav files as carrier or modulator signal sources over using external inputs is that the wav files can be 16 bit signals, while the external inputs are only digitized to 8 bits. Thus higher fidelity and lower noise results can be obtained by using the wav files than with external inputs.

Voiced/Unvoiced Detection

When the V/UV bottom is pressed the vocoder runs a voiced/unvoiced detection process. This process determines whether the modulator signal is a voiced sound (such as a vowel, - e.g. "ee" "aah" "ooh", etc.) or an unvoiced sound (such as a plosive or consonant or sibilant, - e.g. "t" "k" "ss", etc.). When an unvoiced sound is detected by this process, the carrier signal (whatever its source) is replaced by a white noise signal. The reason for using this process is that most carrier signals (such as the internally generated pulse signals) do not have the broad spectral content that is found in unvoiced sounds and hence can not accurately imitate such sounds.

Frequency Modulation

The frequency of the internally generated pulse and chorus signals is modulated via a low frequency square wave when the Yodel button is pressed on. The frequency of this modulation is set by the RATE slider. This results in a kind of a "yodeling" effect, although it is nowhere near as good as the yodel effect heard in Kraftwerk's "Autobahn". Interesting effects can be obtained when this is used in conjunction with the Sample/Hold mode (obtained by pressing the S/H button).

Pitch Tracking

The internally generated pulse and chorus carrier waveforms result in sounds that are very "robotic". Part of the reason for this is that the frequency of these waveforms is constant. A more natural speech sound can be obtained if the frequency of the carrier waveforms could be made to track the pitch of the modulator speech signal.

Please note: accurate pitch tracking is a non-trivial process, however, and good low noise algorithms require a lot of computing power. The Cylonix vocoder uses a very simple process which doesn't track the modulator pitch very well, but provides an interesting sound.

Pressing the Track 💭 button on enables this tracking process.

The pitch tracker works best with slowly changing speech. Don't talk quickly! The tracker seems to sound best when a complex carrier is used (such as the chorus signals) and with a positive spectral shift of two or three places.

Filter Bandwidth

The nature of the sound produced by the channel vocoder is in large part determined by the filters used in the analysis and synthesis channels. In the Cylonix vocoder, three different sets of filters can be used. These are all 8th order Butterworth digital filters, but have different bandwidths.

Pressing the Nrrw we button causes a narrow band version of the analysis and synthesis filters to be used. These filters have half the bandwidth of the usual filters (i.e. they are 1/6th of an octave wide instead of 1/3rd of an octave wide). This produces a somewhat hollow and reverberant sound, due to the increased ripple in the time-domain impulse resonse of the filter.

Pressing the Wide button causes a wide band version of the analysis and synthesis filters to be used. These filters have twice the bandwidth of the usual filters (i.e. they are 2/3rds of an octave wide instead of 1/3rd of an octave wide). This produces a somewhat less distinct and muffled sound than wid\th the normal filter set. Pressing the Comb button, when the wide filters are selected will set the level of every other synthesis channel to zero, resulting in what is effectively a 9 channel vocoder.

Pressing the Norm button causes the normal bandwidth version of the analysis and synthesis filters to be used. These filters have a bandwidth of the 1/3rd of an octave.

The filters are all implemented as 8th order (48 dB/octave cutoff) Butterworth digital bandpass filters having a selectable bandwidth of either 1/6, 1/3 or 2/3 octave. The centre frequencies, (rounded to the nearest Hertz), of the filters are as follows:

108, 137, 172, 217, 273, 344, 434, 547, 689, 868, 1093, 1378, 1736, 2187, 2756, 3472, 4375, 5512

Vocoder afficionados will notice that, unlike most hardware vocoders, the lowest and highest frequency filters are bandpass filters rather than lowpass or highpass. On the low end this is because of the fact that the Creative Labs Soundblaster cards have a low frequency rolloff of around 50Hz, so that there wouldn't be any significant difference between using a lowpass or a bandpass filter. On the high end, the antia-aliasing filtering done by the sound card for a sampling rate of 11,025Khz (which is the rate used in the Cylonix program) enforces a high frequency cutoff of 5,512Khz, So, using a highpass filter would provide no significant difference to the sound than using a bandpass filter.

Modulator and Carrier Gain Controls

The MGain slider controls the level of the vocoder modulator signal. This should be set to as high a level as possible without causing distortion. The center stripe on the slider will glow red when the modulator signal is too high and causing distortion. Note that some microphones have much lower levels than others, and require a higher gain setting. Make sure to adjust the levels using the sound card mixer program before adjusting the vocoder MGain setting!

The CGain slider controls the level of the vocoder carrier signal. This should be set to as high a level as possible without causing distortion. The center stripe on the slider will glow red when the carrier signal is too high and causing distortion.

It is very important to adjust these gain settings to obtain a distortion free sound! Of course, if you like distortion, crank 'em up!

Individual Channel Controls

A detailed diagram of the processing involved in a single vocoder channel is shown in the figure below. There are 18 of these channels, each with band-pass filters (BPF) that are tuned to a different center frequency. In each of these channels there are a number of parameters that can be controlled independently of the other channels. This allows very fine sculpting of the vocoder sound.



As seen in the figure, there are nine different channel operations that are implemented: LEVEL, PAN, DISTORT, DELAY, ECHO, ECHO TIME, ECHO REGEN, SLEW, and WARP. The Carrier Gain and the Modulator Gain controls shown in the figure are not individual channel operations, but adjust the level of the modulator and carrier signals being fed into all of the channels.

Channel Operation Controls:

There are 18 sliders which control the individual channel operation effects for each filter channel. These sliders control only one channel operation at a time. Different channel operations are activated by pressing their respective buttons.

The **Slave** button toggles slave mode, in which moving one of the channel operation sliders causes the other sliders to move as well. If the slider is moved with the left mouse button held down, the other sliders move by the same amount as the slider being manually moved. This can be used to change a channel operation parameter, such as level, by the same amount for all filter channels. If the sldier is moved with the right mouse button being held down, hwoever, results in the other sliders being moved by an amount proportional to their distance away from the manually moved slider. This is useful for making ramps (if the moving slider is at the left or right end), or other curves. You should play around with this to get used to the wide range of patterns that you can easily make in this way. *Keep in mind that if you just want to change the value for a single filter channel, the Slave button should be turned off!*

There are four buttons which set the channel operation sliders into often used configurations. The first three buttons are presets whose function depend on the particular channel operation that is currently selected. There is always a Rand button which randomizes the settings of the channel operation sliders. This is useful for experimenting and fooling around, and when you haven't got a clue of what you are doing.

Level

The output level of each of the synthesizer section filters is modulated by the control signals generated by the analyzer section. This modulation, of course, changes over time as the modulation signal changes. The Cylonix vocoder, however, provides an additional control over the output level of the synthesizer section filters. This can be used to to static shaping of the output spectrum, for example to enhance or suppress a particular frequency.

The channel level slider is used to control the level of the channel independantly of the other channels.

The three preset buttons for the Level operation are "Zero", "Half" and "Comb".

The Zero button sets all channel levels to zero. This is useful when you just want a few of the channels to be heard. You just hit the Zero button and then manually set those few channels to the desired level.

The Half button sets all channels to half level. This is useful in resetting the vocoder to its normal operational mode.

The Comb button zeroes the level on every other channel. This provides a narrow band effect similar to that obtained from some inexpensive hardware vocoders that are on the market (and which shall remain nameless!).

<u>Pan</u>

The synthesizer section output is a stereophonic combination of the 18 synthesis filter outputs. Each of these outputs can be independently panned between the left and right output channels, using the Channel Pan sliders.

The three preset buttons for the Pan operation are "Mono", "Sprd" and "Split".

Pressing the Mono button sets all of the pan sliders to the middle of their range, resulting in a purely monophonic sound.

Pressing the Sprd button sets the even number channel pan sliders to the right and the odd number channel pan sliders to the left. This results in a very wide and open spatial sound.

The Split button sets the pan sliders of the lower 9 channels to the right and the upper 9 channels to the left. This provides a different sort of spatial sound than the Sprd button setting.

Distort

The distort control adjusts the amount of distortion that is applied to the output of the synthesis filter. The distortion is of the "overdrive" type, which is used in many guitar "fuzz" boxes. The effect of the distortion is to produce a harsher, more "metallic" sound.

One can apply the distortion to some channels, while leaving other channels unaffected. This can produce an effect similar to that of aural "exciters" which distort selected frequencies to add extra harmonic frequencies to a sound.

The three preset buttons for the Distort operation are "Zero", "Full" and "Comb".

Selecting the Zero button sets the distortion level on all channels to zero.

Selecting the Full buttons applies maximum distortion to all of the channels.

Pressing the Comb button results in full distortion on alternating channels, with no distortion on the other channels.

<u>Delay</u>

The output of each channel can be delayed in time. Normally, this delay would be something to be avoided, but setting the delay to be different in different vocoder channels can result in some interesting effects. For example, setting the delay to increase from the low frequency channels to the higher frequency channels results in a sound like a "sweep" in filter frequency from low to high.

The delay slider controls the delay time, from a delay of zero to a delay of 737 milliseconds in steps of 5.8 milliseconds.

The three preset buttons for the Delay operation are "Zero", "Up" and "Down".

Pressing the Zero button sets the delay time for each channel to zero.

Pressing the Up button sets the delay time to an upward ramp. That is, the delay for the lowest frequency

channel is zero and the delay increases linearly by a small amount with increasing channel frequency.

Likewise, pressing the Down button sets the delay time to a downward ramp. That is, the delay for the highest frequency channel is zero and the delay increases linearly by a small amount with decreasing channel frequency.

The Up and Down delay ramps produce a "swept" filter effect. These effects are most obvious when used with percussive modulator sounds. Give it a try!

<u>Echo</u>

The echo control sets the amount of a delayed version of the synthesis filter output that is added to the output. This is to be distinguished from the delay described above, since the echo signal contains both the delayed signal as well as the undelayed signal.

The three preset buttons for the Echo operation are "Full", "Zero", and "Ramp".

Selecting "Full" sets the amount of echo on each channel to the maximum. This means that the echo is of the same amplitude as the original sound.

Selecting "Zero" sets the amount of echo to zero.

Selecting the Ramp button sets the echo amount to an upward ramp.

<u>ETime</u>

The ETime control sets the time of the delay used in generating the echoes. As in the Delay control, the ETime value can range from 0 to 737 milliseconds in steps of 5.8 milliseconds.

The three preset buttons for the ETime operation are "Full", "Half" and "Zero".

Pressing the Full button sets the echo time for each channel to its maximum value (about 737 msec).

Pressing the Half button sets the echo time for each channel to half of its maximum value (about 368 msec).

Pressing the Zero button sets the echo time for each channel to, no, not zero, but its minimum value (about 5.8 msec).

<u>Regen</u>

The Regen slider controls the amount of the echo signal that is fed back into the echo input. This feedback allows multiple echoes to be generated. If regen is set to zero there will only be one echo generated. Higher values of regen will increase the number of echoes that are heard. The regen level essentially sets the rate of decay of the echo signal. When the regen slider is at its maximum level, the regen will be greater than one, resulting in a growing echo amplitude. Use this setting sparingly!

The three preset buttons for the Regen operation are "Zero", "Half" and "Full".

Pressing the Zero button sets the regeneration amount for each channel to zero.

Pressing the Half button sets the regeneration amount for each channel to half of its maximum value.

Pressing the Full button sets the regeneration amount for each channel to its maximum value. Please note that this maximum level is slightly greater than one.

<u>Slew</u>

The Slew slider sets the rate at which the analyzer control signal grows with time in response to a change.

The three preset buttons for the Slew operation are "Fast", "Slow" and "Ramp".

Selecting the Fast setting sets the slew for each channel to zero. This means that the channel slew rate

is at its maximum.

Selecting the Slow setting sets the slew for each channel to its maximum value. This means that the slew rate is at its minimum.

Pressing the Ramp button sets the channel slew to an upward ramp pattern.

<u>Warp</u>

The Warp slider sets the amount of warping applied to the analyzer control signal. This results in a compression or expansion of the analyzer control signal. When the slider is set to the middle of its range there is no warping of the control signal. When the slider is above the middle of the range, the control signal is compressed - that is, high control signals values are reduced and low control signal values are increased. When the slider is below the middle of the range, the control signal solutions are increased and low control signal is expanded - that is, high control signal values are increased and low control signal values are reduced.

The three preset buttons for the Warp operation are "None", "Cmpr" and "Expn".

Selecting "None" sets each channel warp level to the middle of its range. This results in a linear mapping being applied to the analyzer control value.

Selecting the "Cmpr" setting sets the channel warp to a level above the middle of its range. This results in a compression being applied to the analyzer control value. That is, low control values are amplified and high control values are attenuated. You should watch your levels in this mode, as the amplification of the low level signals generally results in an overall increase in levels, and clipping can result.

Selecting the "Expn" setting sets the channel warp to a level below the middle of its range. This results in an expansion being applied to the analyzer control value. That is, high control values are amplified and low control values are attenuated.

Input Mode

There are two inputs to the vocoder, the carrier signal and the modulator signal. These can be generated internally (see Carrier Source Selection and the Play/Record sections) or provided by inputs to the sound card's analog-to-digital converter (ADC). The input to the sound card's ADC is usually formed by mixing together a number of sources connected to the sound card, such as the microphone input, the line input, and the CD audio input. The relative level of each of these inputs can be adjusted with a mixer program supplied with the sound card. The ADC of most sound cards is stereo (the program will not work unless the sound card is stereo), providing two possible monophonic signal inputs. The vocoder program will use these two channels to provide the external carrier and modulator signal sources. For example, the left line input could supply the carrier signal and the right line input could supply the modulator signal. A common usage of the vocoder would be to use the sound card microphone input to provide the modulator signal. Unfortunately, on sound cards like the Creative Labs Sound Blaster, the microphone input is monophonic and is spread equally over the two stereo channels. To handle this case, the vocoder program assumes that one of the channels (say the left) contains the carrier signal as well as its portion of the microphone signal. The other channel (say the right) just contains the microphone signal. The carrier signal is then obtained by subtracting the right channel from the left (thereby cancelling out the microphone component). The modulator signal is obtained directly from the channel (in this case the right) which contains only the microphone signal.

There are four different inputs modes which determine how the input channels are assigned to the carrier and modulator signals. These are selected with the following set of buttons:



Note that only one of these buttons is on at a time. The first letter in the button label indicates to the channel used to provide the modulator signal. The second letter indicates the channel used to provide the carrier signal. That is,

L/R indicates that the modulator signal is provided by the left channel, the carrier comes from the right channel.

R/L indicates that the modulator signal is provided by the right channel, the carrier comes from the left channel.

M/L indicates that the modulator signal is taken from the right channel and the carrier is derived from the difference between the left channel and the right channel. In this situation the modulator would be a monophonic signal, such as the microphone input, and the carrier would be connected to the left input.

M/R indicates that the modulator signal is taken from the left channel and the carrier is derived from the difference between the right channel and the left channel. In this situation the modulator would be a monophonic signal, such as the microphone input, and the carrier would be connected to the right input.

It is crucial that you understand how to use the Mixer program supplied with your sound card! Review the mixer documentation. For example, in the Creative Labs "Creative Mixer" program the inputs to the sound card ADC are enabled by selecting the "Recording Controls" and clicking on the "XXX Input Recording Source" box, where "XXX" is one of "MIDI", "CD", "Line", or "Mic". These boxes have a little microphone icon. Set the recording level sliders for the inputs you are interested in to as high a level as possible without causing the vocoder to distort. In the "Volume Controls" settings of the mixer you will want to mute all outputs except for the "Wave Ouput".

Due to limitations in the Creative Labs sound blaster full duplex implementation, external inputs to the vocoder are digitized to only 8 bit precision. The vocoder output is digitized to a 16 bit precision. 16 bit

wav files can be loaded and used as input to the vocoder, however. Some sound cards may be able to digitize both input and output channels at 16 bits. For these cards select the "16 bit input option" in the DEV dialog.

It is strongly recommended that you install the latest full duplex drivers from your sound-card manufacturer. These are usually available from the manufacturers web page. For example, the Creative Labs web site (for SoundBlaster cards) is **www.creaf.com**.

MIDI

Each of the vocoder controls (buttons and sliders) can be associated with a MIDI Controller Message. This can be used either for input or output, or both. As an input a MIDI Controller message can be sent to the program from an external source (such as a MIDI keyboard) and used to adjust the value of one of the programs buttons or sliders. As an output a MIDI Controller message is sent to an external MIDI device when the associated vocoder slider or button is changed.

To setup the MIDI parameters, press the MIDI button.

This will bring up the dialog box shown below:

MIDI Device Selection		
Select the MIDI Input and Output	t Devices	
Select Input Device	Select Output Device	
None AWE64G MIDI In [330] TTS Virtual Piano In	None AWE64G MIDI Synth [620] AWE64G MIDI Out [330] Creative Music Synth [220] Creative MIDI Instrument Map WaveSynth/WG MIDI	
Select Input Channel	Select Output Channel	
Edit Controller Assignments		

You should select the desired MIDI input and output channels. Setting the input channel to OMNI means that the vocoder will respond to controller messages transmitted on any of the 16 possible MIDI channels. If you don't want to receive any messages, you should set the Input Channel to OFF, or set the Input Device to "None". Likewise, if you don't want to transmit any messages, you should set the Output Channel to OFF or set the Output Device to "None".

If a MIDI Input device has been selected, and the input channel has not been set to off, then any NOTE-ON messages that are received on the input channel will be used to adjust the frequency of the internal pulse generator. In this way, an external MIDI keyboard or other source of MIDI NOTE ON messages can be used to set the frequency of the pulse waveform. *Please note that to keep computational loads to an acceptable level it was neccessary to only have a limited accuracy in this frequency setting. Thus the frequency of the internal pulse waveform won't exactly match those of other external sound sources. The frequencies should be close, but audiophiles may notice!*

To associate MIDI Controllers with the vocoder program controls, click on the "Edit Controller Assignments" button in the MIDI DEvice Selection dialog box. This will bring up the MIDI Controller Assignment dialog box, as shown below.

MIDI Controller Assignm	ient	×
OK Vocoder Control Contro	Cancel	
Sample/Hold Butte 💌 10	Off	• 0
Play Button 🔽 56	Off	• 0
Rand Spectrum Bu 🚽 57	Off	• 0
Sibilance 🔹 12	Off	• 0
Channel Slider 1 💌 13	Off	• 0
Off 🗾 🛛	Off	• 0
Off 🗾 42	Off	▼ 0
Off 🗾 🖸	Off	▼ 0
Off 🔽 🖸	Off	<u> </u>
Off 🗾 🖸	Off	▼ 0

There are a total of 20 associations in this dialog box that can be made between vocoder controls and MIDI controller messages. Merely select the vocoder control in the left hand list and type in a MIDI controller number in the right hand list. This number must be between 0 and 121. *Please note that an association of a vocoder control with a MIDI Controller number will be active for both input and output. That is, any MIDI Controller message received with the proper controller number will alter the associated vocoder control and any change in the vocoder control will cause a MIDI Controller message to be transmitted.*

If two or more vocoder controls are associated with the same MIDI Controller number, then when any of these controls are altered, a MIDI controller message is transmitted with that number. Likewise, if any MIDI controller message of that number is received, each of the associated vocoder controls is changed.

If two or more MIDI Controller numbers are associated with a single vocoder control, then that vocoder control is changed whenever any MIDI Controller message is received with one of these controller numbers. When the vocoder control is altered, a MIDI controller message is sent out for each controller number.

Any unused associations should have the vocoder control set to off, so that no unexpected effects take place!

Mix Control

The MIX slider determines the proportion of modulator signal that will be mixed in with the vocoder output signal. Setting its level to the maximum will bypass the vocoder entirely, and output just the modulator signal.

Note that one can also bypass the vocoder, but this time output just the carrier signal by use of the \underline{EQ} mode.

Play/Record

One of the primary intended uses of the vocoder program is to generate wav files of vocoded speech. These can be then put to many different uses, such as in MOD files, or as sound effects in games or merely as sounds associated with various Windows events.

Recording and saving a snippet of vocoder sound to a wav file is very straightforward. To begin, press the REC button. This will begin the recording process. Then create your sound, and press the REC button when you are done. Please be aware that whenever the REC light is on, sound data is being put into an ever growing memory buffer in your computer's RAM, and at some point you will run out of room. At this point the recording process should stop (but the record light will stay on) and your computer just might get hung up.



Once you have recorded your sound bite, you can audition it by pressing the Play button. This will merely play back what is stored in the record buffer. During this playback the normal vocoding process will be suspended. If you press Play while recording, recording will be terminated. Once you are satisfied with your sound bite, it can be saved by pressing the Save button. This will bring up a sampling rate dialog box followed by a file dialog box which will prompt you first for the desired sampling rate to save the date at, and then for the name of a file to save the sound in.

Please note that the internal format of the output data is stereo, 16 bit, 11.025 KHz. You can save the data in one of three different sampling rates, however, these being 11.025 KHz, 22.05 KHz, and 44.1 KHz. The program will interpolate the waveform data to provide the higher sampling rates, and so there is no increased high frequency fidelity at the higher sampling rates.

NOTE: The save function is not available in the demo version!

Saving/Loading Your Settings

In the course of playing around with the vocoder controls you may happen across a configuration that you particularly like. This configuration can be saved for future recall by pressing the SSet button. Pressing this button will produce a dialog box prompting you for the name of a file in which to store the current control settings. Also stored is the spectrum mapping if the Rand button is on. These settings can then be restored at a future time by pressing the LSet button. Pressing this button will produce a dialog box prompting you for the name of a settings the spectrum mapping if the Rand button.

This will also save the MIDI controller mappings.

When you exit the vocoder program, the control settings at the time of exit are stored in a file named "session.set" in the vocoder executable directory. When the vocoder program is started up again, this file will be read in and the vocoder controls set to the state that they had when the program stopped previously.

NOTE: The settings saving function is not available in the demo version!

Sibilance Control

The action of the SIBIL slider is similar to the action of the MIX slider, but differs in that only a part of the frequency spectrum of the modulator is passed through. The portion that is mixed in with the output is obtained by high-pass filtering the modulator signal, with a cutoff frequency equal to that of the highest frequency analysis filter.

The purpose of the sibilance effect is to emphasize the unvoiced (i.e. "s" sounds and consonant sounds) components of speech. Most carrier signals do not have the broadband high frequency components neccessary to adequately represent unvoiced speech. Passing through some of the high frequency components from the modulation signal to the vocoder output partially makes up for this lack, and generally produces a more intelligible vocoder output.

Sound Card Devices and Startup

The Cylonix vocoder requires a sound card to function, preferably a sound card with full duplex capabilities. If your system has more than one sound card you can select between them using the sound card devices button:



Pressing this button will bring up a dialog box, as shown below, which will allow you to select which sound card device you want to use as the audio input and as the audio output. You do not have to select an audio input device if you don't want to, but you must select an audio output device. If you do not have a full duplex capable sound card, select "None" as the audio input device. In this way the single duplex sound card channel will be used for output only. If you do this you will not be able to input any signals from the sound card; you will have to use way files.



The Creative Labs soundblaster cards (up to the AWE64) have one 8 bit channel and one 16 bit channel. The Cylonix vocoder program uses the 8 bit channel for input and the 16 bit channel for output. If you have a sound card that has two 16 bit channels then select the 16 bit input/output option on the DEV dialog. This will then use 16bits for both inputs and outputs. Note that if you select this option when using a SoundBlaster card, the vocoder will not work properly.

You will note that there is about a half second delay between the input and output. This is an unavoidable consequence of the Windows operating system, and is NOT removed in the registered version of the program. Sorry! This delay reduces the usability of the vocoder in live settings (except in the hands of skilled practitioners!) but should not affect its use for generation of wav files, and for just plain goofing around. You can adjust the number of buffers the program uses (and hence the delay) with the slider in the DEV dialog window. If you set the number of buffers to too small a value, the sound will become choppy and the program may hang. If you only select an output device (i.e. you select "NONE" for the input device) you will be able to select a smaller number of buffers than when an input device is selected.

The sound card select window should popup at the beginning of the program. If it does not, wave the mouse around or just click on the DEV button. You will not be able to hear any sound until you select at least an audio output device.

Spectrum Display

The Cylonix vocoder includes a graphic display of the modulator signal energy in each of the 18 channels. This display is similar to spectrum analyzers found in many home stereos.

The display is updated approximately 11 times a second. It includes a fast responding display (the yellow bars) and a time-averaged display (the red lines). The vertical scale of the display is a compressed, approximately logarithmic, function of the channel energy. The display can therefore display a wide range of modulator signal energies.



Spectrum Modifiers

As described in the <u>introduction</u> the vocoder computes the energy of the modulator signal in each of 18 separate frequency bands. These energy measures are then used to modulate the level of the carrier signal in the same 18 frequency bands.

Usually the mapping of analyzer frequency bands to synthesizer frequency bands is the trivial mapping where the energy of the lowest frequency analyzer section filter output is used to modulate the output of the lowest frequency synthesizer section filter, the energy of the next highest frequency analyzer filter output modulates the next highest frequency synthesizer filter output and so on, as illustrated below:

ANALYZER SECTION FILTER	SYNTHESIZER SECTION FILTER	
1	>	1
2	>	2
3	>	3
18	>	18

The Normal Spectrum Mapping

Spectrum Shifts

Some interesting effects can be obtained if we alter the mapping of analyzer to synthesizer filters from the normal case given above.

A simple spectrum modification is to shift the synthesizer frequency bands relative to the analyzer frequency bands. For example, we can shift the spectrum up by 1/3 of an octave with the following mapping.

ANALYZER SECTION FILTER	SYNTHESIZER SECTION FILTER	
NONE	>	1
1	>	2
2	>	3
17	>	18

A Spectrum Shift of 1/3 Octave (one filter width)

There are six buttons on the vocoder display that can be used to set the amount of spectrum shift. These are:



Only one of these buttons can be on at one time (e.g. the -1 button in the situtation depicted above). Pressing one of these buttons will result in the application of the associated spectrum shift. A shift of 0 implies the normal spectrum mapping. A negative shift indicates that the analyzer section filter energies modulate synthesizer section filters of lower frequency, while a positive shift indicates that the analyzer filters are mapped to synthesizer filters of higher frequencies. In the case of a negative spectrum shift of N filter widths, the N highest frequency synthesizer section filters will not be modulated (their output will be zero). This is because there are no analyzer filters of high enough frequency to provide the modulation signals. Likewise, in the case of a positive spectrum shift of N filter widths, the N lowest frequency synthesizer section filters will not be modulated

In general, positive spectrum shifts cause the resulting vocoded speech to sound more female, and more "Donald-Duck-ey", while negative spectrum shifts resulting in more male sounding and muffled speech.

Spectrum Inversion

An interesting effect can be obtained by inverting the mapping between the analyzer and synthesizer sections. That is, the highest frequency analyzer section filter energy is used to modulate the the lowest frequency synthesizer section filter, the next highest frequency analyzer section filter energy is used to modulate the next lowest frequency synthesizer channel and so on, as shown in the following figure:

SYNTHESIZER SECTION FILTER	
>	18
>	17
>	16
>	1
	SYNTHES > > >

The spectrum is inverted when the INV button is pressed. Please note that any spectrum shifts will be applied *after* this inversion process. The sound that results from spectrum inversion is usually quite unintelligible (but interesting). The vocoder sometime distorts when in this mode; if this happens reduce the setting of the gain slider.

Random Spectrum Mapping

Another interesting spectrum modification is a random scrambling of the analyzer-synthesizer mapping. This can be obtained by pressing the Rand button. Doing so creates a randomization of the mapping. The scrambling is such that no two analyzer outputs are connected to the same synthesizer control input. Each time the Rand button is pressed a different scrambling pattern is generated. If you find a pattern that you like you should save it using the <u>SSet</u> button (Save Settings).

Linear Spectrum Mapping

Pressing the LIN button sets the spectrum mapping to the normal mapping of each analyzer section to the corresponding synthesizer section of the same frequency band.

Spectrum Swap

Pressing the Swap we button swaps adjacent spectrum values. This thickens the sound up slightly.

Spectrum Enhancement

Pressing the Enh button applies an enhancement operation to the spectrum. This operation enhances high spectrum values while diminishing adjacent spectrum values. You can see the results of this operation has on the spectrum by looking at the spectrum display.

Spectrum Smoothing

Pressing the Smooth button applies a smoothing operation to the spectrum. You can see the results of this operation has on the spectrum by looking at the spectrum display.

Spectrum Mapping Display

Pressing the TDisp button toggles the display between the Modulator Spectrum level display and the Spectrum Mapping Display. The Spectrum Map shows a matrix of channel map indicators. Each column of the matrix corresponds to a synthesizer channel, while each row corresponds to an analyzer channel. When the Spectrum Map is displayed, the user can adjust the mapping by moving the map indicators with the mouse. To do this, click, with the left mouse button, on the channel column whose mapping you want to change. The map indicator for this channel will turn red. Hold down the left mouse button and move the cursor up or down until the indicator is at the desired

location.

Spectrum Sampling and Holding

Pressing the Hold button will sample the analyzer control signals at the moment of the button press and will hold these values fixed until the Hold button is pressed again.

Pressing the S/H substantial button will *periodically* sample the analyzer control signals and hold them until the next sampling instant. The rate at which the analyzer output is sampled is set by the RATE slider.

Using the Vocoder as an Equalizer

The vocoder can also be used in an Equalizer mode, by pressing the EQ button. When this button is pressed the analyzer control signals are all set to the same constant value, independent of the modulation signal. The synthesizer spectrum can then be shaped by using the <u>channel level</u> sliders, resulting in an 18 channel equalizer!.