

SOUNDSHAPER: How to create hybrids of two sounds

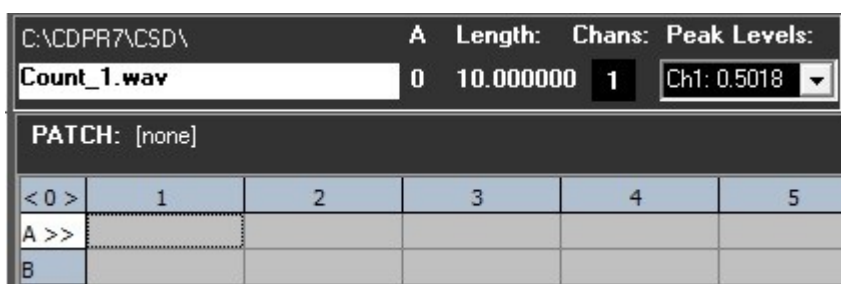
INTRODUCTION

Soundshaper is a control interface for the Composers' Desktop Project suite of sound manipulation tools. CDP processes are all 'off-line': the processing is not done in real-time; instead it typically takes a soundfile (.wav) as an input and produces an output file. For spectral processes, the input should be a spectral analysis file (.ana) and the output is also a spectral file.

Some processes require a secondary input, while others (such as mixing) can take several input sounds. This short tutorial looks at how to select secondary files to create hybrid sounds.

SINGLE-SOURCE PROCESSES

First of all, we need to understand how Soundshaper deals with processes involving just one input soundfile.

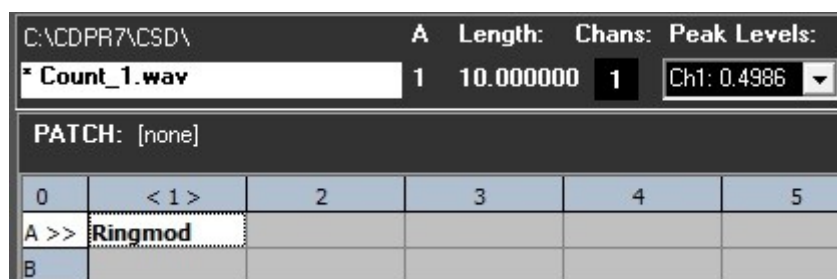


The screenshot shows a portion of Soundshaper's Patch Grid. Row A has a soundfile loaded into it on column 0, which is for sources. The source's name is shown as "Count_1.wav" (someone counting from 1 to 10), length 10", mono and with peak Level of just over 50%. (The file's waveform will also be displayed graphically towards the top of the page and there are "tape-transport" controls to control its playback.)

Each row represents a possible chain of processes from left to right along the row. The rest of Row A and all of Row B are currently empty. The current row is identified by the >> arrows, while the current column number is highlighted by < > (here, < 0 >).

To put the source file in cell A_0, first click on that cell to highlight it (if it's not selected by default). Load the file by drag+drop from an open folder, or by selection from the File Menu, or choose from a pre-populated drop-down list of favourite files shown in the Files box (if you have any).

Using the process menu, we might select PITCH | RING-MODULATE. This leads to a parameter page where a frequency of 100Hz will give the sound a suitably artificial quality. Clicking 'RENDER' ('OK' in older versions), returns us to the Main page. The Patch Grid now looks like this:



The current sound is now our processed one in cell A_1, ready to play. To compare it with the original, click back on cell A_0 and play that instead. Soundshaper always processes the last sound in the currently highlighted row. So to add a further process to the Ringmod sound, select the process and the result will be returned to cell A_2. (Note that we could have chosen a different row for this chain of processes, by highlighting Row B, for example.)

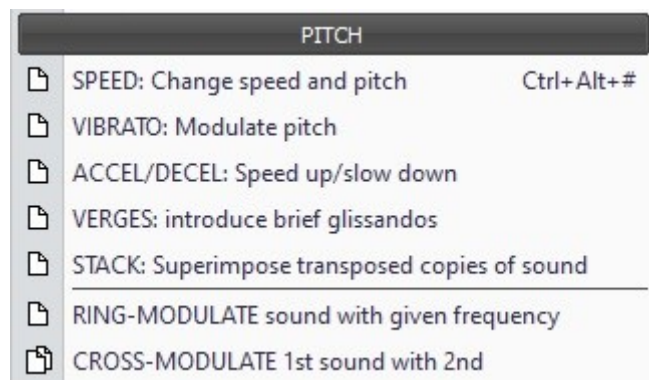
2.

CROSSMOD: A HYBRID USING TWO SOURCES

Now we'll use a slightly different process, **CROSSMOD**, in which our source sound is multiplied by a second sound, instead of by a single frequency as in **RINGMOD**. (Ringmod and Crossmod are both described in the CDP documentation under Modify Radical.)

Firstly, we delete the RINGMOD cell by highlighting it and selecting CTRL+Z for Undo, or menu UNDO | Undo Ringmod, or hit the DELETE KEY. Row A again looks like the first screenshot above.

How would we know that CROSSMOD needs a second source? Soundshaper currently deals with this in two ways. If we look at the menu list, the icon for CROSSMOD shows two files (instead of one as for Ringmod):

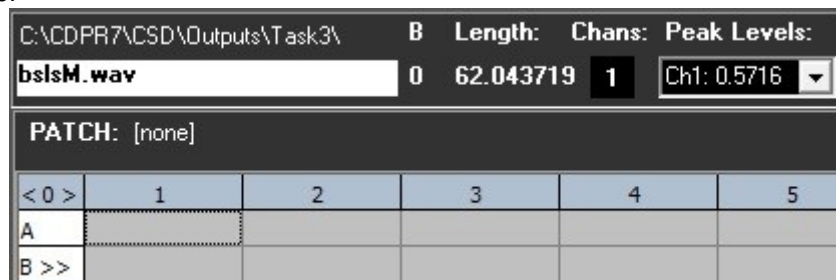


The other way is to show Quick Help: see Menu HELP | Show Quick Help or hit the F1 key. Now, when you run the mouse down the contents of a process menu, a short description is given in a separate window, plus the IN/OUT requirements. For CROSSMOD, this lists two .wav inputs and one .wav output:

```
CROSSMOD Cross-modulate one sound by another
Two input soundfiles are multiplied, creating
complex sidebands
-----
-
IN 1: .wav, IN 2: .wav  OUT: .wav
```

(The Quick-Help window self-closes when you select a process.)

Now we need to select the second source. Any row will do, except Row A which is hosting the primary source. Click on cell B_0 and as before select a sound by drag+drop or from the File menu etc. It should look something like this:



Cell B_0 now contains sound b1s1M.wav, which is the sound of the ocean (as supplied with the forthcoming Keele tutorials); it's much longer than Count_1, but that doesn't matter for CROSSMOD. Highlight Row A again and select the process: Menu PITCH | CROSS-MODULATE.

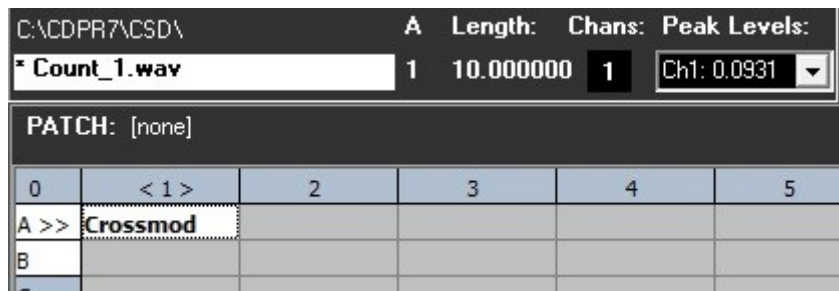
3.

Now a red **ADD INPUT** box appears for you to select the second source. To do this, simply click on the cell B_0 and this will be shown as the selected second source in the box. Click the box's OK button to proceed (or CANCEL if you wish to cancel the selection and process).



If you're not expecting it, this box could easily be missed, despite being bright red!

CROSSMOD has no parameters, so clicking OK processes the sound straight away and the result appears in Cell A_1, as it did with RINGMOD:

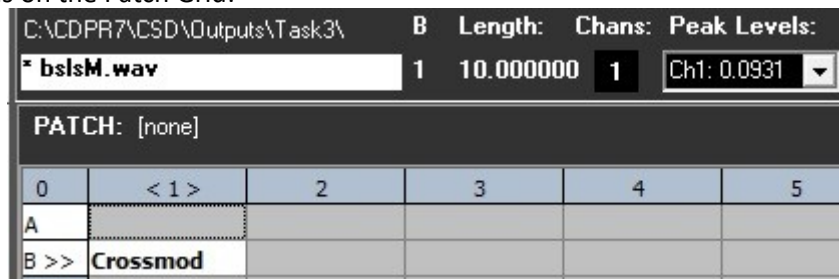


Another way of selecting the second sound – less recommended – could be useful if the sound is not already on the Patchgrid. Starting from the first screenshot, where only the primary source is on the grid in cell A_0, select the process and the ADD INPUT box appears. Clicking the white **wav** label opens a file selector; navigate to the second source's folder and select it. Soundshaper puts the sound into the next available row (in this case B) and selects it.

Note that the sound is shown in the ADD INPUT box by its temporary file name ~B_0.wav, not BelsM.wav. The reason for this is to do with saving and running patches – so that a saved patch remains portable and not tied to specific files in specific folders – but that is another topic altogether.

Swapping the primary and secondary sources

The B_0 cell could equally have been made the primary one and cell A_0 the secondary, so that the result would look like this on the Patch Grid:



In the case of Crossmod, the result would sound identical, but with most processes it would be very different. **The general idea within CDP is that the secondary sound is altering the primary one.** Key examples are imposing an envelope on a sound, or vocoding it.

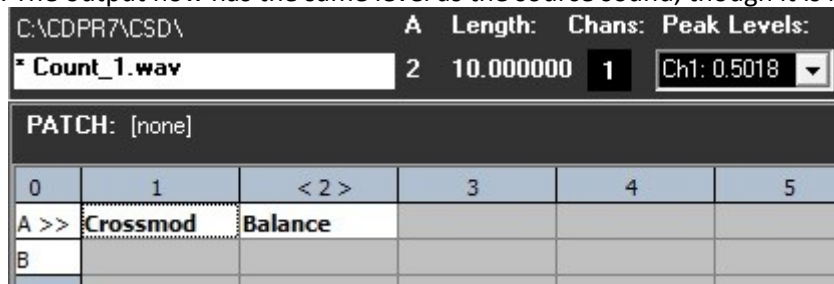
A word of caution: CDP expects the two files to be of compatible type. Processes won't work if one is a 24-bit soundfile and the other is 16-bit, for example. If you're not sure, click on the cell and select CTRL+I (*sndinfo props*) or CTRL+ALT+1 (*sfprops*) for file information.

4.

Selecting a second sound from the same row

Within Soundshaper, the primary source will always be the last cell of the selected row, whereas the secondary cell can be any one, even an earlier one from the same row, as in this next example.

In the CROSSMOD example, the output sound was significantly quieter than the original. We could guess a level to boost it, or use **BALANCE** to bring its peak level up to that of the original. BALANCE needs a second input for comparison. Select Menu LEVEL | BALANCE and this time select the original sound in cell A_0 as the second source. The output now has the same level as the source sound, though it is now in cell A_2:



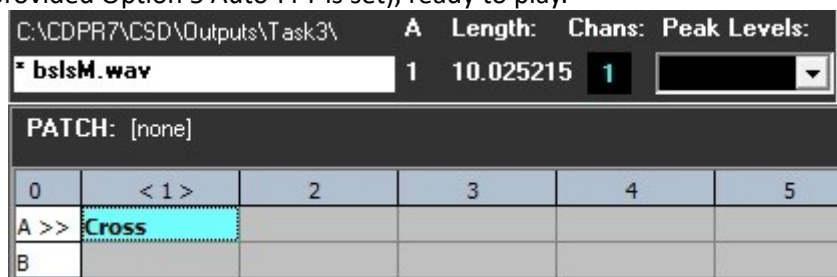
Spectral Hybrids

Several CDP processes use spectral files in the frequency domain to create hybrids. A simple example is vocoding, or the more direct **CROSS**, in which the frequencies in the primary source are replaced by those of the second one. For a “good” result, it’s normal for the modulating second sound to be rhythmic and the primary one to be relatively static. For this reason, the two sources have been re-loaded with “Count” in B_0 and the sea (bslsM.wav) in Row A. (This is not strictly necessary, as we could simply select A_0 as the secondary sound, processing a primary one in B_0.)

CROSS is selected from the menu SPECTRAL | FORMANTS | CROSS. It has only one parameter – the degree of replacement, which for clarity is best set to 1 (100%). As before, we click cell B_0 as the secondary source. Soundshaper auto-converts the .wav file into a spectral one (.ana) and the red box shows this:



The primary input A_0 is also auto-converted behind the scenes. The result is a rather gurgly counting. The output cell is coloured aqua to indicate a spectral output, but that too, is normally converted back to .wav in the background (provided Option 5 Auto-FFT is set), ready to play.



CDP spectral processes (and their derivatives using formant files) are natively mono-only, but Soundshaper can handle stereo and multi-channel files by channelising them, processing each channel separately and then merging the outputs. A selected stereo secondary file may show a channel suffix like _c1:



5.

This is normally of no interest to the user, as long as the process doesn't fail because of some CDP restriction or other. Given the size of the CDP suite, that is unfortunately a possibility. The workaround would be to convert the file to mono before processing.

OTHER EXAMPLES

Envelope Following

The second sound's envelope can be imposed on, or replace, that of the source. This works well with drums imposed on a more static sound and appears to change the drum sound. Some trimming of the length and level may be needed.

Pitch Processing

CDP can extract pitch, process it and then combine the result with the formants (spectral envelope) extracted from the same or another sound. Where the same sound is used, Soundshaper has a built-in composite process (see SPECTRAL | PITCH/FREQUENCY | QUANTIZE, RANDOMIZE etc.) which does pitch and formant extraction and the final combining automatically, with no secondary selection needed.

However, the following screenshot shows the procedure the long-way round:

0	1	2	3	< 4 >
A >>	GetPitch	PchQuant	MakeSpec	Balance
B	GetFmts			

The source vocal sound in A_0 is copied to B_0 (select A_0, click Alt+C to copy, click cell B_0 to paste). Extract the pitch from A_0 (Menu PITCH DATA | EXTRACT PITCH | GET PITCH). Also extract the formants from B_0 (Menu SPECTRAL | FORMANTS | FET FORMANTS). Process cell A_1 (extracted pitch) to quantize it (Menu PITCH DATA | PROCESS Extracted Pitch | QUANTIZE). The process uses a pitch template of MIDI values; the Soundshaper default is a series of semitone steps. The quantized output in A_2 is then combined with the original formants using MAKE (Menu PITCH DATA | MAKE SPECTRUM, which also available under FORMANTS). The secondary input selected for this process is cell B_1, the formant file. The output in A_3 is a spectral file. BALANCE (A_4) has been added to bring the level up to that of the source A_0 (which is selected as the second file).

You wouldn't need to do it this way unless the formants were taken from a different sound – which could even be a hybrid derived from the same source.

* POSSIBLE BUG *

At the time of writing, certain apparently standard .wav files are currently not working with **repitch getpitch** (CDP's pitch extractor function), while others are perfectly OK. This affects the March '23 compilation of **repitch.exe**. An earlier compilation of 2 Jan. '23 seems to work fine, so as a temporary fix, I'm making that available at: <https://www.soundshaper.net/downloads/repitch.zip>

CDP is aware of the fault and when fixed, this paragraph will be removed.
