CUBE 2

Spectral morphing resynthesizer

Operation manual

VirSyn Software Synthesizer

Harry Gohs

Copyright

Trademarks

Caution !

Careless operation of this software can produce excessive sound pressure levels which can be harmful to your ears and can damage your loudspeakers or headphones. To avoid such damages, reduce the main output level before operating the software. VirSyn Software Synthesizer will not be liable to any damages whatsoever arising out of the use or abuse of this software.

04/2005

Contents

1 Introduction

1.1   Welcome to the CUBE !

Congratulations and thank you for purchasing CUBE !

CUBE is a multitimbral software synthesizer bringing you the power and complexity of additive synthesis to your computer studio. This kind of synthesis is well reknown for its potential capability of creating virtually any sound you can think of ! But its also well reknown for its huge amount of parameters you have to control for reasonable results. Therefore only a few synthesizers exists which uses this mighty synthesis technology.

With CUBE this prejudice will no longer be true. Specific sound morphing capabilities enables you to easily create complex evolving sound scapes, vocal phrases or stunning rhythms simply by choosing up to four sound sources from a predefined set. Then add one of the predefined envelopes and a unique, complex sound preset containing thousands of parameters is created ! If you want to dive deeper into the inner secrets of this synthesis monster you can manipulate each parameter individually to get precise control over the sound. A number of high level parameters such as Brightness and Harmonic spread control hundreds of synthesis parameters in a way modeled after the inner workings of natural instruments.

CUBE gives you also a new method of getting randomly generated sound patches in a way you never experienced before: try out our "Slotmachine" ! Every single try gives you a new unique sound patch, really every try, promised…

Since version 1.5 CUBE entered the realm of additive resynthesis. What does this mean ? The term "additive" comes from the fact that any sound can be reproduced by adding together a sufficient amount of simple sine tones with varying amplitudes and frequencies. This is true even for noisy sounds: white noise for example can be broken down to a mix of an infinite number of sine tones spread across the whole frequency range, all having the same amplitude. But the drawback of this synthesis method is the enormous amount of parameters which must be defined to get an acceptable result: hundreds of envelopes for the amplitudes and frequencies of the partials have to be programmed ! For this reason CUBE 1.0 already offered you a set of macro parameters and easy graphical editing methods to cut down the number of parameters to a managable amount. Since CUBE 1.5 you can go one step further: now you can import any sample and CUBE will automatically analyze this sample and extract the amplitude/pan and frequency envelopes for the most prominent partials it needs to reproduce the sound with its synthesis engine. This is what the term "Resynthesis" stands for. And even more – you can have up to four of them in a single preset and the morphing process of CUBE 1.0 is available for these new type of sources too.

What's new with CUBE 2 ?

The main new feature of CUBE 2 is the ability to manipulate the analyzed sound spectra in a variety of musically useful ways: Imagine for example drawing a comb filter structure freely into the spectrum ! And you can do this not only in the sine spectrum, but also with the residual noise spectrum and the stereo panning. The analyzer of CUBE 2 improved a lot and allows now specialized algorithms for different types of sounds. This includes harmonic / inharmonic and noisy sound types. The new residual noise generator of CUBE 2 allows you to create evolving noise spectra in a way which was never possible before.

Here are almost all of the new features of CUBE 2 at a glance:

- Improved Spectral analyzer with selectable algorithms optimized for different sound types
- Graphical editor for direct manipulation of the spectral and spatial structure of the sound.
- Improved morphing algorithm
- Residual noise generator
- Volume envelope now independent from the Morph envelopes
- Switch off morph envelopes for manual morph control
- Modulation range for morph modulation added
- New modulation sources "Breath" and "Expression"
- Editor for MIDI assignment of user interface controllers
- Legato mode for monophonic sounds
- Reset function for source edits and resynth sources.

1.2   Installation PC

For proper operation of CUBE the PC system requirements are:

- **Pentium III/IV, Athlon XP (SSE required !) with min. 600MHz**

- **Microsoft Windows XP.**

- **min. 512 MB RAM.**

- **min. screen resolution 1024X768 / 16-bit colors.**

- **ASIO or DirectX compliant Soundcard.**

- **400 MB free space on your hard disk.**

To install CUBE, insert the installation CD into your CD-ROM drive. The installation program starts automatically after inserting the CD-ROM. If for some reason it does not, you can start the installation program manually by opening the CD-ROM with the explorer and starting „setup.exe". The installation program leads you through the rest of the installation .

After starting the program the first time you will be asked for the registration code which is printed on the CD cover. After entering this code successfully the program is licensed on your computer. Please keep the registration code together with the license number in a safe place ! You will need them for every update and reinstallation of CUBE.

For the installation as a VST2 plugin you have to copy the file CUBE.dll from the installation folder of CUBE manually into the vstplugins folder of your favourite VST2 compatible host program. If you encounter problems entering the registration code within Cubase SX, please make sure that you switched on the Option "Plugin receives key commands" in Cubase SX before registration.

1.3 Installation Mac

For proper operation of CUBE the Mac system requirements are:

- **G4 (Altivec required !) with min. 400MHz**
- **Mac OS X 10.2**
- **min. 512 MB RAM**
- **min. screen resolution 1024X768**
- **CoreAudio (OS X)**
- **400 MB free space on your hard disk.**

To install CUBE, insert the installation CD into your CD-ROM drive. Start the installation by double clicking the Installer icon. The installation program leads you through the rest of the installation. The Installer will install the stand-alone version, the VST and AudioUnit version of CUBE.

After starting the program the first time you will be asked for the registration code which is printed on the CD cover. After entering this code successfully the program is licensed on your computer. Please keep the registration code together with the license number in a safe place ! You will need them for every update and reinstallation of CUBE.

If you encounter problems entering the registration code within Cubase, please make sure that you switched on the Option "Plugin receives key commands" in Cubase before registration.

1.4   Screen layout

This is CUBE's main screen shown here for explaining the terms used for each user interface element:

sound sources A/D   bank select   part select   program select   file menu   "slot-machine"

2D envelope editor   main envelope editor   sound sources B/C   arpeggiator
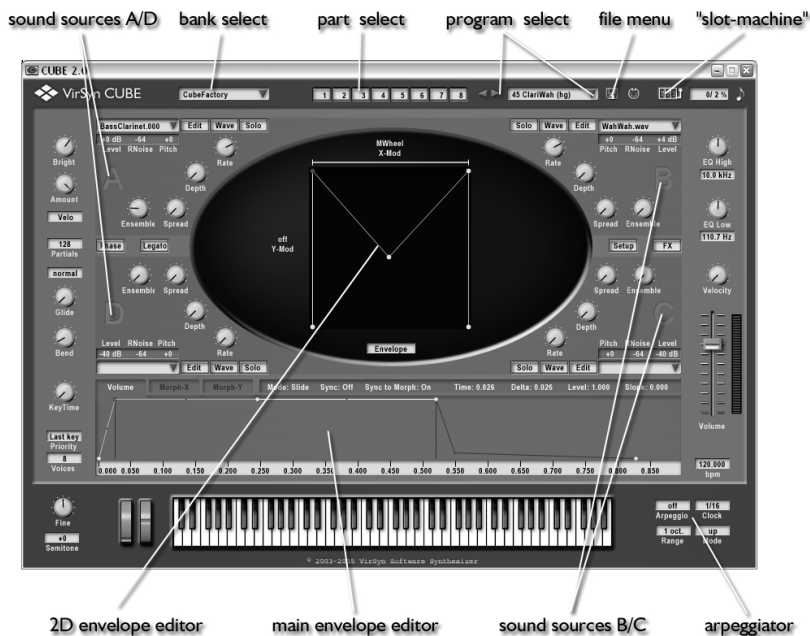
1.5   Basic operation

1.5.1   Parts and Voices

CUBE is a multitimbral synthesizer. Up to 8 parts can be used to play back up to 8 different sound programs simultaneously. Each part can receive MIDI messages from a user selectable MIDI channel. If you use the same MIDI channel for more than one part you can easily create stacked sounds. In the standalone version of CUBE the MIDI input device can be selected on the setup page. Only one part can be edited at one time. To select the part for editing just click into one of the 8 part buttons on the top of CUBE's window. These part buttons also have a built in light flashing on incoming MIDI notes.

Each part can have up to 32 voices, which means the number of simultaneously playable notes. The audio output of all parts is mixed down to the stereo audio output selected on the setup page in the stand-alone version of CUBE.

When used as VST or AudioUnit instrument the part outputs can be routed to up to 8 stereo channels which can be mixed down separately by the host program. The routing works in two ways:

-    if "Master" is selected as output for a part on the Setup page, this part is routed – together with the reverb return signal – to the stereo channel 1.

-    if "Out2"- "Out8" is selected as output for a part on the Setup page, this part is directly routed to the VST stereo channel 2 – 8.

The setup page can be opened with a click on the "Setup" button:

1.5.2 Banks and Programs

On startup of CUBE a default sound bank is loaded. This bank is a file located in CUBE's home folder and is named "CubeFactory.fxb". If you like to use your own default sound bank you can replace this file with a sound bank of your own. Each sound bank can contain up to 64 programs. The default sound bank delivered with CUBE is a selection of sounds collected from the six main sound banks delivered with CUBE containing about 500 ready to use sound programs. These sound banks are also located in CUBE's home folder. You can load one of these banks by clicking into the bank selection field on the left side of the part buttons or by selecting "Load bank" from the file menu which can be opened by a click on the disk button located at the right side of the program selection field:

This menu also offers you commands to:

- Save bank to file

- Load preset from file

- Save preset to file

- Copy preset to clipboard

- Paste preset from clipboard

- Clear preset for setting an empty default program

- Randomize

You can select one of the 64 programs by clicking into the small triangle inside the program field. To edit the name of the current program just click into the program field and enter the new name. You can exit the field after editing with the Enter key or through clicking on anything outside the program field.

You can zap through the programs of a bank by clicking into the left/right buttons on the left of the program field.

The randomize command has the same effect as clicking on the "slot machine" button on the upper right of the picture above: Every click selects sound sources and envelope settings randomly from the programs in the actual loaded sound bank to create new programs with a near 100 % usability guarantee !

1.5.3   General parameters

Now follows a description of the more common sound parameters which are typical for most synthesizers.

On the left side of CUBE's window you find the following parameters:

- Brightness and Amount
  These will be explained in the
  next chapter because they are not general at all…

- Glide
  Smooth transition of the pitch between two notes played
  in succession. The time to reach the final pitch can be
  adjusted in the range of 0 seconds to 132 seconds. This
  effect is typical for many lead sounds on monophonic
  synthesizers, but can be used here in polyphonic mode
  also ! In "normal" mode all notes played will have the glide
  effect. In "fingered" mode only legato played notes will
  have this effect.

- Bend
  Adjusts the bend range of the pitch bender. You can set
  the range from +/-1 up to +/-48 halftones.

- KeyTime
  This parameter controls all envelope times in CUBE
  0 % means that the envelope times doesn't change with
  the note played. Values greater than 0 % let the envelope
  time getting longer for notes below standard pitch and
  faster for notes above.

- Priority
  The key priority controls which note will be played if you
  press more than one key in monophonic mode. With key
  priority "Last key" always the last note will be played. With
  key priority "Low" or "High" the lowest or the highest key
  will be played.

- Voices
  Here you can limit the maximum number of voices which can be
  played simultaneously. Voices set to 1 means monophonic mode.

On the right side of CUBE's window you find the following parameters:

- EQ High
  High pass shelving filter with an adjustable range
  of +/- 40dB. The cutoff freqency has a range from 4 kHz
  to 16 kHz.

- EQ Low
  Low pass shelving filter with an adjustable range
  of +/- 40dB. The cutoff freqency has a range from 20 Hz
  to 200 Hz.

- Velocity
  This is a built in modulation of Volume by MIDI velocity:
  0 dB means no modulation
  60 dB means maximum modulation, the volume varies
  by 60 dB depending on MIDI velocity.

- Volume / Level meter
  The main part volume slider has a range from –60 dB to
  +20 dB. The level meter is showing the actual level of
  the signal. If the signal is above a certain level the
  output will get distorted in a soft way.

- Bpm
  The tempo in beats per minute used by the internal
  arpeggiator and the syncable envelopes and delay
  effects. In the stand-alone version the tempo can be
  adjusted manually. In the VST and AudioUnit versions of
 CUBE this value is automatically synchronized to the
  host.

On the bottom left side of CUBE's window you find the following parameters:

- Fine tuning
  This parameter adjusts the fine tuning for the whole part. The range is from – 50 cent to + 50 cent (one semitone).

- Semitone
  This parameter is for coarse tuning in the range of +/- 48 semitones.

- Modulation wheel
  This can be used to modulate the morphing envelopes and the Brightness parameter.

- Pitch bender
  The range for the pitch bending effect can be adjusted for each program separately with the bend range parameter already described.

On the bottom right side of CUBE's window you find the following parameters:
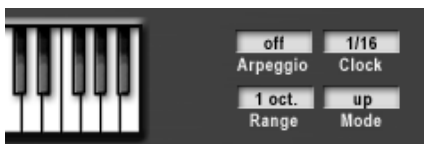
- Arpeggio
  For each part you can switch on an independent arpeggiator. With the hold mode the arpeggio doesn't stop when all keys are released.

- Clock
  Here you can select the base note duration and predefined rhythms.

- Range
  The range for the arpeggiated notes from 1 to 8 octaves

- Mode
  Controls how the hold chords are broken into single note sequences. Possible values are "up", "down", "alternate" and "random".

1.5.4   MIDI learn and others

Many parameters of CUBE can also be adjusted by using MIDI controller messages. You can easily assign a parameter to a MIDI controller by the following procedure:

- right click / ctrl click on the parameter you want to assign.

- Choose "MIDI learn" from the popup menu.

- Turn the knob on your MIDI controller box to send some controller messages to CUBE.

- Now you will see the knob in CUBE turning in sync to the external controller box.

To cancel an already made assignment just select "MIDI forget" in the same menu. All assigments are saved in a special file "VSMIDI.vsm" in CUBE's home folder.

On the upper right side of CUBE you see the actual number of partial tones generated and the CPU load in %. The small note symbol on the right is the "MIDI panic" function: a click on this will eliminate "hanging notes".

You can click on the MIDI connector button to see all MIDI controllers which are associated to a knob or selection on the user interface and you can even change the controller numbers directly.

On the upper left side you can make a right click on the company logo to get a menu full of surprising effects...
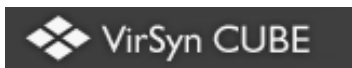
If you make a right click any where else in the background you will get a help menu from where you can open the about box or open this manual and a link to our website.

1.5.5 Microtuning

Microtuning gives you the possibility to use other tunings than welltempered. In CUBE you can change the tuning for each part separately. In CUBE's installation folder is a new subfolder "Tuning" created for this purpose. CUBE looks there for Scala compatible tuning definition files (.scl, .tun) named "part1.scl", "part2.scl" and so on. If you want to change the tuning for part 1 for example you simply have to put the corresponding Scala file into the subfolder "Tuning" and rename it to "part1.scl". On the next start of CUBE it will then use this tuning for part 1.

1.5.6 Editing commands

The Mac and the PC version have the same functionality, but they differ slightly in the use of the mouse and in certain key strokes according to the usual user interface on their respective platform.

Editing of parameters is done straight forward by selecting a knob or slider with a click of the mouse and then dragging the mouse up or down to change the value. If you hold the shift button pressed while dragging the mouse most parameters are changed in finer steps. You can also change the value of parameters with the key commands below. The actual value of a parameter is shown as soon as you move the mouse over the parameter you want to investigate.

The following table shows the corresponding key strokes and mouse usage on the Mac:

PC	Mac	Function
Right Mousebutton	**ctrl+ Mousebutton**	Context menu
ctrl+ Mousebutton	**alt+ Mousebutton**	Special drawing in graphic parameter editor
↑	↑	Increments selected parameter
↓	↓	Decrements selected parameter
Pg ↑	**pg up**	Increments selected parameter coarse
Pg ↓	**pg dn**	Decrements selected parameter coarse
Pos 1	**Home**	Sets minimum value for selected parameter
End	**end**	Sets maximum value for selected parameter
Del	**Backspace**	Sets default value for selected parameter
Alt+F4	**Command+Q**	Quits CUBE

Please note that the general availability of the keyboard shortcuts above depend on the used Host program: with Cubase you have to enable the option "plugin receives key commands" to enable the key functions. With Emagic Logic currently CUBE have no key functions at all.

# 2 Quickstart

2.1 The easy way

If you just want to jump into the water making some new sounds without bothering too much, this is the right chapter for you !

The simplest way to get some new sounds is selecting some of the predefined sound sources from the four "Custom source" menus grouped around the "oval" window. Then select one of the predefined envelopes which can be selected in the context menu of the envelope:

1. Click into one of the four source selection fields.

2. Select a predefined source

3. Repeat this for other sources too
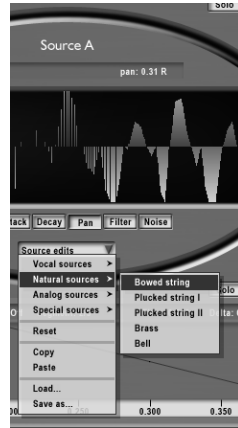
4. Right click (Mac: Ctrl click) into the envelope field below the oval window.

5. Select one of the predefined envelopes

6. If you want another sound go back to 3. or 4.
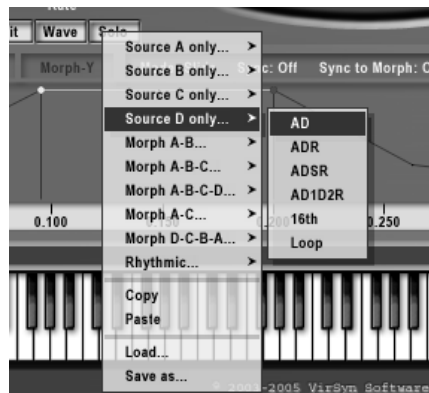
3 CUBE inside

3.1   Additive synthesis

CUBE is based completely on additive synthesis. The general concept behind additive synthesis is that every sound in the world can be reproduced exactly by simply mixing the output of many oscillators each producing only simple sinewaves named partials. The amplitude and frequency envelopes of these oscillators are the only parameters you will need to modify to get any sound you can ever think of ! Most traditional synthesizers are working just the other way: they take oscillators producing complex waveforms like sawtooth with a rich sound compromised from hundreds of sine waves and then trying to shape and modify this base sound to get the desired result. With additive synthesis you can control the timbre of a sound much more precisely by changing the levels of each partial separately.

You can ask now that if it's so simple to create any sound with additive synthesis why aren't there more synthesizers using this simple synthesis technique ?

The answer is that the method is simple but since you need hundreds of partials with hundreds of envelopes you get stuck in such a huge amount of parameters that you will need years to get one sound programmed. Even worse if you want to control sounds drastically by simple modulations with key velocity, modulation wheel or others you also have to control this huge amount of parameters !

CUBE introduces a new approach making the access to this migthy synthesis technique very easy and transparent. CUBE's concept and the power of todays computer technology eliminates many of the limitations which were inherent in earlier implementations of additive synthesis. Observations on the sound generating mechanism of natural instruments helped in creating high level parameters modeling important sound aspects like the Brightness or the Inharmonicity of natural sources. Even the liveliness of sounds can be controlled with CUBE's unique high level parameters.

This should be enough theory for the moment, you're already curious…

3.2 Global sound parameter

On the upper left side of CUBE's window you will find some very useful global sound parameters:

- Brightness / Amount
  This parameter controls the brightness of the sound of all sources. Its characteristic is modeled from the behaviour of natural instruments played loud or soft. In contrast to the well known simulation of this phenomenon by modulating the cutoff frequency of a low pass filter in most synthesizers, CUBE uses a low pass filter with a cutoff frequency in sync with the pitch of the played note and changing the slope of the filter in a range from 0 to 30 dB/Octave!
  This is a much more realistic simulation. The Brightness can of course be modulated with an adjustable amount (hence the Amount knob) and by a source which can be selected in the field below the "Amount" knob.

- Phase
  With "Phase" switched off all sine oscillators are running "freely". This has the effect that the same note value played repeatedly can sound slightly different from each other. This is the preferred mode for most sounds. With "Phase" switched on the sine oscillators starts synchronously with maximum level on each note start. Then repeated notes all sound equal. This is the preferred mode for drum sounds. Note that this mode can lead to unexpected clicks in the sound!

- Partials
  This parameter defines the maximum number of partials used. This heavily affects CPU usage while playing low pitched notes !
  The range for this parameter is from 4 partials to 512. There is also a special parameter value called "auto" which has the effect of "thinning" out the number of high ordered partials. This can be used for special effects or just to reduce processor usage…

- Legato
  Switch legato on to suppress retriggering of envelopes while playing legato notes. This works only if you set the number of voices to 1.

3.3   Sound sources

The heart of the sound engine of CUBE is surprisingly simple: it's a set of up to 512 sine oscillators per voice with controllable level and frequency. The output of each sine oscillator is given a pan position and summed up for each part.

A sound source is a set of parameters used to drive this engine, here is a complete list of parameters you can set for each sound source:

1.   High level parameters

- Rate/Depth
  Pitch LFO rate and depth

- Ensemble
  This parameter simulates the effect of more than one instrument playing the same tune together. Behind the scenes this parameter controls up to 512 LFOs modulating each partial individually !

- Spread
  This macro parameter adjusts the frequency ratios between the partials: With spread set to 0 the partials are perfectly harmonic, the frequency of the partials then have the ratios 1:2:3:4 and so on. When the spread parameter is below 50 the upper harmonics are getting more and more detuned in the same way a piano string exhibits. With the spread parameter going beyond 50 the partials are getting very inharmonic with a bell like characteristic.

- Pitch
  The base pitch for this source. Range is +/-48 semitones.

- Level
  This is the overall level of the source.

- RNoise (residual noise)
  This is the relative level of the residual noise part in relation to the sine partials. Works only with resynthesis sources.

2.   Partial parameters

The following parameters are on the "EDIT" page:

- Partial levels
  This defines the static level for each partial group.

- Attack/Decay time
  Every partial group have a simple Attack/Decay envelope generator.

- Pan position
  The position of each partial group in the stereo panorama.

- Filter
  A filter with arbitrary characteristic and special morphing capabilities

- Noise
  Adds a noise spectrum to each partial. The noise can have an arbitrary spectrum.

Often you will need to edit one of the high level parameters Rate/Depth or Ensemble and Spread for all four sources simultaneously to suppress the morphing of these parameters. This can be achieved easily by holding the Ctrl-Key (Mac: Alt-key) down while dragging the knob with the mouse.

Every program can use four sound sources labelled A,B,C and D. These are the basic building blocks used by the morph envelopes described in the next paragraph.

If you click into the selection field a file dialog will open where you can select an analysis file for resynthesis. For more details go to the chapter "Resynthesis". If the selection field is empty CUBE is working without analysed data and uses a constant default instead.

The Solo function is very helpful during design of a new preset or the examination of existing ones. With Solo turned on for any of the sources A-D the morphing process is disabled and you hear only the contribution of the source you switched to "Solo"

3.4   Spectral morphing

One of the key ideas behind CUBE is the spectral morphing. Instead of burdening the sound programmer with drawing hundreds of envelopes to control the dynamic spectral characteristic of a sound you can simply morph between these four sound sources controlled by a complex 2 dimensional Envelope having up to 64 time level segments. In addition to this you can assign a modulation source like velocity, aftertouch,… to both morph dimensions.

With CUBE 2 you can also switch of the morph envelopes completely by selecting the option "Manual" in the field below the morph square. Then you can manually morph with the mouse or with MIDI controllers.
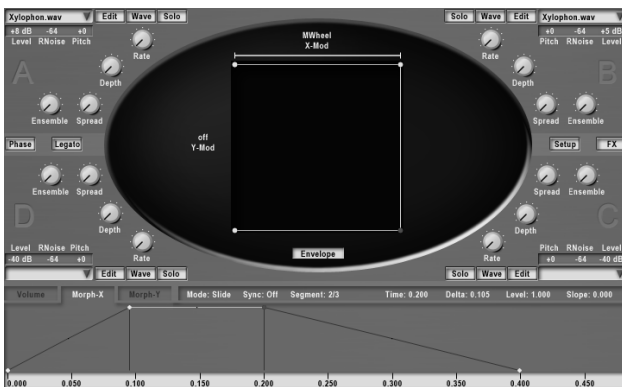
In the inner square of the oval window above you see the 2-dimensional envelope which controls how the four sound sources are used for the final sound. It is surrounded by the high level parameters of the four sources A,B,C and D. The lines in this window determines the amount each source is participating in the final sound for a given time point. The two morph envelopes can also be drawn in the second envelope editor below the oval window. They are always in sync with the editings inside the cube window. In the Morph-X and Morph-Y envelope editors you can also define the time points for each segments and other attributes you can't change in the 2D envelope editor.

What happens in the example above ?

The envelope starts with the first envelope segment (attack) in the upper left corner therefore source A is dominating the sound for 100 %. In the next 200 ms the sound is moving from source A to source B in the upper right corner and all parameters are smoothly morphed from A to B. The next envelope segment (decay) morphs the sound smoothly from source B to source C because the line is drawn form the upper right corner down to the lower right one. After 400 ms

the sound no longer moves because the sustain phase (indicated by the black horizontal line in the envelope) is reached. The sustain phase of the sound is dominated by source C. After releasing the key the sound is smoothly morphed from source C to source D and the sound ends after 600 ms.

The morphing controls all parameters of the sound sources. If you set for example the detune parameter of source A to +12 semitones and of source B to 0 semitones and the spread parameter of source A to maximum and B to 0 then morphing from source A to B will start as an inharmonic sound with the first partial one octave above normal pitch sliding down to a harmonic sound with normal pitch.

The morphing envelopes in X- and Y-direction can both be modulated by one of the following modulation sources:

- Keytrack
- Velocity
- Aftertouch
- Modulation wheel
- Breath
- Expression

These modulations works as follows:

Take for example a morphing envelope which starts at source A and ends at source B. If you then assign the modulation wheel as a modulation source for the X-axis (X-mod) the morphing depends on the setting of the modulation wheel as follows:

- mod wheel at minimum
  only sound source A is in effect, no morphing
- mod wheel at maximum
  morphing from A to B, same as without modulation

If you assign the modulation wheel to the Y-axis (Y-mod):

- mod wheel at minimum
  morphing from D to C
- mod wheel at maximum
  morphing from A to B, same as without modulation

With CUBE 2 you can specify the modulation range of the MIDI controllers by dragging the endpoints of the gray line shown on the left and upper side of the cube.
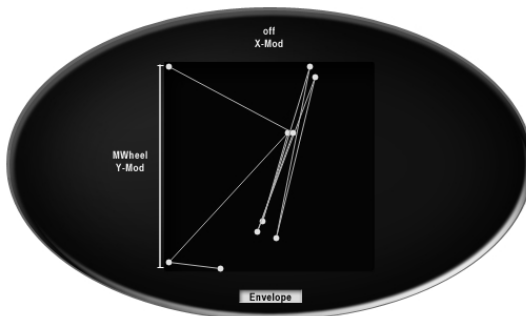
3.4.1 Envelopes

CUBE provides you with three multi segment envelopes:

- Volume
- Morph X
- Morph Y

The envelopes can have up to 64 time/level segments giving you all from simple Attack/Decay style envelopes up to complex rhythmical patterns which can even be synchronized to MIDI tempo.

The three envelopes are normally not independent from each other: They always have the same number of segments and the time when a specific segment starts is always the same for all three. With CUBE 2 you can make the Volume envelope be independent of the morph envelopes by switching off the option "Sync to morph".


2D morph envelope

The two Morph envelopes can be edited in two ways: you can edit the levels – which are in effect the current position of the sound in the morphing plane in X- and Y-direction – with the 2D envelope editor inside the cube of the central oval window. This is a convenient way to change the morphing characteristic without changing the timing of the envelopes. In the main envelope editor right below the oval window you can select one of the three envelopes for editing and there you can change all parameters including the timing and the slope of the envelope segments which is not possible with the 2D Envelope.

The main envelope editor figured below provides you with a context menu for copying and pasting envelopes through the clipboard and for loading and saving your own envelope presets. In addition to this you can simply select there from about 64 predefined envelopes.

Editing the envelopes is an easy and straightforward process:
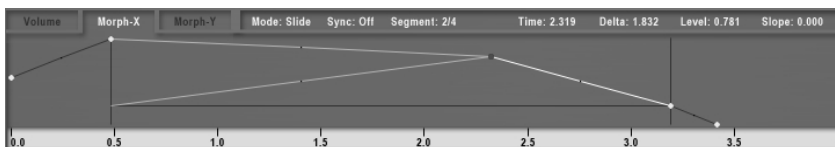
You can change the time or level of any segment by clicking on the small yellow circle at the end of each segment, it turns red and you can drag it into the desired position. Clicking the small black dot within the segment changes the slope. All editing take effect immediately !

With this click you also have selected this segment and then you can use alternatively the value fields for delta, level and slope for editing these values. Click into the appropriate value field and drag the mouse up or down for changing the values.

To insert a new segment click into the desired position in the envelope window while holding the shift-key down.

To delete a segment click into the small circle representing the segment you want to delete while holding the shift-key down.

The two black vertical lines are the loop marker positions for the envelope. They are managed automatically by CUBE. The looping of the envelope between these two lines come atomatically in effect as soon as more than one envelope segment is inside the marker positions.

You can scroll and zoom in the envelope window concurrently the following way: click on the time bar below the envelope window and drag up/down for zooming and drag left or right for scrolling in the window.

If you want to edit a segment without changing the positions of other segments switch off the slide mode with a click into the mode field above the envelope.

To create tempo synchronized envelopes click into the sync field which will show then the value "On" and the time bar will show you then note values instead of absolute time.

The "KeyTime" parameter located on the left side of the envelope window provides you with the possibility to modify the envelope times dependent on the pitch of the currently played note: high pitched notes get then shorter envelope times and low pitched ones get longer envelope times. This feature is only available if the tempo sync for the envelopes is switched off !

3.4.2   Edit windows

The following chapters describe the "inner" parameters of a sound source. To access this parameters click on the corresponding "Edit" button for one of the sound sources A-D. Then you will get access to six different graphical editors manipulating different aspects of the sound source.

If you click into the selection field "Source edits" you will get a menu of functions to set all parameters at once:

- a list of predefined source presets from where you can create your own sounds right out of the box.

- Copy and paste gives you the possibility to exchange source presets across different sources and programs.

- Load and save functions allows you to create your own library of source edit presets for later reuse.

- Reset sets default values in a way that you can hear resynthesized samples unmodified.

3.4.3   Partial parameters

Partial level editor
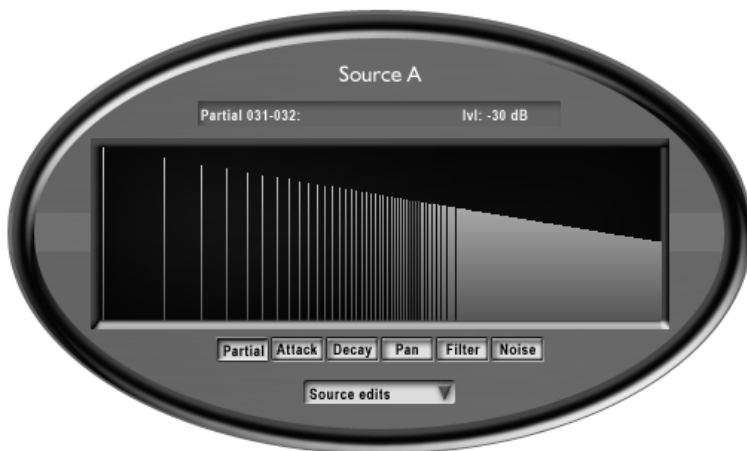
In this editor you can directly draw the base level of each partial group for up to 512 partials. The first 16 partials can be edited individually, all other partials can be edited only in groups of 2, 4, and so on. When you move the mouse over the partial level editor the nearest partial group below the mouse cursor is selected and the actual level setting for this group shown in the info field on top of the window. On the left side of the window are the low pitched partials, on the right side the high pitched ones. Please remember that the levels edited here are not the final partial levels because they dynamically are modified by the following editors. You can select also some predefined waveforms like sawtooth from the context menu which pops up with a right click into the editor window. Here you will also find functions for copying and pasting partial level settings between other sources or presets.

Partial attack/decay time

Each partial has it's own simple envelope generator with simple Attack/Decay characteristic. More complex envelopes are achieved by using the main envelopes for morphing between the four sound sources. The default setting here is attack time zero resp. decay time "endless" for all partial groups. You can reset to this default by selecting the "Reset" entry from the context menu of this editor. There you can also copy and paste the attack/decay settings between other sources or presets.

The decay settings you see in the example picture above will give you a sound of a plucked string: the higher partials are decaying much faster than the lower ones, which is a main characteristic of any plucked string.

Partial pan position
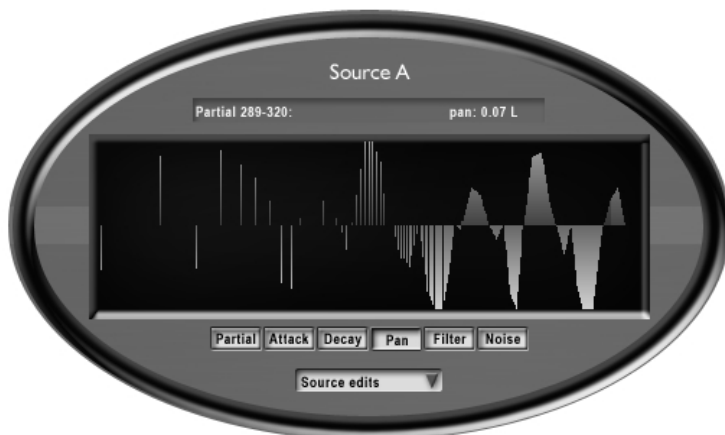
In the partial pan editor you can define the pan position for every partial group. The default setting here is mid position for all partial groups. You can reset to this default by selecting the "Reset" entry from the context menu of this editor. There you can also "copy", "paste" or "paste reverse" the pan settings between other sources or presets.

This function allows you to spread a single sound over the stereo panorama and make arbitrary complex pan movements when morphing different sound sources without using any effects !

For a more traditional pan effect set the pan position for source A all to maximum left position and for source B to maximum right position. If you morph between both sources the sound will pan with all its partial from left to right.

Morphing Filter

This is one of the really complex features of CUBE. Let's begin with the simple part: the filter without considering the morphing capabilities. In contrast to the partial editors described before this filter acts independently of the partials. It works like a classical filter or in analogy to natural instruments: It's the body of the instrument influencing the sound by characteristic body resonances. Traditional analog filters have types like lowpass, highpass and so on. The filter in CUBE can have an arbitrary filter transfer function where you can simulate all possible filter structures. For each frequency range you can define how strong the input signal will be attenuated. On the left side you can define the attenuation for low frequency partials and at the right side for the high frequencies. You can define the filter transfer function by simply drawing the filter curve inside the editor window. You have additional drawing modes if you press the Ctrl- or Shift-Key (Mac: Alt- or Shift-Key) while drawing with the mouse. And of course there is also a context menu allowing you to reset the filter and copy and paste filter settings between sources and programs.

The example above is showing you the filter characteristic of the vocal trakt from a soprano singer while singing an "a". You see the typical formant structures which are very important for the identification of a vocal sound. In the source presets you'll find a complete set of vocals for soprano and bass.

Now coming to the morph points. With the help of 8 morph points which are shown as small blue circles you can get the filter moving. The first and last morph points are always fixed to the leftmost and rightmost position, the other 6 morph points can be freely positioned in between. If you want to morph from one vocal lets say "A" to another vocal "U". Both of these vocals have three different characteristic formant frequencies which you can draw as bandpass filters in source A for the first vocal and in source B for the second vocal. Then drag the first three morph points on top of each formant in both sources.

Morphing between source A and B will then smoothly change the formant frequencies from those for the first vocal to those from the second one ! This is exactly what happens in a "real" speaking or singing voice: The formant frequencies shift from one vocal to the next.

You can try this easily with the predefined sources delivered with CUBE:

1. Select "Bass A" from the source select menu of source A

2. Select "Bass U" from the source select menu of source B

3. Select the envelope "Morph A-B.../ADSR" from the context menu of the envelopes

4. Now you should hear "Auhh…" like sound if you play with the keyboard!

Noise spectrum

The noise spectrum editor gives you the possibility to add noise to your sound. The amount of noise added can be individually adjusted for each frequency range. Just draw the desired noise spectrum in the editor window. What actually happens is that the partial frequencies are modulated randomly. And the amount of modulation depends on the frequency of the partial which is taken to look up the noise amount in the noise spectrum. For this reason the noise is not independent from the partials but are somehow "connected" like in many natural sound sources. Therefore you won't hear any noise if there are no partials active in the frequency region where you have drawn a high amount of noise in this editor. This also have the consequence that the noise color is "moving" with the base pitch of the sound. With CUBE 2 you can also add a noise spectrum independent from the partials with the new "Residual noise generator" described in the chapter "Resynthesis".

The default setting here is no noise modulation in any frequency range. You can reset to this default by selecting the "Reset" entry from the context menu of this editor. There you can also copy and paste the noise settings between other sources or presets.

4 Tips & Tricks

4.1   Delayed vibrato

Suppose that you want to create a vocal sound with a pitch vibrato slowly increasing during the attack phase. In modular synthesizers you can do this with a modulation of the pitch LFO level by an envelope. CUBE has no modulation matrix, how can sounds like this be created ?

For most questions of this type the answer will be the same:

Use the morphing of sound sources !

Here is a step by step instruction how it can be done:

1. Make an empty preset through the file menu of CUBE: click on the small disk icon in the tool bar of CUBE and select "Clear preset".

2. Select from the menu of sound source A the preset "Soprano A"

3. Do the same for sound source "B"

4. Now set the pitch LFO rate for source A and B to the same value around 6 Hz. You can do this in one step if you hold the Ctrl-key (Mac: Alt-key) down while dragging the "Rate" knob with the mouse.

5. Set the LFO depth parameter of source A to a small value, the depth parameter for source B should remain zero.

6. Now select the envelope "Morph-X" and set the attack time (the first envelope segment) to a value around 1 second through dragging the second envelope point a bit to the right.
    Then set the level of the first envelope point to maximum.

7. Press a key on the keyboard: The sound starts now with source B without vibrato and smoothly morphs into source A with vibrato through the morphing process !

4.2   Metallic morph

This example shows you how to morph from an inharmonic metallic sound to a harmonic vocal "Ahh".

Here is the step by step instruction:

1.   Make an empty preset through the file menu of CUBE: click on the small disk icon in the toolbar of CUBE and select "Clear preset".

2.   Go to the partial editor of source A and select "Sawtooth" from the context menu there.

3.   Select from the menu of sound source B the preset "Soprano A"

4.   To make the "Ahh" more prominent go to the partial editor of source B and select "Impulse" from the context menu there.

5.   Turn the "Spread" knob of source A to maximum

6.   Select the "Morph-X" envelope. Set the attack time (the first segment) to a minimum.

7.   Drag the level of the last three envelope segments to maximum. Now this envelope should look like this:
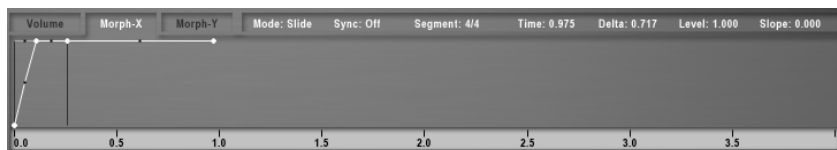
8.   Press a key on the keyboard: The sound starts now with a metallic sound smoothly changing to a vocal "Ahh" !

5 Resynthesis

5.1   The Basics

Where does Resynthesis fit into the concept of CUBE 1.0 ? First let's have a look how additive synthesis works there without considering resynthesis. This make it easier to understand how the resynthesis of the new version fits into this sheme.

With CUBE 1.0 you have a set of up to 512 sine tone generators which represents the synthesis engine. Each of this generators have the simplest possible envelope for amplitude / pan and frequency: The amplitude is always maximum level, the pan position is in the mid of the stereo panorama and frequencies are in strict harmonic proportion and doesn't change over the time. The only way to make the sounds "move" is to change the high level parameters. You define the overall level and pan for each partial separately in the partial editors. There you can also define the morphing formant filter and the noise modulation. The only way to break the harmonic character of the sound is to use the "Spread" parameter which smoothly turns the harmonic ratios of the partial frequencies into an inharmonic structure. Finally you can use the morphing between four complete parameter sets to make more radical changes in the movement of the sound. This parameter manipulations result in quite complex envelopes for level, pan and frequency of each partial. Why did we made it this way instead of simply offering 512*3 separate envelope editors for the individual partials ? The reason is that it's much easier to change the characteristic of a sound in a musical meaningful way through the high level parameter CUBE offers.

CUBE 2 now adds another, much deeper level of manipulation of the analyzed sample. Imagine for example drawing a comb filter structure freely onto the spectrum ! And you can do this not only in the sine spectrum but also with the residual noise spectrum and the stereo panning.

The analyzer of CUBE 2 improved a lot and allows now specialized algorithms for different types of sounds. This includes harmonic / inharmonic and noisy sound types. The new residual noise generator of CUBE 2 allows you to create evolving noise spectra in a way which was never possible before.

But remember that the basis for all sound manipulations in CUBE 1.0 are the constant envelopes of the partials. And this is exactly the point where the new resynthesis module fits neatly into the existing concept. When you import a sample the underlying constant envelopes are replaced by the time varying envelopes which results from analyzing the sample. All the macro parameter editors are then applicable to resynthesized sounds the same way as before !

After all this theory it's now time to get into the new parameters for the resynthesis module and how you can use them.

5.2   Wave analyzer

The Wave analyzer module is the heart of the resynthesis and has only a few
parameters you have to handle. For each source A-D you'll find a new button
"Wave" which opens the corresponding wave analyzer/editor window for this
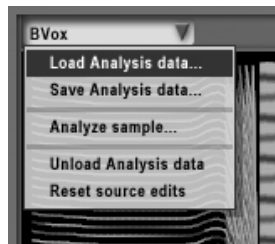source:
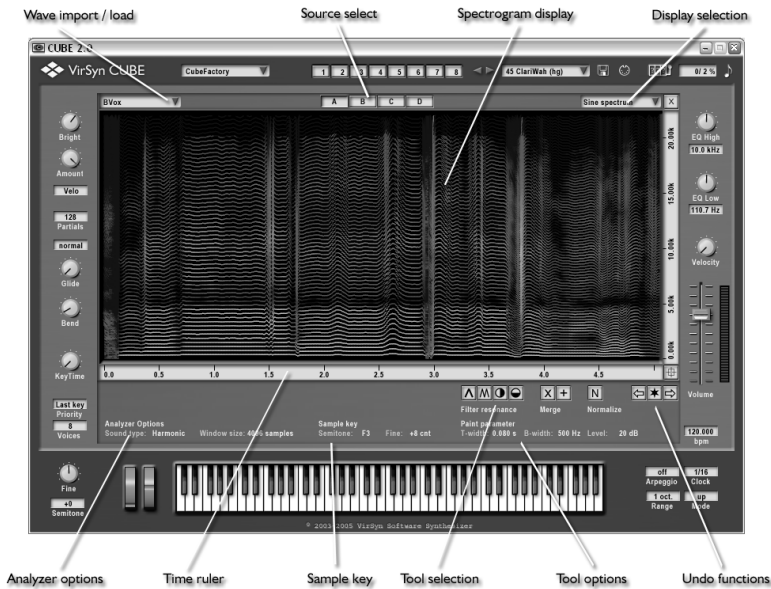
The wave import/load menu allows you to import sample files or to load already
analyzed data files into the source. You can switch between the four sources
with the source select button. With the display menu you can select which
aspect of the analyzed sample is displayed in the spectrogram display window.
You can zoom and scroll this window both on the time and frequency axis. Below
the spectrogram window you find the analyzer options, a toolbar to specifiy the
drawing operation you want to perform together with some drawing parameters.
On the right side you have access to a convenient undo functionality for your
spectral drawings.

Wave import / load menu

- **Load Analysis data**
  This opens a file dialog from where you
  can select an already imported analyser
  data file. It is important to note that CUBE
  doesn't use sample files directly, they must
  be imported before they can be used. The
  Analyser data file reside in the subfolder
  "Data" in CUBE's installation folder and
  have the extension ".vrd". The name of the
  imported sample is also shown in the
  selection listbox.

- **Save Analysis data**
  This saves the analysis data for the current source. You only need to
  save if you made changes to the spectrum with the drawing functions
  described below.

- **Analyze sample (Import)**
  This opens a file dialog where you can select a sample file in **WAV**
  format or **AIFF** format. This sample file is then analyzed by CUBE and the
  resulting analyser data is automatically saved in a file with the same
  name and with the file extension ".VRD" into the "Data" folder. The
  name of the imported sample is also shown in the selection listbox.

- **Unload Analysis data**
  This function resets the resynthesis source.

- **Reset source edits**
  This function resets all source edits applied to this source.

**Display selection**

- **Time warp envelope**
  Displays the sine spectrogram together with
  the time warp envelope. Here you can
  modify the time warp envelope.

- **Sine spectrum**
  Displays the sine spectrogram showing the
  partial tone tracks in time and frequency.
  The color of the lines indicate the relative
  level.

- **Residual spectrum**
  Displays the residual (noise) part of the spectrum. Note that the residual
  spectrum is not available with all import options.
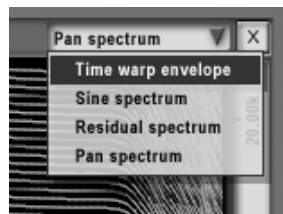
- **Pan spectrum**
  Displays the stereo panorama position for each partial together with it's temporal evolution. The line color indicates the position: gray lines means sound is centered, red lines means sound is on the left and blue lines means sound is on the right.

Analyzer options

- **Soundtype**
  For best import results you can select here the type of sound you want to import before actually analyzing the sample. Here is a list of the available options:

  1. Harmonic
  Works best for harmonic pitched sounds. Samples imported with this option give you a very clean sound over a broad pitch range. But this option needs a clean sample without extreme modulations or too much effects. The noisy part of the sample is separately analyzed in the residual noise spectrum.

  2. Inharmonic
  Use this option whenever the previous doesn't give a good result. The disadvantage is that this option uses much more partial tones for reconstruction of the original sample. The noisy part of the sample is separately analyzed in the residual noise spectrum.

  3. Noisy
  Is essentially the same as 2.) but the residual noise contains the complete analyzed spectrum and not only the residual. This option is the best for sounds with high noise content.

  4. Mixed
  This enables the analyzer used by CUBE 1.5 and is best for very complex sample material. With stereo samples this option also tries to map the stereo information into the pan spectrum. Note that this option doesn't deliver a residual noise spectrum.

- **Sample Key**
  The analyzer tries to figure out the sample or root key of the imported sample. With this parameter you can modify the estimated key. The proper key is important because you have to "tune" the imported sample if the estimation of the analyzer fails. This setting also influences the way morphing is done between different sources: For "best" results all sources should have the same tuning. But it's also very interesting to experiment with different settings of the sample key: the morphing process will change in a very subtle manner.

5.3 Spectrogram editor

The spectrogram editor is the heart of the wave analyzing part of CUBE. Here you can have different views on the analyzed sample. There are three types of data you can view and edit here selected with the "display selection" field on the upper left:

- **Sine spectrum**

  The sine spectrum shows the temporal evolution of the frequency and level of each partial. The partials are shown as lines in the time/frequency plane. The level of the partial is shown as color ranging from yellow (high level) to blue (low level).

- **Residual spectrum**

  The residual spectrum represents the noise part of a sound and shows the temporal evolution of the level of each of fixed 256 frequency bands each with a bandwidth of approx. 80 Hz. The noise energy level for each band is shown as color ranging from yellow (high level) to blue (low level).

- **Pan spectrum**

  The pan spectrum shows the temporal evolution of the pan position of each partial. The partials are shown as lines in the time/frequency plane. The actual pan position is shown as color ranging from red (left) over grey (center) to blue (right).

The time ruler on the lower side and the frequency ruler on the right site of the display window can be used to scroll and zoom the spectrum display simultaneously: just click into the rule and move it up/down to zoom and right/left to scroll the window.

The most interesting aspect of these displays are the drawing functions you can use to manipulate the spectrum in a musically meaningful way. The following pages describe the tools you can use to "draw" your own sound creations.

5.3.1 Spectrum tools

You can select one of four main drawing tools to manipulate the spectrum from the toolbox above:

With the "Filter resonance" tool you can draw a single bandpass filter into the spectrum. The paint parameter are used to specify the time smearing, the filter bandwith and the filter level. Negative filter levels results in a bandstop.

With the "Comb filter" tool you can draw a bunch of bandpass filters at once into the spectrum. The paint parameter are used to specify the time smearing, the filter distance and the filter level. Negative filter levels works as expected.

With the "Frequency contrast" tool you can enhance or reduce the amount of frequency modulations in the spectrum. The paint parameters T-Width and B-Width are used to specify the area which gets influenced, a positive factor enhances the modulations and negative factors reduce the modulations.

Works like the frequency contrast but this time on the Level of the spectrum.

The way the two filter tools work is also influenced by the merge/add option:

Merge means that the filter is not added and comes only in effect if the filter level is above the current level in the spectrum.

Add means that the filter is added into the current levels in the spectrum and therefore always affects it independent of the level settings of the filter.

All the drawing functions above can be used for the sine spectrum, the residual spectrum and the pan spectrum in the same way.

The actual drawing can be done in two different ways:

- click where you want to start the filter and drag the mouse for arbitrary filter evolution.

- Shift+click on a start point and release the mouse at the end point of a virtual line. The filter effect is then drawn in the same way as if you had drawn a perfect line between these two points.

After you have done some drawings this way you may want to normalize the level of the complete spectrum to get maximum level without distortion. This can be done with the "Normalize" button.

CUBE provides you also with a sophisticated UNDO function for the drawing operations. The UNDO function is managed by the buttons below:

⇦ ★ ⇨

The "*" button saves the current state into the UNDO buffer. You can use up to 99 states at once.

The left arrow restores the latest saved state.

The right arrow let you redo the last undone state.

5.4   Time warp envelopes

The time warp envelopes control the way the analyzed sample is scanned. They have the same functionality and user interface like the Morph X and Y envelopes. They can have up to 64 time/position segments which are also syncable to the MIDI clock. Right after loading analysis data or importing a sample file this envelope is reset to a default setting. With this default setting the resynthesized sound is played in a natural way for the start of the sample to the end in the original speed. All sounds are also automatically looped with this default envelope.
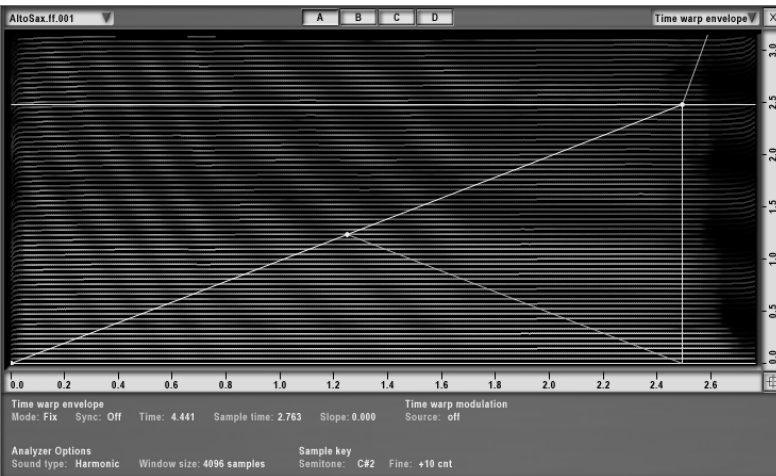
But how does this envelope work ?

The horizontal x-axis represents the relative position or time in the analyzed sample. The y-axis is the real time axis. The straight line from the bottom left edge to the top right means that the analyzed sample is scanned from the beginning to the end. If you stretch the envelope in vertical direction the analyzed data is scanned more slowly: This works like time stretching but this one is without changing the pitch of the resynthesized sound ! And in contrast to other techniques this time stretching is done very naturally without audible artifacts. Another example is playing the analysis data backwards in time: To achieve this just let the envelope go from upper left to lower right ! And now try to imagine what you can do with 64 syncable segments…

It is also possible to modulate the scan position with different modulation sources like the modulation wheel. One application of this modulation is scratching through your analysis data: this is like scratching vinyl but without pitch change. Try for example importing a simple loop sample and change the envelope like sketched below and assign the modulation wheel as time warp modulator.
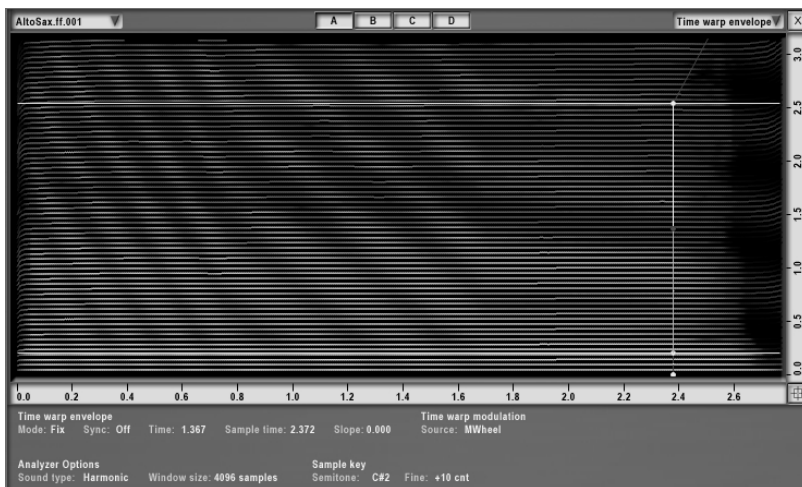
5.5   Simple Tutorial

This tutorial gives you a step by step instruction for importing and resynthesizing your first sample.
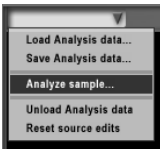
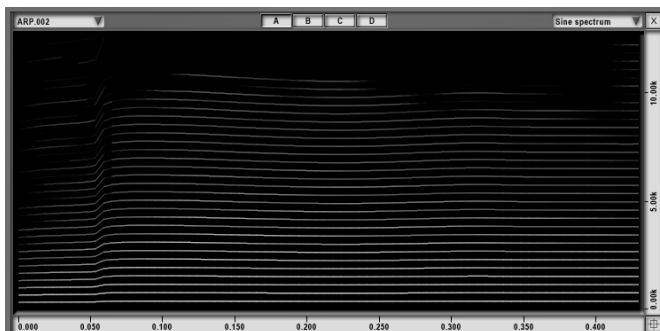1.   Start with an empty preset by selecting "Clear preset" from the file menu:
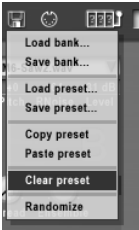
2.   Click on the "Wave" button of Source A to open the Wave analyzer window:

3.   Select "Analyze sample…" from the Wave import/load menu and select the sample file you want to analyze:

4.   After a short while of processing you should see the spectrogram of the analyzed Sample:

5.   OK, now you should be able to test the imported sound !
Now we come to the fine tuning.

6.   Click on the "A" or the close button to get back to the main window. Now you can see the volume and morph envelopes instead of the spectrum display. If you want to get exactly the same volume envelope as the original sample, you have to modify the default envelope to a "gated" one: minimum attack and sustain with maximum level.

7.   With an empty preset, the number of partials is set to 32. This is usually not enough for a realistic result. Try higher values up to 512 and you will notice that the sound quality and characteristic is changing a lot with this value, but keep in mid that this also uses more CPU power.

6 Effects

6.1 Overview

If you click on the FX button on CUBE's window you'll get the effects page:

The effect section of CUBE contains one global Reverb and three part effects which can have different settings for each part.

- Global reverb
  The global reverb is shared by all parts and exists only once. You can use it for all 8 Parts. For every part you can only adjust the effect send level for the reverb. The reverb effect output is added to the main output. The other reverb settings are saved in the sound bank and not individually for each preset !

- Part effects
  Each part have its own part effects Overdrive/Delay/Modulation. These are designed as insert effects, the input coming from the additive synth engine goes through each of them in sequence and the output of the last effect (Modulation) is routed to the channel output of the part. The part effect settings are stored together with the sound program.

6.2   Reverb

The reverberation effect is used to simulate the sound reflections off walls in different kind of rooms. This reflections disperse the sound, enriching it by overlapping the sound with its reflection. The reverb effect is global, this means only one reverberation type can be used at one time, but every part has an individual amount of reverberation set in the program. The effect parameters can be set individually, or be selected from a list of 32 predefined presets.

-     ON/OFF
      Switches the reverb on and off. If the reverb effect is not used, it should
      be switched off to lower the processor usage.

-     Preset
      Selects from 32 global presets.

-     Delay
      Sets the amount of time that elapses between receiving a direct sound
      and its first reflection. This time correlates with the perceived size of the
      simulated room. The adjustable range is from 0 ms to 100 ms.

-     Time
      The reverberation time indicates the time required for a sound to die
      away to –60 dB of its amplitude after its source shuts off. This time
      correlates with the perceived size of the simulated room. The adjustable
      range is from 50 ms to 330 seconds. In the uppermost position, labeled
      „endless" the sound never dies out.

-     Size
      The room size parameter sets the average distance between the
      reflecting walls. The adjustable range is from 1 meter to 100 meters.

-    Damping
     The damping parameter simulates the different absorption of high
     frequencies by the wall material. The adjustable range is from 0 to 1.
     Due to damping the reverberation Time (as defined page 39) is shorter
     for high frequencies than for low frequencies. As higher the damping as
     faster the high frequencies will decay.

6.3 Overdrive

The part overdrive effect provides you with different types of distortion from hard clipping to soft distortion simulating old tube amplifiers.The sound balance can be adjusted using a combined High/Lowpass filter.

- ON/OFF
  Switches the part overdrive on and off.

- Drive
  Sets the amount of distortion by amplifying the input signal.
  Range is from -12 dB to +48 dB.

- Cutoff
  This parameter sets the crossover frequency of the distortion filter. The distortion filter consists of a lowpass filter and a highpass filter with a cutoff frequency equal to the crossover frequency. The relative amount of lowpass and highpass filtered signal power is controlled by the Tone parameter.

- Tone
  This parameter controls the relative amount of lowpass and highpass filtered signal power. The cutoff frequency of both the lowpass and highpass is controlled by the Cutoff parameter.

- Select Overdrive type
  Selects between four distortion characteristics:
  hard, soft, tape and tube

6.4 Delay / Echo

The part echo/delay effect provides you with three different types of delay effects using two delay lines. The delay times can be synchronized to an external MIDI clock. The delayed signal can be feed back to the input for echo effects. The delay time can be synchronized to the MIDI clock or to the internal arpeggiator tempo to realize rhythmic echoes.

- ON/OFF
  Switches the part echo/delay effect on and off.

- MIDI
  Switches the synchronization to the MIDI clock on and off.

- Delay 1
  Sets the delay time of delay line 1 ranging from 0 ms to 728 ms. If MIDI sync is activated the range is from 1/32 triplet to 1/4.

- Delay 2
  Sets the delay time of delay line 2 ranging from 0 ms to 728 ms. If MIDI sync is activated the range is from 1/32 triplet to 1/4.

- Feedback
  This parameter controls the feedback of the delay lines. This controls the decay of the amplitude of the echoed sound.

- Cutoff
  A lowpass filter is build into the feedback loop of the delay. This parameter controls the cutoff frequency of this feedback filter. This results in successive echoes sounding darker.

-       Level
        The level parameter controls the output level of the effect. The
        adjustable range is from 0 to 1.

-       Select effect type
        With this listbox you can select between three different delay line
        configurations:

        Stereo delay
        Both delay lines are used independently for the left and right input.

        Mono delay
        The delay lines are connected in series for monophonic dual tap delay
        effects.

        Cross delay
        The feedback loops of the two delay lines are cross coupled. This
        means the delayed sound bounces between the stereo channels.

6.5   Modulation effect

The part modulation effect provides you with three different types of modulation effects:

- Chorus
  The Chorus effect have the same structure as the global Chorus effect: four Delays modulated by four independent LFOs.

- Flanger
  The Flanger is composed of two independent delay lines with feedback for the left and the right channel respectively. The delay time of both delays is modulated by one LFO with adjustable frequency. For typical flanging effects the delay time should be below 5 milliseconds.

- Phaser
  The Phaser is built of an 8-pole allpass filter, which gives you a very short and frequency dependent delay time. This delay time is modulated by an LFO.

- ON/OFF
  Switches the part modulation effect on and off.

- Select modulation effect type
  With this listbox you can select between the three modulation effects: Chorus, Flanger and Phaser.

- Speed
  Sets the average speed (frequency) of the LFOs modulating the delay time. The adjustable range is from 0.011 Hz to 75 Hz (cycles/second).

- Delay
  This parameter sets the average delay time of the four delay lines. The adjustable range is from 0.1 ms to 93 ms for the chorus and the flanger effect. For the phaser effect the delay time is considerably smaller and depends on the frequency spectrum of the input signal. The value for a typical chorus effect should be between 10 ms and 20 ms. For a typical flanging effect the value should be below 5 ms.

- Depth
  This parameter controls the depth of the delay time modulation by the LFOs. The adjustable range is from 0 to 1.

- Feedback
  The feedback parameter controls the amount of positive or negative feedback for all four delay lines. The adjustable range is from -1 to 1.

- Level
  Sets the output level of the part modulation effect. The adjustable range is from 0 to 1. With a value of 1 the output is a 1:1 mix of the original and the effect sound.