



CDP SPEC Functions - Basic Spectral Utilities

(with Command Line Usage)

List of basic SPEC (Gain/Edit) utilities

ANALJOIN

Join analysis files together

BARE

Zero the data in channels which do not contain harmonics

CLEAN

Remove noise from PVOC analysis file

CUT

Cut a section out of an analysis file, between *starttime* and *endtime* (seconds)

GAIN

Amplify or attenuate the spectrum

GATE

Zero all channels (in all windows) whose amplitude lies below the threshold

GRAB

Grab a single analysis window at the point specified

MAGNIFY

Expand (in duration) a single analysis window at time *time* to duration *dur*

ANALJOIN JOIN – Join analysis files together

Usage

analjoin join *inanalfile1 inanalfile2 [inanalfile3 ...] outanalfile*

Example command line to join analysis files:

```
analjoin join file1.ana file2.ana file1join2.ana
```

Parameters

inanalfile1 – first input analysis file

inanalfile2 – second input analysis file to join with first

[*inanalfile3 ...*] – third (or more) input analysis files to join

outanalfile – output analysis file with all inputs joined together: spliced end-to-end in the order of input

Understanding the ANALJOIN JOIN Function

ANALJOIN is a straightforward join utility. Note that there is no splicing window function, which does not apply in the spectral analysis data domain. For technical reasons, it may not always be possible to achieve a smooth join.

Musical applications

This enables us to join files in the spectral domain without having to return to the time domain to do so.

The process can be particularly useful when wanting to create analysis data with a great deal of variation in it, for example to create slow transitions by time-stretching such a file.

End of ANALJOIN JOIN

SPEC BARE – Zero the data in channels which do not contain harmonics

Usage

spec bare *infile* *pitchfile* *outfile* [-x]

Parameters

infile – input analysis files made with PVOC

outfile – output analysis file

pitchfile – a binary pitch data file

The *pitchfile* must be extracted from your *infile* analysis file before running SPEC BARE so it is ready to be used as an input.

This is done with **REPITCH GETPITCH**, normally using the **-z** flag to mark any unpitched material

-x more body in resulting spectrum

Understanding the SPEC BARE Process

SPEC BARE zeroes all the non-partial data, retaining the harmonics of the original tone. The harmonics are the integer multiples of the fundamental. When the original is a complex sound, the actual result is dependent on what the GETPITCH process works out this fundamental to be.

If the original sound is a clearly pitched tone, it is probable that harmonic partials will be captured in some channels. SPEC BARE will zero out the data in the channels which do not contain harmonic partials, to provide a good, clean version of the original for further processing.

If the original sound is a complex tone with many inharmonic partials, then the SPEC BARE process may not work. The windows will be flagged as having no pitch. However, I (Ed.) have used it on a donkey bray and found that the result somehow brought out an overtone series. This was then further emphasised by applying HILITE PLUCK to the SPEC BARE output.

Musical Applications

As implied above, the musical uses of SPEC BARE begin when another process needs to work with clearly identified partials content, such as emphasising pitch content, overtone series, or harmonic colouration.

Some of the processes which may benefit from the use of the BARE Process are (use BACK in your Browser to return here to SPEC BARE):

BLUR BLUR
MORPH MORPH
HILITE PLUCK
HILITE TRACE
PITCH TUNE
HILITE FILTER
FORMANTS VOCODE

Also see: **REPITCH GETPITCH** and **PITCH ALTHARMS**

End of SPEC BARE

SPEC CLEAN – Remove noise from PVOC analysis file

This process may have been superseded by the more recent function [SPECNU CLEAN](#).

Usage

spec clean 1-2 *infile nfile outfile skiptime* [-**gnoisgain**]
spec clean 3 *infile nfile outfile freq* [-**gnoisgain**]
spec clean 4 *infile nfile gfile outfile* [-**gnoisgain**]

Modes

- 1** deletes a channel (after *skiptime*) FROM THE TIME its level falls below the (*noisgain* adjusted) maximum level seen for that channel in *nfile* TO THE END of the file
- 2** deletes channel (after *skiptime*) ANYWHERE its level falls below the (*noisgain* adjusted) maximum level seen for that channel in *nfile*, RESTORING data when it rises above that level
- 3** deletes channel as in MODE 2 but ONLY for channels of frequency > *freq*
- 4** deletes channel EVERYWHERE, whose level in *gfile* is ALWAYS below the (*noisgain* adjusted) maximum level seen for that channel in *nfile*

Parameters

infile, *nfile* and *gfile* are all analysis files made with PVOC

nfile and *gfile* should be cut (using [SPEC CUT](#)) from *infile* to show typical noise (*nfile*) and good signal (*gfile*)

outfile – output analysis file

skiptime – (seconds) may be set to time at which good source signal level has been established

freq – frequency in Hz above which noise is to be removed

noisgain – multiplies noise levels found in *nfile* before they are used for comparison with *infile* signal: (Default 2)

Understanding the SPEC CLEAN Process

There are now several ways to approach the removal of unwanted noise, including [SPECNU CLEAN](#) and other functions within SPECNU.

SPEC CLEAN works by comparing the level (for each channel in turn) in *nfile* (the noise file) with the level in the sound to be cleaned. To find the noise threshold in any channel, i.e., the level below which we decide the data is merely noise and not significant signal, the program looks at the whole *nfile* and takes the maximum level it finds (in each channel) as the noise threshold (for each particular channel).

If *noisgain* is then set as e.g., 2, then the noise threshold (for each channel in turn) is set 2-times louder than the threshold extracted from the noise file.

In Mode **1**, as soon as the level in (the particular channel of) the file to be cleaned falls below the threshold (for that channel), the channel data is eliminated from that time to the end of the file. This is therefore useful for cleaning up the end of a file where the sound source is perhaps fading away, and falls below the noise threshold.

In Mode **2**, the channel data is eliminated ANYWHERE it falls below the noise threshold (for that channel), but the data is restored at any time it rises above that threshold again.

Mode **3** works like Mode **2**, but only on a specified high frequency area.

Mode **4** attempts to look for channels which contribute 'nothing' to the spectrum of the sound required. It does this by comparing a noise segment with a good signal segment, and deciding which channels are 'redundant', i.e., are always below the noise threshold.

The key to its use therefore lies in selecting effective sections of the soundfile for these reference levels. Listen carefully to the source soundfile, noting where appropriate sections begin and end, and use these times to cut the input (analysis) file with SPEC CUT to produce the required comparison files.

This may take some trial and error. Listen for 'bubbling' effects; these may occur if the noise level or the noise comparison sample isn't quite right to achieve the desired noise reduction.

Musical Applications

This process can be useful in tidying up the signal of e.g., an analogue recording. However, the user is cautioned that the methods used here are relatively simple and are not sufficiently 'intelligent' to professionally digitally remaster analogue recordings.

Also see [SPECNU CLEAN](#) for a more recent version of this process.

End of SPEC CLEAN

SPEC CUT – Cut a section out of an analysis file, between *starttime* and *endtime* (seconds)

Usage

spec cut *infile outfile starttime endtime*

Parameters

infile – input analysis file made with PVOC

outfile – output analysis file

starttime – time in seconds at which the cut is to begin

endtime – time in seconds at which the cut is to end

Understanding the SPEC CUT Process

SPEC CUT makes it possible to extract a section of a PVOC analysis file and save it as a new analysis file. The cut points are specified in seconds.

Musical Applications

The primary application is simply to shorten an analysis file without having to go back to the original soundfile, cut that and re-analyze. Listening to the original soundfile should give a rough indication of where the cut points might best be located.

It is also used to select portions of an analysis file containing 'typical noise' (the *nfile*) and 'good signal' (the *gfile*) for use with SPEC CLEAN. To do this, one uses a standard sound editor, blocking out sections of the original source sound, listening to them, and making a note of the time points. These time points are then used with SPEC CUT directly on the analysis file for that soundfile – the times of source and analysis match.

End of SPEC CUT

SPEC GAIN – Amplify or attenuate the spectrum

Usage

spec gain *infile outfile gain*

Parameters

infile – input analysis file made with PVOC

outfile – output analysis file

gain – numerical float value > 0.0

gain may vary over time

Understanding the SPEC GAIN Process

GAIN multiplies the amplitude of each partial with the *gain* factor. This will **attenuate** the amplitude if less than 1 and **increase** the amplitude if greater than 1. For any soundfile, there is an optimal gain factor, above which distortion is created.

Musical Applications

GAIN can be usefully applied to bring a sound up to full volume before running another process, or restoring it to full volume if it has lost amplitude as a result of the process. However, it does need to be used with some caution. Large or repeated amplifications of a poor signal will also amplify discontinuities and increase digital noise.

Sometimes PVOC, the *Phase Vocoder*, recommends an increase or decrease in gain, esp. the latter, if an overflow has occurred. SPEC GAIN is used to make this adjustment.

Excessive application of gain on a poor signal can introduce digital noise by creating sharp changes in amplitude levels.

End of SPEC GAIN

SPEC GATE – Eliminate channel data below a threshold amplitude

Usage

spec gate *infile outfile threshold*

Parameters

infile – input analysis file made with PVOC

outfile – output analysis file

threshold – lowest acceptable amplitude level (Range: 0 to 1)

threshold may vary over time

Understanding the SPEC GATE Process

SPEC GATE very simply cuts out low level signal.

Musical Applications

It is intended as a facility by which various unwanted artefacts in the sound can be removed from the analysis file. The use of the phase vocoder with various types of input sound can be somewhat unpredictable, with the result that 'artefacts' may be produced, often heard as a light swishing sound.

End of SPEC GATE

SPEC GRAB – Grab a single analysis window at time point specified

Usage

spec grab *infile outfile time*

Parameters

infile – input analysis file made with PVOC

outfile – output analysis file

time – time location in seconds of the window to grab

A time beyond the end of the file will grab the last window in the file.

Understanding the SPEC GRAB Process

SPEC GRAB grabs a single window from an existing analysis data file. The idea behind this is to be able to take the information from a single window for various purposes.

Musical Applications

Two applications have been implemented so far: [SPEC PRINT](#), which prints the analysis data to the screen or to text file; and [MORPH GLIDE](#), which interpolates between two single analysis windows extracted by SPEC GRAB. [SPEC MAGNIFY](#), which time-expands the data of one window to any duration, is another application of this idea, but it selects the window for itself.

End of SPEC GRAB

SPEC MAGNIFY – Expand (in duration) a single analysis window at time *time* to duration *dur*

Usage

`spec magnify infile outfile time dur`

Parameters

infile – input analysis file made with PVOC

outfile – output analysis file

time – time location in seconds where the single window to expand occurs in the analysis file

dur – the required duration of *outfile*; it MUST BE > the analysis window length

Understanding the SPEC MAGNIFY Process

This function was previously named SPECWEXP. It synthesizes a steady sound from the analysis data at any single window within the analysis file. Listen for material in the source which might benefit from being drawn out for a longer period of time. A wide range of frequencies and rapidly changing timbral content are two things to listen for.

Note that the resynthesized sound output will need to be topped and tailed. The CDP functions [HOUSEKEEP EXTRACT 3](#) or [TOPNTAIL2](#) can be used for this, but many sound editors also have this facility.

Musical Applications

The purpose of SPEC MAGNIFY is to generate a potentially interesting sound by expanding a moment of apparent timbral interest.

End of SPEC MAGNIFY