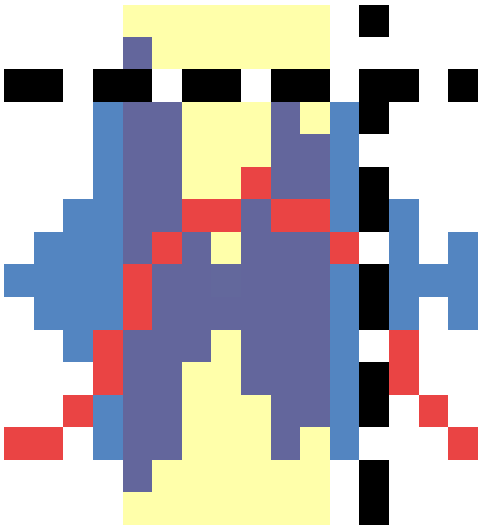

The LAoE tutorial serie - How to use layers

© 2002 oli4, Neuchâtel Switzerland
first edition, april 2002

author: Olivier Gäumann
rue Matile 6
2000 Neuchâtel
Switzerland

e-mail: laoe@oli4.ch
web: www.oli4.ch/laoe

wordprocessor: OpenOffice 641

First of all, excuse my awfull english, it's not my native language! The author is the owner of this document. **It is allowed to copy this document or a part of it**, but with the restriction to keep the name of the author and this copyright text on each copy.

LAoE is a rich featured graphical audiosample-editor, based on multi-layers, floating-point samples, volume-masks, variable selection-intensity, and many plugins suitable to manipulate sound, such as filtering, retouching, resampling, graphical spectrogram editing by brushes and rectangles, sample-curve editing by freehand-pen and spline and other interpolation curves, effects like reverb, echo, compress, expand, pitch-shift, time-stretch, and much more...

Table of Contents

1.Introduction 3
2.Why working with layers? 3
3.The clip-layer-channel concept of LAoE 3
4.Manipulate layers 3
Create new layers 4
Copy, paste, duplicate and delete 4
Move layers up and down 4
Merge layers 5
5.The same is valid for channels 5
Manipulate channels 5
Convert from stereo to mono 5
Convert from mono to stereo 6
Why using more than 2 channels 6
6.Using masks 6
What is a mask? 6
Edit a mask 6
When using masks? 7
7.The layer-plot 7
View-options 8
8.Working as with multiple audio tracks 8
Copy, perform effect and mask 9
Keep layer-information separated 9
9.Working with parameter-layers 10
What is a layer-type 10
Edit a parameter-curve 10
Envelope 11
Resampling with doppler-effect 13
Convert progressively to mono 14
Other variable effects 15
10.Making superposing measures 15
Finding clicks and pops 15
Show a RMS mean curve of the clip 16
Measure the transfer-function of a filter 16

1. Introduction

The LAoE tutorial serie is a serie of articles focussed to LAoE-users. These articles have been written with the idea to bring LAoE to a larger public, because LAoE isn't easy to use without documentation. At the time of writing this article, no documentation was existing about LAoE.

In this article, we will learn the basic concept of LAoE, the clip-layer-channel concept, what's the purpose of the layers, how to manipulate them, how to manipulate channels too, using masks, some effects, like amplitude, math, spectrum analysis... etc. We will look around, not focussing exclusively on layers, but also largely the domains around, the ones who get in touch with layers. This article is based on the current version of LAoE v0.4.05. So when you read this lines, maybe a newer version of LAoE will be out, and the graphical user interface will probably have minor differences to the figures of this document.

2. Why working with layers?

There are many reasons why working with layers. Layers are aequivalent to "tracks" in a multitrack music editor. You can separate the sound in independent parts, which are mixed together when playing, but keeping separate when editing. You can create free editable parameter-curves to use with effects that allows variable parameters. You can extend layers for calculation and analysis. You can exchange whole layers and channels between several audioclips. And if you've read this article, you'll see, LAoE is great...!

3. The clip-layer-channel concept of LAoE

After launching LAoE, the small LAoE-main-frame appears, it contains many buttons to start the most common plugins. In LAoE, most of the functionality like open-dialog , save, select, zoom, amplify... etc. is called a plugin. Open a new clip with the help of the open-dialog , and then open the layer-stack. The clip and the layer-stack plugin are visible in separate window each. Now you should see the windows like in the figure below.

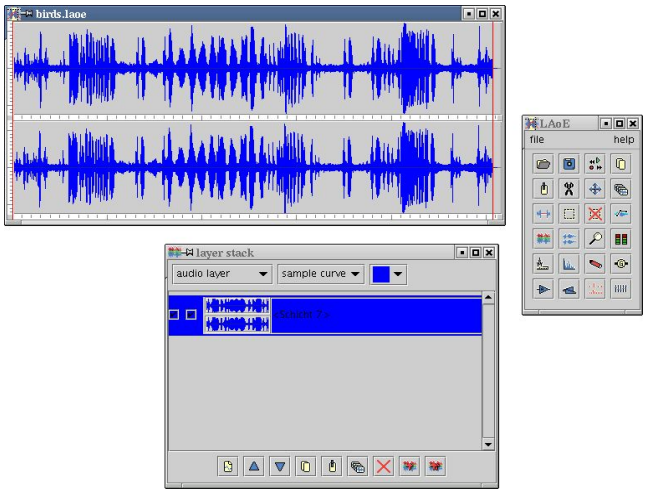

In the example, birds.laoe is a stereo-clip, it contains two channels. Each channel is output separately on the audio-system (sound-card). A mono-clip contains only one channel. Most of the audio-systems support only one or two channels. LAoE has no limitations of number of channels for editing purposes, but it can only play the number of channels your audio-system supports. Now what is a layer? The clip of this example contains only one layer. If you open an ordinary soundfile like a wavefile (.wav) you will only have one layer. This layer contains the two channels. A clip may have an unlimited number of layers. Each layer contains its own audio-recording, this allows to keep different recording parts separately, superposed, like in a multitrack system. All audible layers are mixed together when playing. Since the layers must not contain the same numbers of channels, channel 1 of layer 1 is mixed with channel 1 of layer 2. Channel 2 of layer 1 is mixed with channel 2 of layer 2, and so on. If layer 2 contains only one channel, channel 2 is not mixed at all.

A clip contains one or more layers, and each layer contain one or more channels. A clip with zero layers doesn't exist, a layer with zero channels too. A clip is the whole sound-entity, a layer is a mixed part of the sound, a channel is a separately processed sound-in/output.

4. Manipulate layers

The sample-curves of a clip with multiple layers are painted one on the top of each other in the clip-frame. You can differentiate them by different colors. The layer-stack plugin shows another view: each layer is painted separately in the layer-stack. The position of the layers in the layer-stack has no signification. All of them have the same importance or priority. When clicking on a layer, this will become the selected layer, recognizable in blue. The selected layer is always painted on top in the clip-frame. The layer-type, layer-plot and layer-color on the

top of the plugin and the buttons on the bottom are always related to the selected layer. The layers can be manipulated as a whole part in the layer-stack plugin.

Create new layers

Click on new layer to create a new layer. A dialog appears, you can modify the number of channels and the sample length of the channels. It proposes the same values as the existing layers.

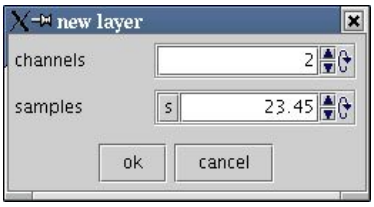

After accepting, you have a clip with two layers, the original one <Schicht 7> and the new empty layer <layer 5>. You may change its name, type any text in the layer-name field and type return.

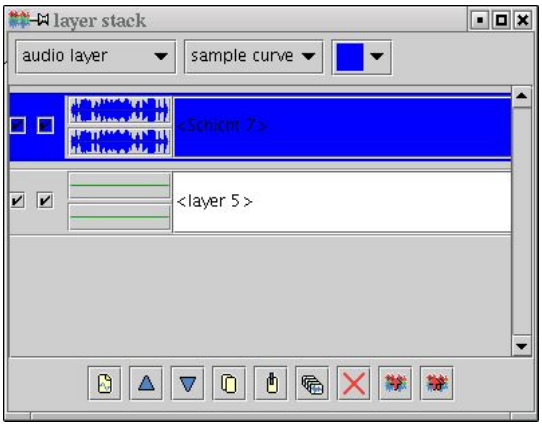

This new layer is now ready to use. At the time it contains zero-data, silence from begin to end. So when playing, there will be no difference compared to the original clip.

Copy, paste, duplicate and delete

Copy and paste layer works with a special clipboard of the layer-stack. First select a layer by mouse, then click the copy button. When pasting a layer, the layer is placed at the location of the actually selected layer, but without replacing it. The layers may be copied/pasted between different clips. Duplicate layers works inside a single clip, and it is a shortcut for copy and paste. The delete button removes the selected layer, except if there is only one layer, this will remain. The next figures shows an example, where a second clip has been opened, a mono audio clip (green). The layer of the second clip has been copied and pasted into the first clip birds.laoe. Since the source was a mono-clip, only one channel has been pasted (don't worry, you may duplicate it to both stereo-channels with the help of the channel-stack).

Pay attention to the fact, that copying layers between different clips works fine, if both clips have the same samplerate and samplewidth. If not, the samplerate can only be adjusted to one of both samplerate-settings, and the data of the smaller samplewidth will be too small to see (or hear). If the samplerates and samplewidths are different, it is still possible to convert them later, each layer separately.

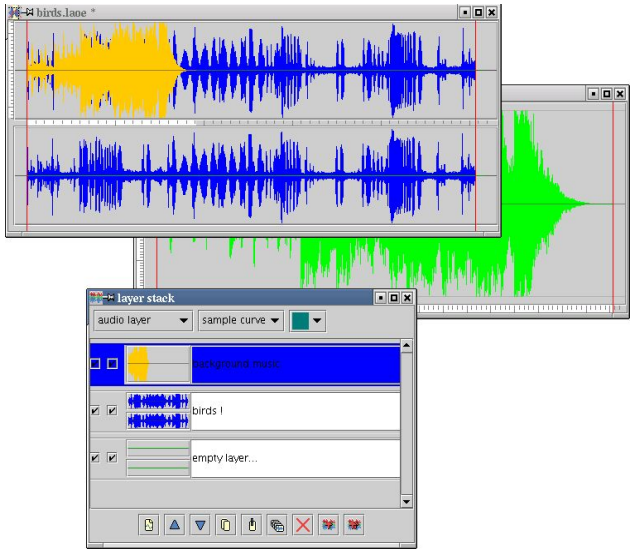

Move layers up and down

The selected layer may be moved up and down in the layer-stack. This has no effect to the hearable result, it is only a graphical feature to re-arrange the layers.

Merge layers

If you have two or more layers, and you want to mix them together (e.g. to save the clip in a standard format like a wavefile), you can do it with the merge-buttons. The merge up layer button merges the selected layer with the upper layer in the layer-stack, so the two layers are replaced by one single layer. All other layers remain untouched. The merge all layers button merges all layers of the clip into one single unique layer. When merging layers, the samples are mixed together, and no normalisation and no saturation is performed. That means, no sound-information is lost, but the sample

volume may be louder than the allowed sample-width. To avoid this, check if you have clipped samples with the measure-plugin or normalize the merged layer with the amplify-plugin.

5. The same is valid for channels

As mentioned above, there is a channel-stack. This is an aequivalent tool for manipulating channels. Channels may be created, copied, pasted, duplicated, deleted and merged too. To select a certain channel x in the layer y, you have first to select the layer y in the layer-stack, and then go back to the channel-stack. This will be updated, the channels of the selected layer are shown only. Now you can select the channel x by clicking with the mouse. The position of the channels in the channel-stack has the following signification in a stereo clip: the top channel is the left stereo-channel, the bottom channel is the right stereo-channel.

In the example below, the <channel 1> has been duplicated. So the background music will be centered on both channels of the stereo clip, but still in mono, since both channels contain exactly the same samples.

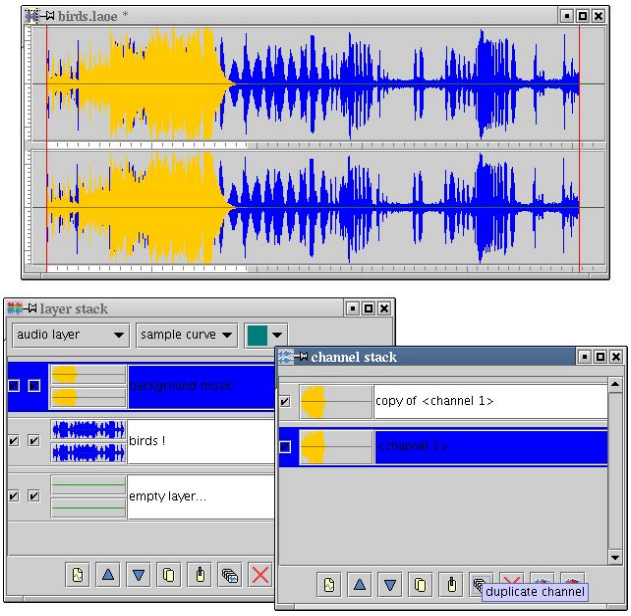

Manipulate channels

The most common channel manipulations are conversions from mono to stereo or inverse. Layer manipulations are more often used than channel manipulations.

Convert from stereo to mono

To convert a stereo-clip to mono, open the channel-stack and merge the two channels to one. Now the volume may be too loud, so use the amplify-plugin to adapt the volume, e.g. normalize the channel's volume. Both origin channels have been mixed in equal parts on merging, so sound-components of both channels will be heared. By the way: there is a special plugin narrow-wide, which allows to mix the two stereo-channels in a way it sounds wider or narrower, in a constant or variable manner. The figure below shows the relevant plugins for converting to mono.

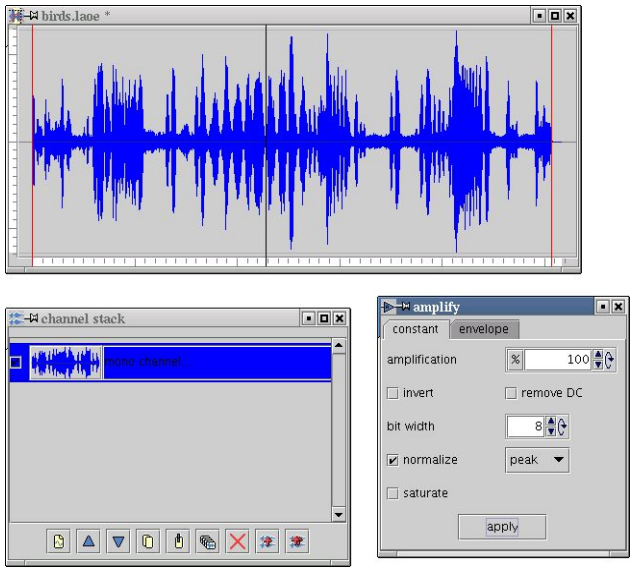

Convert from mono to stereo

First of all, duplicate the unique channel. Now you have two channels in the concerned layer. This is not enough, because two identical channels do not give the effect of a stereo sound. Two methods are possible in LAoE to generate a stereo sound. The first consists of delaying one channel compared to an other about 10 to 50 miliseconds. This gives a little bit space in the sound. To do this, select a range of this length at the beginning of one channel only, and cut this range. The second method is to use the narrow-wide plugin and widening these two channels.

Why using more than 2 channels

It is rarely needed to use more than two channels. Maybe in future, if multi-channel soundsystem like surround-sound or quadrophony is used. At the time multiple channels may be used for editing purposes, like copy-paste, or user-defined mixing of channels. As mentioned above, in theory LAoE supports any number of channels,

but the soundsystem doesn't.

6. Using masks

Have you ever seen a multi-track music-editor, where each track has an editable segmented volume-curve (envelope-curve)? The mask-concept in LAoE is here for the same purpose.

What is a mask?

A mask (or better a volume-mask) is a kind of logical envelope-curve, which defines the output-volume in function of time, which is applied to a channel, without changing the original samples. This mask is applied when playing the channel and when merging the channels/layers. This curve consists of linear segments, editable by mouse. The value of the most left point continue to apply the samples until the left end of the channel, the same is valid for the right side. The vertical scale is defined as follows: 0% (silence) at the visible bottom side of the channel, 100% (original volume) at the visible top side of the channel. So the vertical scaling is independent of the zooming of the samples, and the volume range is limited from 0% to 100%. The example below has a mask which fades in in linear manner, then (in the middle part) the volume is almost full, decreasing very few, and at the end fades out in an exponential way.

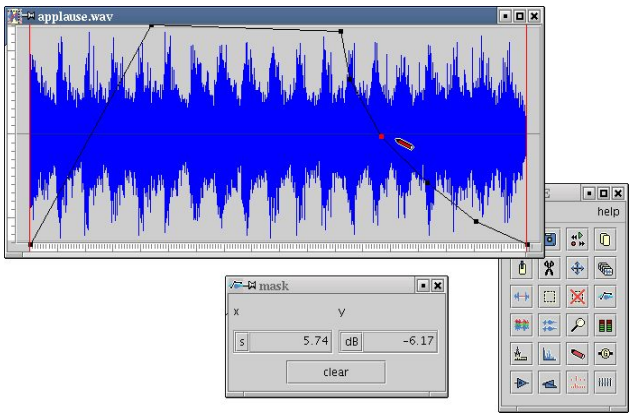

Edit a mask

Start the mask-plugin and the pen-cursor appears. Now click on the channel, so a point appears at the mouse's position. Click a second and third time, and more points and line-segments appear. Each point is movable by mouse when dragging and removable when clicking with pressed shift-key. The current selected point is painted in red color and its coordinate values are displayed in the mask-plugin frame.

If you have a stereo-clip and you want the same mask on both channels, you will need to copy-paste the mask, avoiding to redraw two similar masks. The mask of the last channel, which was touched by mouse, is copied to a special clipboard using copy mask plugin. Now touch the channel, where the mask should be copied in, and activate the paste mask plugin. If you want to remove all masks of a layer at once, activate the unmask layer plugin. This is a short-cut, it avoids to remove each point of each mask individually.

When using masks?

In general, the masks are useful if volume should be modified multiple times, because the quality doesn't decrease since no sample-information is lost. Typical applications where masks are helpful:

- background music with variable volume in function of the foreground sound (simply change background mask when adapt changes to the foreground sound)
- introduce an effect in a duplicated layer, and let the freedom to mix original and modified layer where you want and how you want

7. The layer-plot

LAoE supports two different plot-types, the classical sample-curve plot and the spectrogram plot. The approach of these views are very different. The sample curve plot shows the audio samplepoints in form of a curve in function of time. Louder parts have a bigger vertical variation than silence parts, which normally around the zero-line. This is the classical form found in most of the audioeditors. The figure below shows the sample-curve plot.

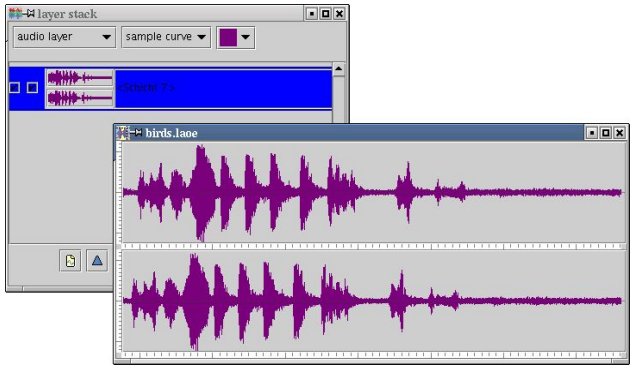

It is the default plot in LAoE. It is the only plot capable of plotting multiple layers (one of the top of the other). The layer-color is choosable in the layer-stack plugin. There are just a few colors, and some main colors are not

supported, to avoid color-conflicts with other graphical components like masks, selections, loop-points... etc. The sample-curve plot is zoomable in vertical and horizontal axis, and is optimized to zoom very deep into the sample-dimension. If you have to edit samples by samples, the sample-curve plot is ideal. Most of the plugins work with this plot.

The spectrogram plot is very different. The horizontal is still the time-axis, the vertical is the frequency-axis and the color-intensity shows the amplitude. You see a kind of spectrum that varies in function of time. Dark color means big amplitude, few color means little amplitude. Dark color at lower location means the sound has lower frequency-parts here, dark color at upper location means the sound has higher frequency-parts here. This view shows the timbre of the sound variing in time. It is not very precis when zooming deep in time, it is not as exact as sample-curve view and you cannot edit sample by sample. But it shows different components of the sound separately, even if they appear at the same time. As an example, you see the human voice separately from music, or a bird separately from another. This is very interesting, it would never be possible in the sample-curve plot. LAoE has a spectrogram-filter, where these different components may be isolated from the rest or filtered out.

Look at the figure below, where you distinguish three or more different birds, one of them has a very low frequency range, the others share the same frequency range, but have different timbre-shapes. This figure shows the same soundclip as the figure above. It's interesting to compare both plot-types.

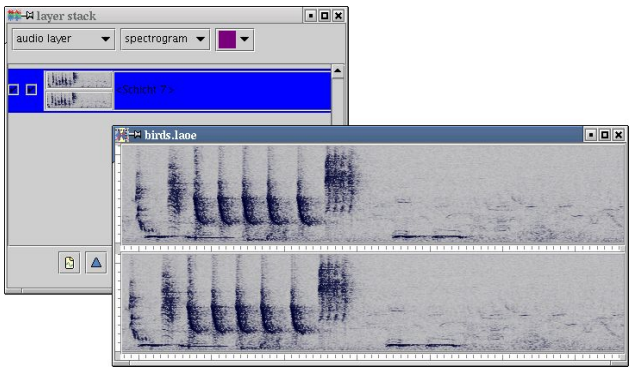

The spectrogram plot is also zoomable in horizontal and vertical axis, but the resolution is limited, but configurable (see the view-options plugin). This plot exists in many soundeditors, but most of them don't allow zooming in both axis or editing. There are just a view plugins which require the spectrogram plot. The spectrogram plot has no transparency and fills the whole channel-area, so all channels/layers below are not visible in the clip-frame.

View-options

The unit of the vertical and horizontal scale and of the selection is configurable with the view-options plugin. In addition, the spectrogram plot has some settings, like the FFT-length and window. Spectrogram-resolutions are always a compromise between frequency-resolution and time-resolution. You cannot have both at the same time. You may have better frequency-resolution and lower time-resolution, or in opposite, better time-resolution and lower frequency-resolution. Vary a little bit the FFT-length and compare the plot results. The longer the FFT-length, the slower the plot-update because the amount of calculations increase very fast. In general, the Hamming-window is the best, but try also the Blackman-window. Rectangular window is very bad.

The brightness setting of the colors may help a little, when the amplitudes are very small or very big. The color-scale begins at zero and ends at the maximum value of the clip (related to the sample-width). The brightness setting of 50% generates a linear relation between value and color. When greater than 50%, the color gets darker very quickly, when lower than 50%, it get darker very late in the value-range.

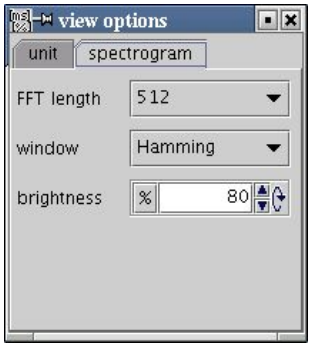

8. Working as with multiple audio tracks

Lets look at an example of a multilayered clip with 3 layers containing differend sound-components. The music-layer uses a mask to reduce the volume when the speaker says hello! in the hellooo-layer, and fades out at the end of clip.

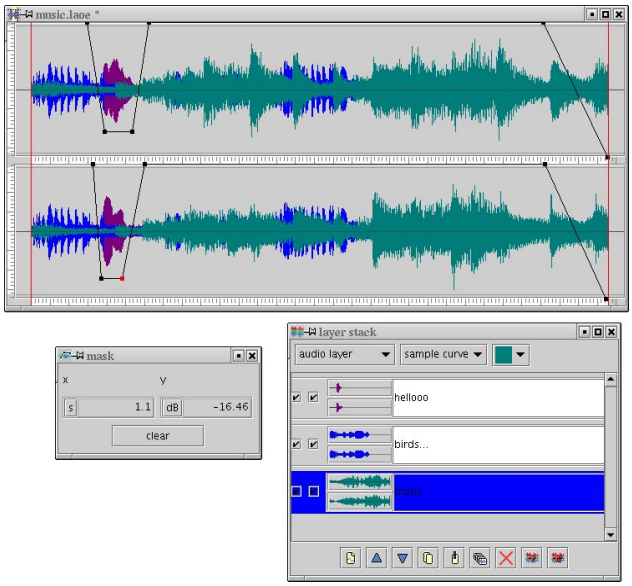

The birds-layer also uses a mask to fade in at beginning, increasing volume and fade out at the end. The hellooo-layer doesn't use any mask, so the volume keeps 100% all the time.

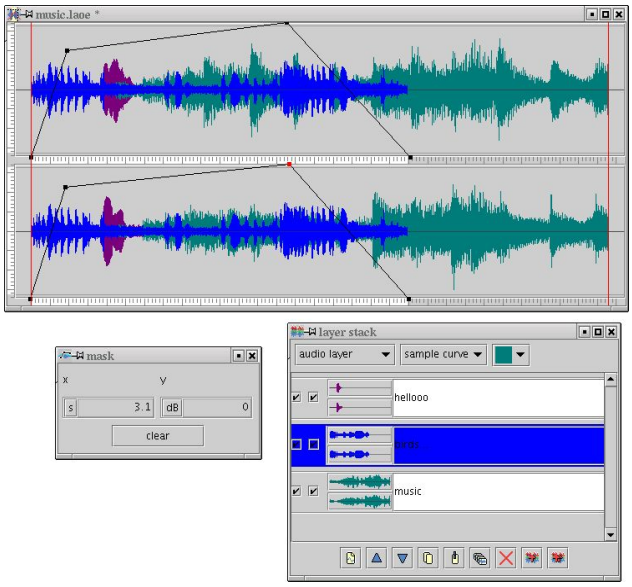

All 3 layers are painted one of the top of the other, painting the selected layer on the top. The layer-stack represents them separately, where the layer-name is editable, with the visible- and audible-checkbox. All layers are audio-layers (audible), painted as sample-curves, each in a different color. Each layer is editable separately, manipulations and effects may be performed individually, and the 3 layers are played as a common mixed result. If you want to hear an individual layer, play with the audible-checkboxes, enable only the concerned layer, disable all others. If you want to see only one layer, play with the visible-checkboxes.

Copy, perform effect and mask

This example uses layers because of another reason: an effect will be performed on the original clip, but timely limited, only in a certain time-range. This time-range is not well-defined, and we want to finetune this range later. The easiest way to go is to duplicate the original layer, so we get two identical layers. Now the effect (e.g. reverb) is performed on one layer only. So we have the original layer and a layer with the effect. This allows to switch from one layer to the other at any time-position. This is done with the help of the masks. We can define the mask on both layers, one mask is the complement mask of the other. At the places we want to hear the original sound, the original layer is masked at 100% and the effect-layer at 0%.

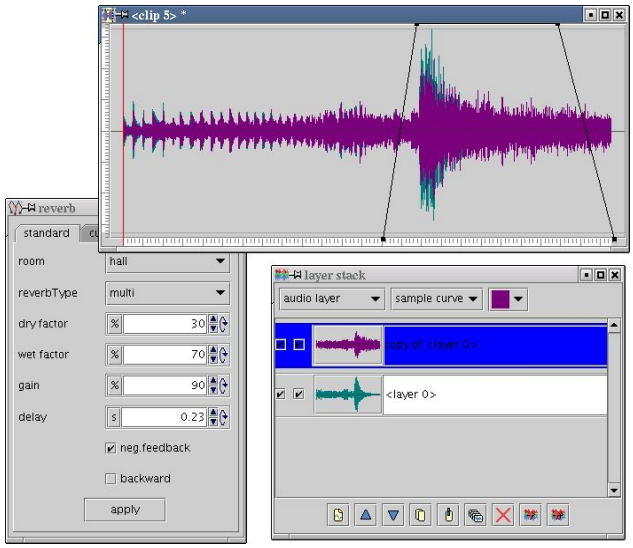

And at the places we want to hear the effect, the effect-layer is masked at 100% and the original layer at 0%. The transition from original to effect may be of any shape, freely editable graphically. Here we have a smooth linear transition.

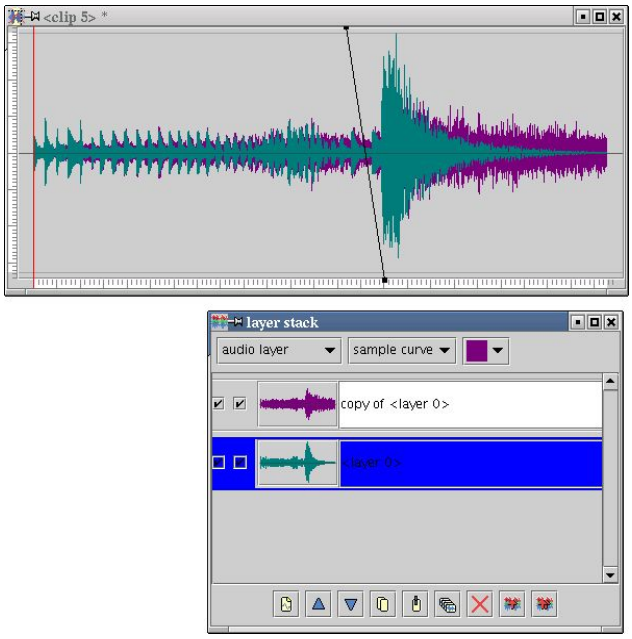

If the range of effect is not satisfactory, we can reshape the masks without changing the precious sample-informations. Instead of switching from one layer to the other, we could imagine a mix of both with any variable factor.

Keep layer-information separated

At the time, only the .laoe-fileformat saves the layer-informations, so if you want to keep the layers separately saved, use the native .laoe-fileformat. At the other hand, if you want to save the result in a standard-fileformat as . wav, merge all layers before saving. But don't forget, the layer-separation will be lost, only one layer with all sound-components mixed will remain.

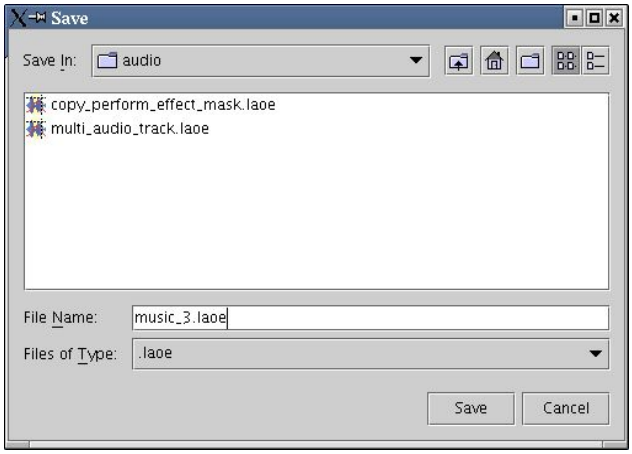

9. Working with parameter-layers

What is a layer-type

Layers may contain other stuff than audio-information, layers may contain parameter-curves. Parameter-curves are very useful when performing variable effects over time, e.g. variable resampling to create a doppler-effect. LAoE supports two different layer-types: audio-layers and parameter-layers. Audio-layers are audible, parameter-layers are not audible. Audio-layers are painted fully, but only the borders of parameter-layers are painted, they are transparent inside. There is no other difference between the two layer-types. Both are editable the same way, effects may be performed the same way on both layer-types. The figure below shows a clip with two different layer-types, an audio-layer and a parameter-layer.

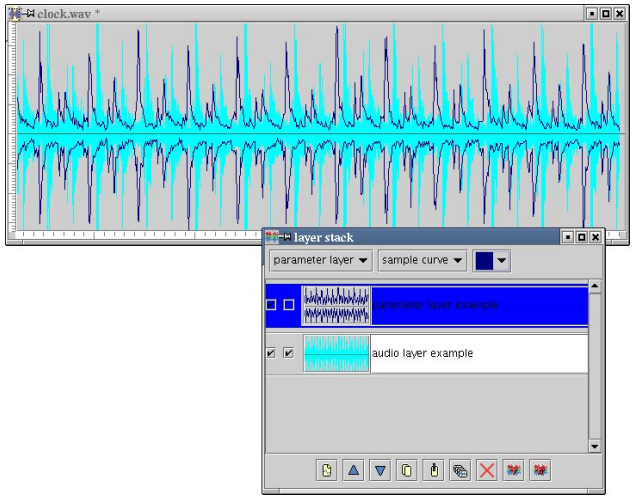

Edit a parameter-curve

So let's edit a parameter-curve. Many tools are provided to edit a parameter-curve. We begin with generating a constant curve with the value of 1. To do this, we use the signal-generator plugin.

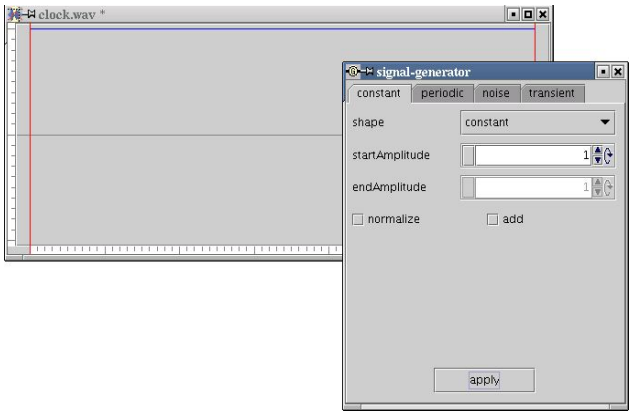

Now we want to add a section, where this curve will grow a little bit in linear segments, let's say to the value of 1.2. We create this segments using the free-generator plugin, so we can work easily with the mouse. The corners of the segments are draggable by mouse, and the coordinate-values of the actual point are visible in the plugin-frame. You can snap the point to the existing line when pressing the ctrl-key. We need this feature for the two extreme points.

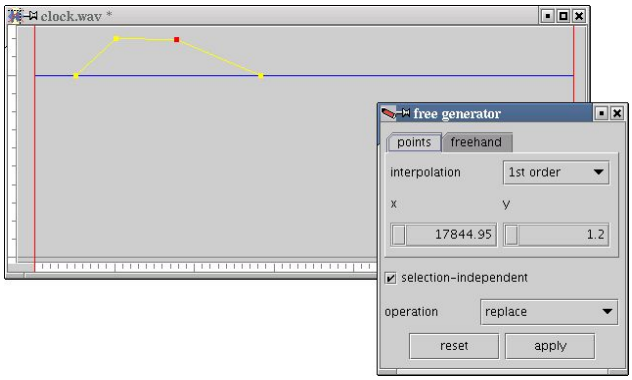

Once the segments are ok, we just have to apply them to the current parameter-layer. Now we want to shrink the value in another section to 0.8 using a freehand-shape. Just draw the freehand-line with the pen, with the freehand-tab selected in the plugin. To avoid too much aliases, we smooth the curve before applying.

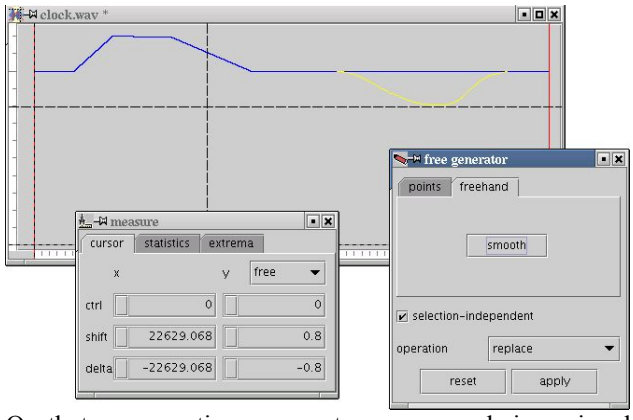

On that same section we want a superposed sinus-signal with an amplitude of 0.06 and a frequency of 3Hz. This may be done with the help of the signal-generator plugin again. Define the section with a selection, create smooth borders of that selection with the intensity-curve (this is possible if the intensity-tab in the selection-plugin is selected). So the sinus-amplitude will grow smoothly in the borders. Create the sinus-signal with the given parameters, but don't forget to click to the add-checkbox if you don't want this section completely replaced by the sinus at offset zero.

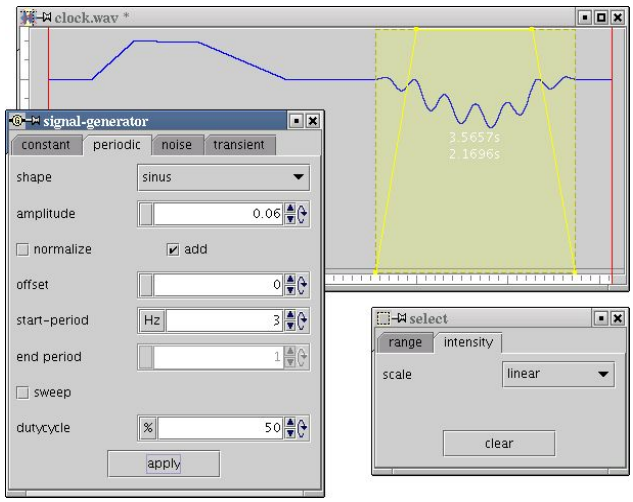

Let's say we need a second parameter-layer with a complementary curve of 1 - x. To do this, Add a new parameter-layer with a constant value of 1 (see above how to proceed with the signal-generator plugin).

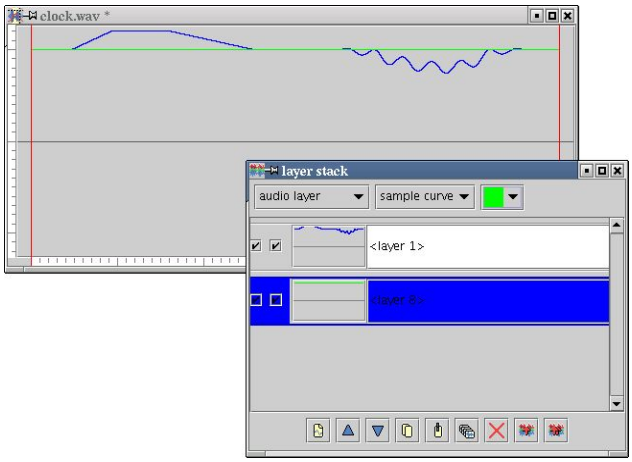

Select the layer with this constant curve and perform a substraction variable in time f(time) in the math plugin. The second operand must be defined in this plugin. The result is written into the selected layer (overwriting the first operand, the constant curve).

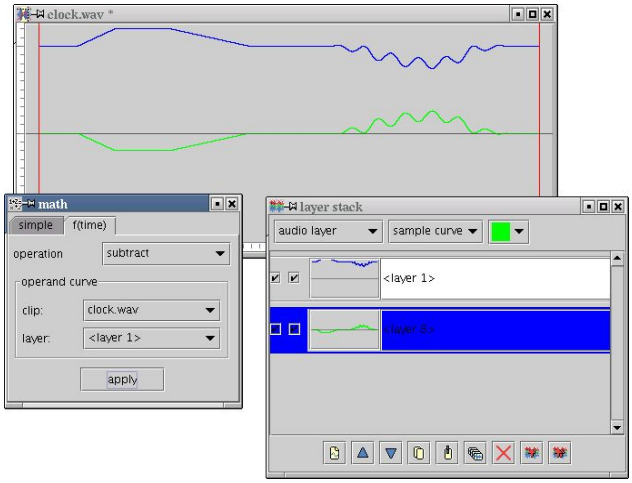

So we have created two non-trivial parameter-curves with the help of different plugins. Ok, this was just an exercise, but we could use them now to perform other operations and effects...

Envelope

A more practical application of the parameter-curve: we want to envelope definitely an audio-layer. We could simply use a mask, but since we want to perform the envelope definitely, we can use the amplify plugin and its envelope-functionality. First, we create a new layer, select it and perform a constant neutral value of 1 to this layer. When applying 1 to the envelope-function of the amplify plugin, the volume isn't changed (it's multiplied with 1). Once the constant value has been generated with the signal-generator plugin, we use the free-generator plugin to generate segmented lines. In this example, the fade-in at the beginning is linear, and the fade-out at the end contains multiple segments, to get a kind of exponential shape. To get more precisely, zoom in as much as possible to the used vertical range. The segments have been created on the first channel only.

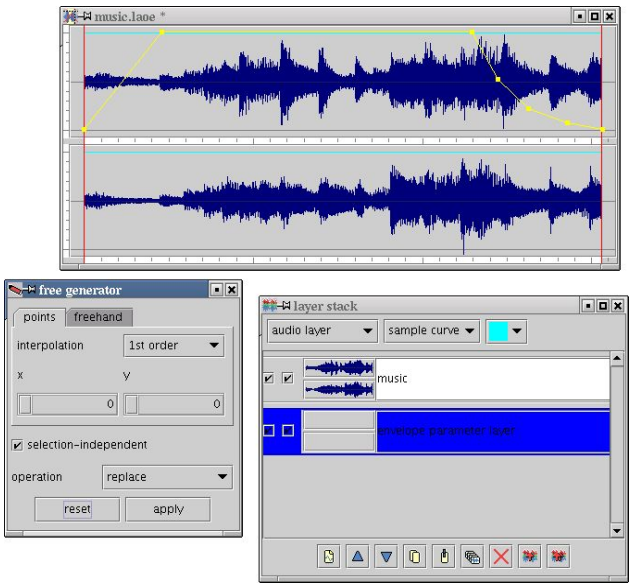

After applying the segments, and if we are happy with the result, we simply have to delete the untouched channel of the envelope parameter layer and duplicate the first channel, to get two identical channels. We have to provide the two channels, because LAoE performs envelope individually in each channel. We could use different shapes for the two channels, but we don't do it in this example.

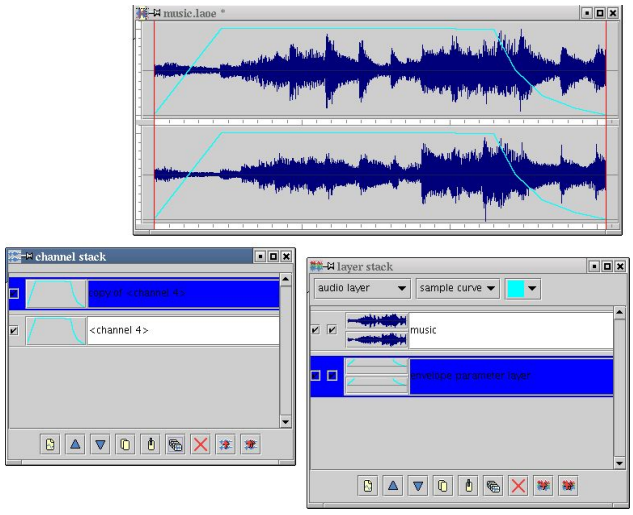

After creation of the envelope parameter layer, we select the music layer again, we continue to work on this original layer. Open the amplify plugin, select the

envelope-tab, and choose the envelope curve. Since this curve could come from any clip, we have to select the clip and the layer of the parameter curve. On applying, the envelope-function is performed on the selected layer (music), using the parameter layer choosen in the plugin. The envelope curve keeps untouched.

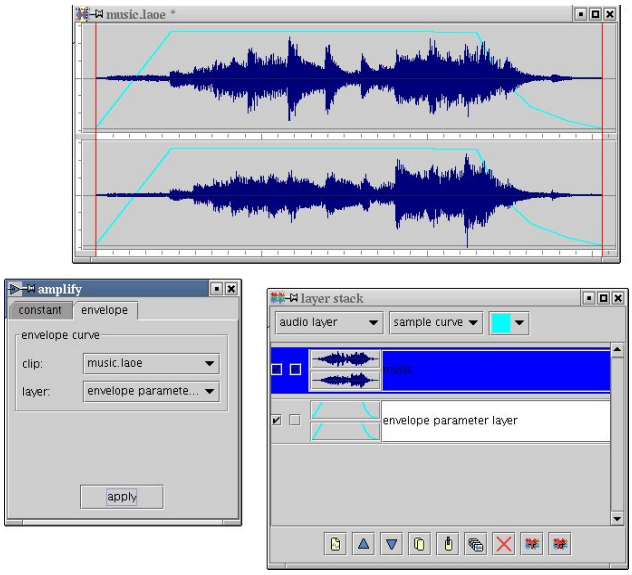

Look at the music layer. The envelope has been proceed, the samples of both ends are reduced progressively to zero. Enveloping is possible with any shape, reducing or increasing the volume. This would not be possible with the mask-concept, masks cannot increase volume.

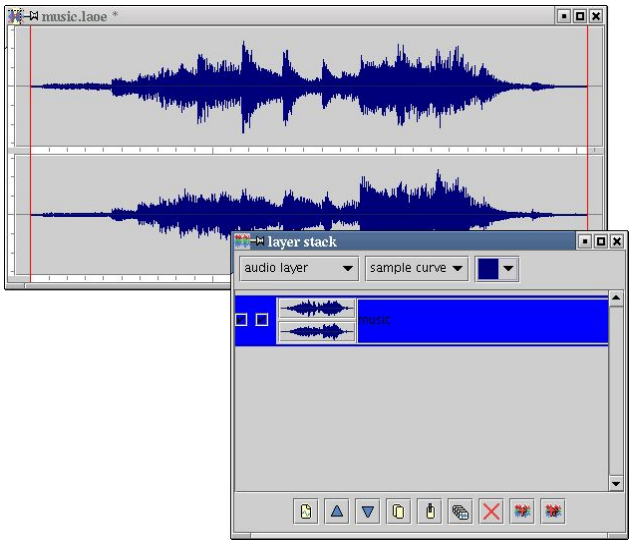

If you are happy with the result, you can remove the envelope parameter layer, you don't need it anymore. Don't forget you can go back any editing step if you're not happy or if you made mistakes. Use the undo-stack plugin for this purpose. The enveloping is done now.

Resampling with doppler-effect

We want to perform variable resampling, to get a kind of Doppler-effect on the sound. First a new layer is created. The signal-generator plugin is used to create a constant (neutral) value of 1. After that, we use the free generator to vary this constant curve near the value of 1, between 1.05 and 0.85. It is important to avoid extreme values, keep in the range of 0.1 .. 10. It is rarely needed to resample with bigger factors. To be as precisely as possible when working with the mouse, zoom as much as possible into the needed y range, setting individual y zooming. We have choosen a spline-curve in the free-generator, this gives a very smooth curve. The spline-curve is edited in the first channel only.

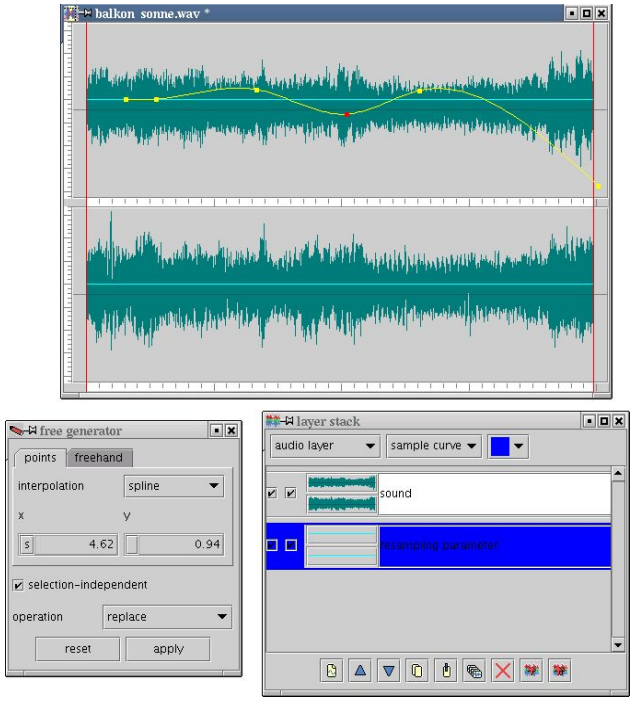

The second channel is just a duplification of the first. The parameter layer is now ready to be used for variable resampling. Open the resample plugin, choose the f(time) tab, this mode allows to use a layer as resample curve. The curve is interpreted as follows: 1 is the neutral value, values above are resampled to give higher frequencies as the original, values below are resampled to give lower frequencies as the original. The audiosamples in a given time location are resampled with the values of the parameter curve of the same time location.

In the figure below the resampling has been performed yet, you can see the sound-layer has become a little longer because of the resampling.

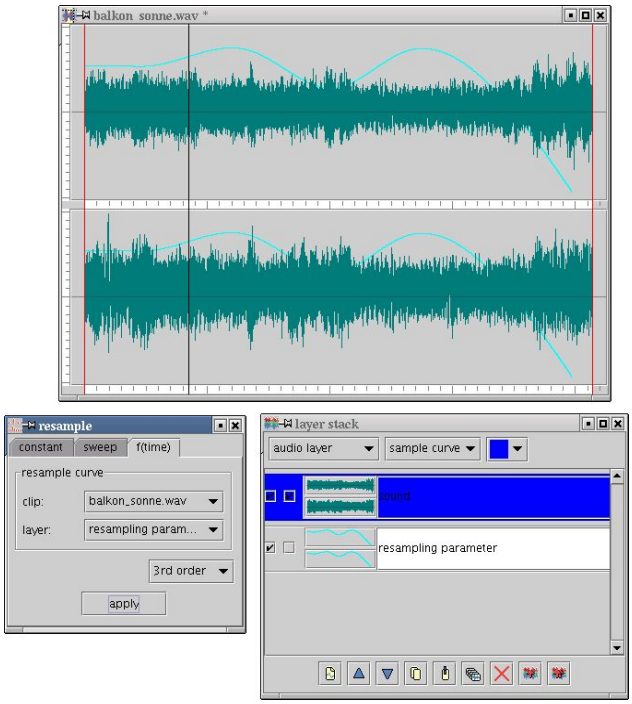

Convert progressively to mono

Another example: a clip containing stereo-sound should progressively convert to mono at the end of the clip. This is possible with the narrow-wide plugin, using a single channel as parameter. So create a new layer with just one channel.

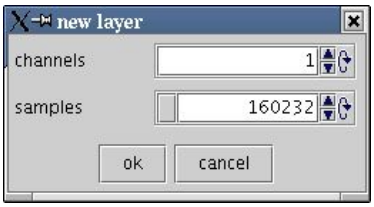

The narrow-wide curve has following value-range: 0 narrows totally to mono, 1 has no effect, 2 wides totally to stereo. Since we have a stereo-clip and want to convert to mono, the parameter can begin with 1 and progressively near to 0. How we proceed: create a constant curve with the value of 1 on the whole channel, then select the end-range, create a constant value of 0 there.

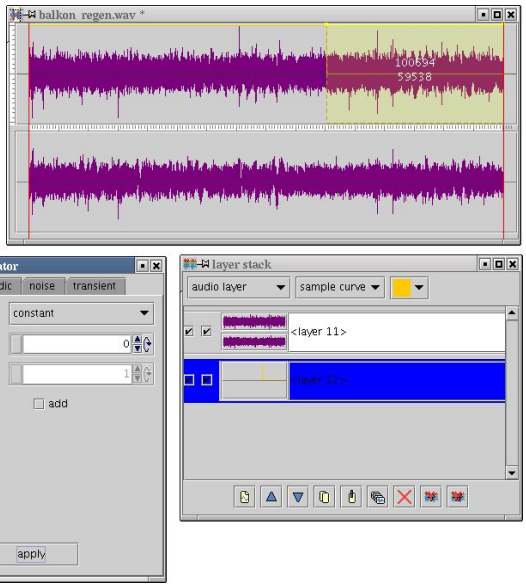

So the end-parts of the curve are prepared. Now let's make the transition part. Create a smooth curve with the free-generator plugin, using the spline interpolation. The end-points have been snapped to the curve using the ctrl-key to minimizes aliases.

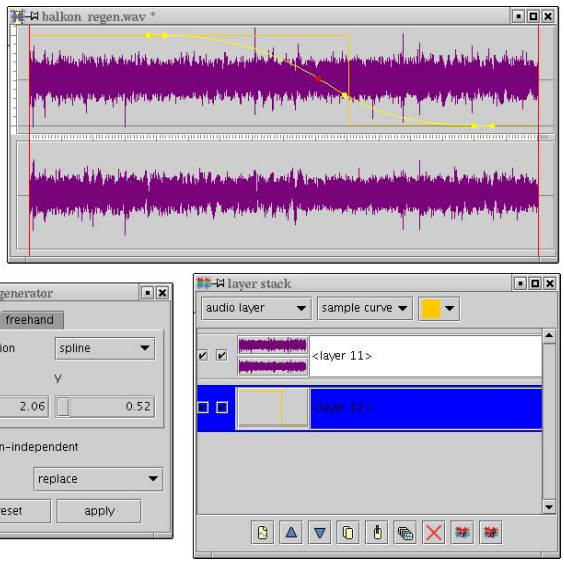

The narrow-wide parameter is now ready, a curve beginning with 1 (no changes) and ending with 0 (totally mono) with a smooth, progressive transition has been created. Select the <layer 11> which contains the original sound, open the narrow-wide plugin, select the f(time) mode, choose the parameter channel and perform the effect. You can individually perform the narrow-wide effect on the left (top) channel and right (bottom) channel. To have full effect, you have to enable both channels.

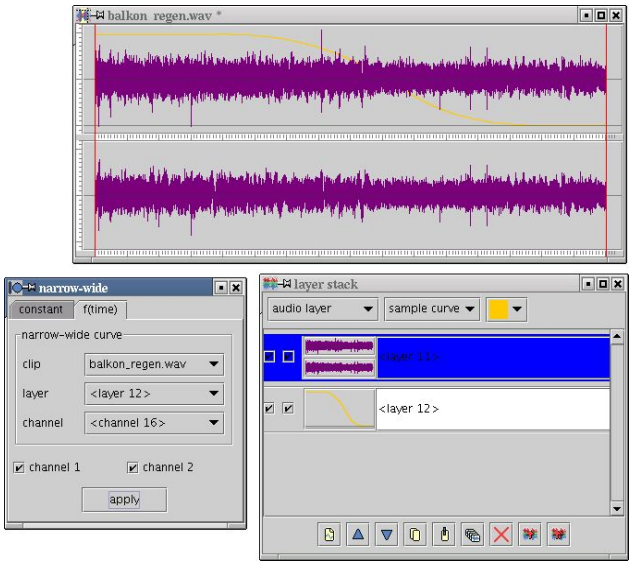

You cannot hear the result in this document, but you can see, the channels are still different at the beginning and they have the same shape at the end now.

Other variable effects

Many plugins allows to use parameter curves, to apply a variable parameter in function of time or in function of something else. With the help of the different generator-plugins it is possible to create almost any possible parameter curve.

10.Making superposing measures

We have seen layers containing sound to be mixed, and layers used as parameters. This chapter shows another application of layers, layers used for measurements.

Finding clicks and pops

Disturbing clicks- and pops-noise from old vinyl-records or from badly manipulated digital sound is not recognizable very easily in classical sample-curve view. Ok, you could see it when zooming very deeply in, but this doesn't help finding them in houndred-thousands or millions of samples of a clip. An easy method to track them is described here. Make a duplification of the original sound-layer, and work with it (select it). The idea is to perform a derivation of the sample-curve, so you can see the slope of the sample-curve. Since clicks often have a bigger slope than the original sound, this method helps detect the clicks.

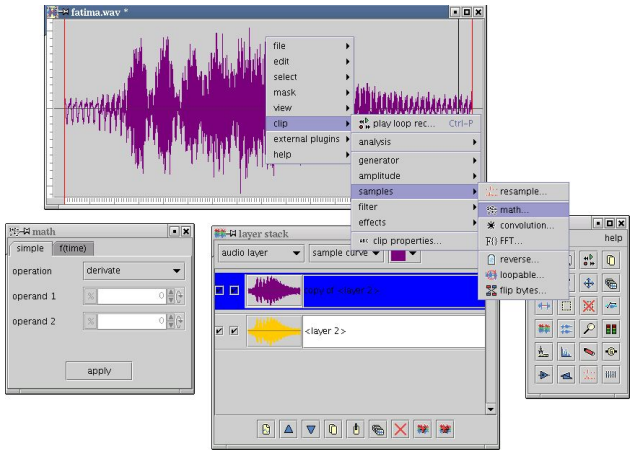

The figure below shows the derivation obtained with the math plugin. The deformed original sound is still recognizable in the result, and additionally you see three peaks, positive and negative ones. The sign of the peaks is not relevant, they sound the same way. So this clip contains probably three disturbing clicks. If no peak is visible clearly, try to perform the derivation a second and third time, the peaks increases on each step. Don't forget to autoscale after each step.

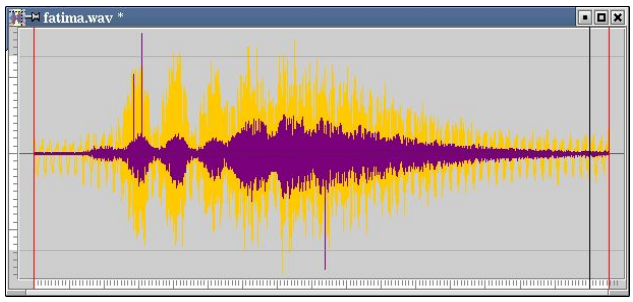

Once the result is satisfactory, we can zoom precisely where the clicks are located, without the need to search for them. The figure below shows the third negative click.

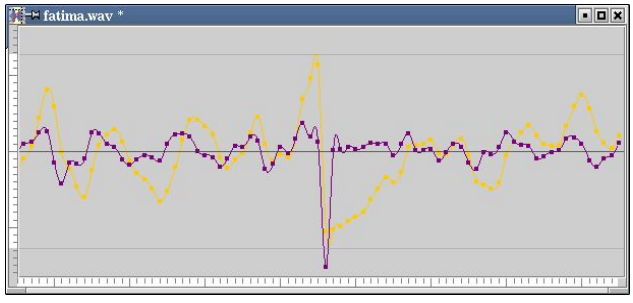

Now we could filter them out using a filter or the freehand pen of the free-generator.

Show a RMS mean curve of the clip

We need to know the RMS average of a sample-curve, in function of time. The math plugin offers a moving RMS average, where the weight of average is configurable in the first operand. The bigger the value, the flatter the result.

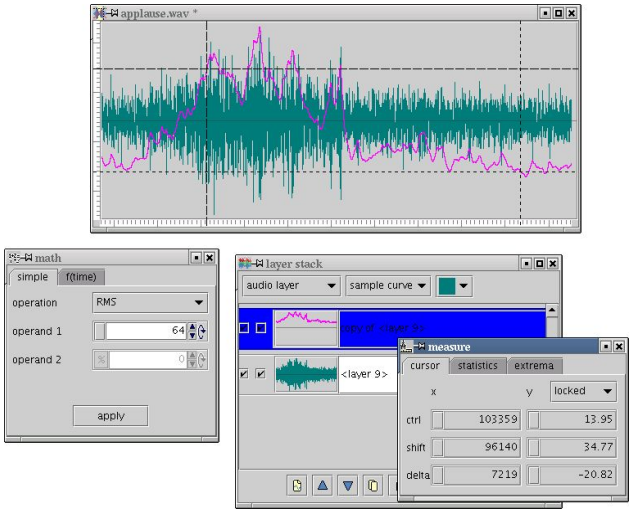

Create a duplification of the original layer, and perform the RMS average on it. After zooming as desired, you have the result superposed to the original sample-curve. The measure-plugin shows you the numerical values of the result.

Measure the transfer-function of a filter

It sounds very strange! LAoE is capable to measure a transfer-function of a filter, this is due to the fact, that LAoE allows to edit spectrums. Given a filter as a black box, observing the input- and the output-spectrum, the relationship between them is called transfer-function. In our example, the clip birds_mono.laoe is the original clip, the input. The clip birds_mono_filtered.laoe has been filtered, so it is the output of the filter. Say the filter is unknown, a black box. It is not exactly true, it is just an exercise, but I'll show you the filter at the end of this chapter. Now let's create the spectrums. Open the spectrum plugin, choose the Blackman window, create a new spectrum clip, select (focus) the original clip and update the spectrum. It's important to focus the original clip before update, because the spectrum plugin works with the focussed clip as input. Now the spectrum clip contains the spectrum of the original clip. If you can't see the spectrum, it's because the spectrum is zoomed anywhere. Autoscale the spectrum, and you will see the complete spectrum. The first spectrum has been created, now do it again for the output-clip birds_mono_filtered.laoe. Create a new spectrum again, select the output-clip, update the spectrum and autoscale it. You have now 4 clips, the input, the output, the input-spectrum and the output-spectrum, as in the figure below.

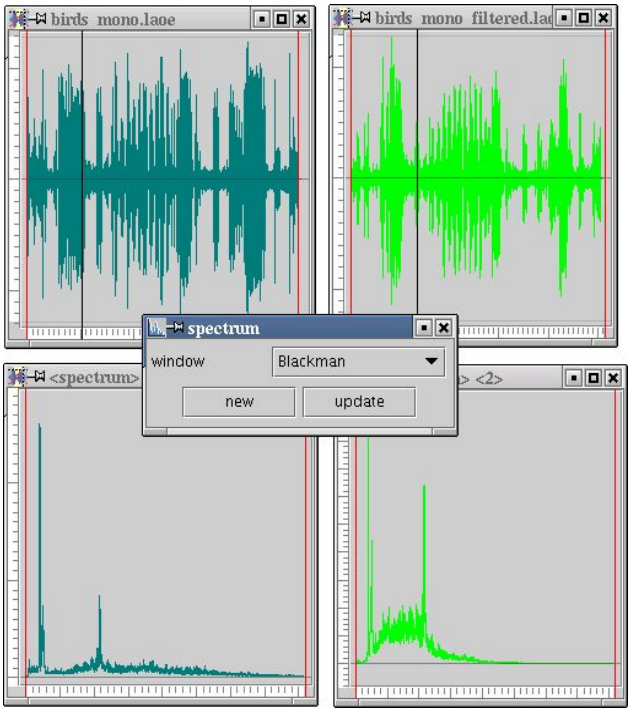

The two spectrums have visible differences. We could visualize them in a unique clip, working with copy and paste layers in the layer-stack plugin. This step is not required here, but the two superposed spectrums give an interesting view. Now we can create the transfer-function. Take the output-spectrum and divide it by the input-spectrum. To do so, select the ouput-spectrum, open the math-plugin and select the divide-operation in the f(time) tab. In opposite to the simple tab, the operations here uses complete curves as operands and not only constants.

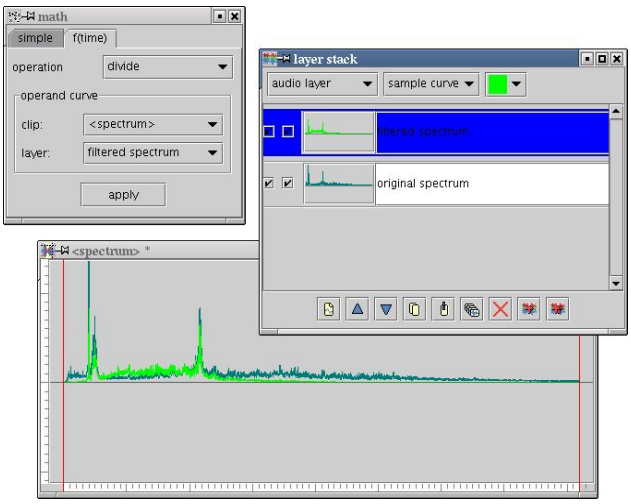

Select the input-spectrum as operand and perform the division. The operand is not changed, the selected layer will be overwritten with the result.

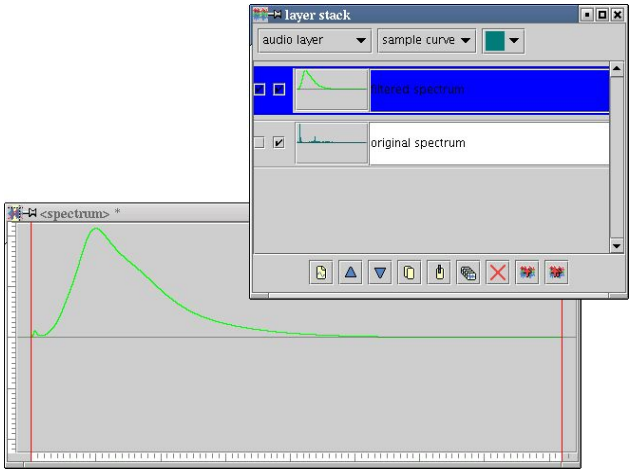

The transfer-function is now visible in linear scale, the other spectrum is not used anymore, so we delete that layer. Since the spectrums are typically shown in dB-scale, we will convert the linear scale in dB-scale using the math-plugin again. The final result is visible in the figure below. We recognize two peaks, one at 83Hz and the other at 1347Hz. The second peak is about 9dB higher than the first. The measure-plugin shows the numerical values with the help of its two cursors. The peak-mode helps to snap to the peaks.

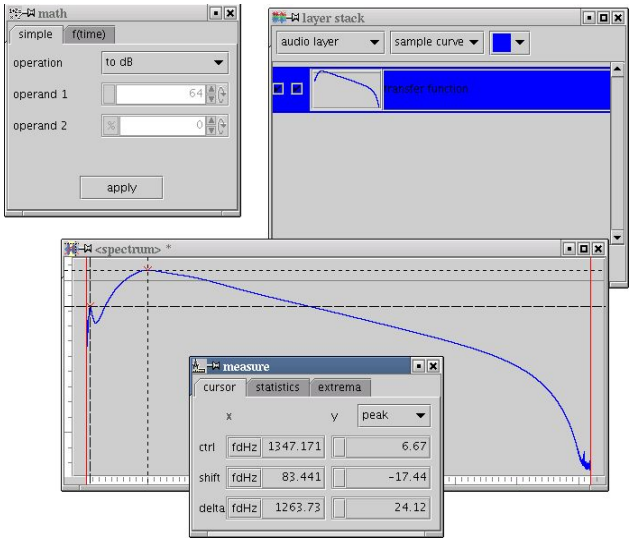

The transfer-function may be edited, zoomed and saved as any other sound-clip. So what about the blackbox filter? Here it is, I used the equalizer-plugin with the settings visible in the figure below. We recognize the setting corresponds to the transfer-function: both have peaks at about 83Hz and 1347Hz (the horizontal scale is linear in the clip and logarithmic in the equalizer), and the second peak is bigger than the first. The transfer-function was calculated correctly.

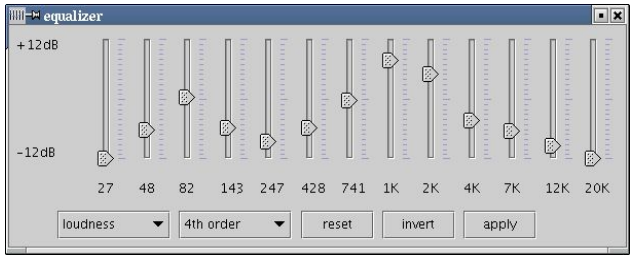

That's all folks!