

FFmpeg Resampler Documentation

Table of Contents

- 1 Description
- 2 Resampler Options
- 3 See Also
- 4 Authors

1 Description# TOC

The FFmpeg resampler provides a high-level interface to the libswresample library audio resampling utilities. In particular it allows one to perform audio resampling, audio channel layout rematrixing, and convert audio format and packing layout.

2 Resampler Options# TOC

The audio resampler supports the following named options.

Options may be set by specifying *-option value* in the FFmpeg tools, *option=value* for the aresample filter, by setting the value explicitly in the `SwrContext` options or using the `libavutil/opt.h` API for programmatic use.

`ich, in_channel_count`

Set the number of input channels. Default value is 0. Setting this value is not mandatory if the corresponding channel layout `in_channel_layout` is set.

`och, out_channel_count`

Set the number of output channels. Default value is 0. Setting this value is not mandatory if the corresponding channel layout `out_channel_layout` is set.

`uch, used_channel_count`

Set the number of used input channels. Default value is 0. This option is only used for special remapping.

`isr, in_sample_rate`

Set the input sample rate. Default value is 0.

`osr, out_sample_rate`

Set the output sample rate. Default value is 0.

`isf, in_sample_fmt`

Specify the input sample format. It is set by default to none.

`osf, out_sample_fmt`

Specify the output sample format. It is set by default to none.

`tsf, internal_sample_fmt`

Set the internal sample format. Default value is none. This will automatically be chosen when it is not explicitly set.

`icl, in_channel_layout`

`ocl, out_channel_layout`

Set the input/output channel layout.

See (ffmpeg-utils)the Channel Layout section in the ffmpeg-utils(1) manual for the required syntax.

`clev, center_mix_level`

Set the center mix level. It is a value expressed in deciBel, and must be in the interval [-32,32].

`slev, surround_mix_level`

Set the surround mix level. It is a value expressed in deciBel, and must be in the interval [-32,32].

`lfe_mix_level`

Set LFE mix into non LFE level. It is used when there is a LFE input but no LFE output. It is a value expressed in deciBel, and must be in the interval [-32,32].

`rmvol, rematrix_volume`

Set rematrix volume. Default value is 1.0.

`rematrix_maxval`

Set maximum output value for rematrixing. This can be used to prevent clipping vs. preventing volume reduction A value of 1.0 prevents clipping.

`flags, swr_flags`

Set flags used by the converter. Default value is 0.

It supports the following individual flags:

`res`

force resampling, this flag forces resampling to be used even when the input and output sample rates match.

`dither_scale`

Set the dither scale. Default value is 1.

`dither_method`

Set dither method. Default value is 0.

Supported values:

`'rectangular'`

select rectangular dither

`'triangular'`

select triangular dither

`'triangular_hp'`

select triangular dither with high pass

`'lipshitz'`

select lipshitz noise shaping dither

`'shibata'`

select shibata noise shaping dither

`'low_shibata'`

select low shibata noise shaping dither

`'high_shibata'`

select high shibata noise shaping dither

`'f_weighted'`

select f-weighted noise shaping dither

`'modified_e_weighted'`

select modified-e-weighted noise shaping dither

`'improved_e_weighted'`

select improved-e-weighted noise shaping dither

`resampler`

Set resampling engine. Default value is `swr`.

Supported values:

`'swr'`

select the native SW Resampler; filter options `precision` and `cheby` are not applicable in this case.

`'soxr'`

select the SoX Resampler (where available); `compensation`, and filter options `filter_size`, `phase_shift`, `filter_type` & `kaiser_beta`, are not applicable in this case.

`filter_size`

For `swr` only, set resampling filter size, default value is 32.

`phase_shift`

For `swr` only, set resampling phase shift, default value is 10, and must be in the interval [0,30].

`linear_interp`

Use Linear Interpolation if set to 1, default value is 0.

`cutoff`

Set cutoff frequency (`swr`: 6dB point; `soxr`: 0dB point) ratio; must be a float value between 0 and 1. Default value is 0.97 with `swr`, and 0.91 with `soxr` (which, with a sample-rate of 44100, preserves the entire audio band to 20kHz).

`precision`

For `soxr` only, the precision in bits to which the resampled signal will be calculated. The default value of 20 (which, with suitable dithering, is appropriate for a destination bit-depth of 16) gives SoX's 'High Quality'; a value of 28 gives SoX's 'Very High Quality'.

cheby

For soxr only, selects passband rolloff none (Chebyshev) & higher-precision approximation for 'irrational' ratios. Default value is 0.

async

For swr only, simple 1 parameter audio sync to timestamps using stretching, squeezing, filling and trimming. Setting this to 1 will enable filling and trimming, larger values represent the maximum amount in samples that the data may be stretched or squeezed for each second. Default value is 0, thus no compensation is applied to make the samples match the audio timestamps.

first_pts

For swr only, assume the first pts should be this value. The time unit is 1 / sample rate. This allows for padding/trimming at the start of stream. By default, no assumption is made about the first frame's expected pts, so no padding or trimming is done. For example, this could be set to 0 to pad the beginning with silence if an audio stream starts after the video stream or to trim any samples with a negative pts due to encoder delay.

min_comp

For swr only, set the minimum difference between timestamps and audio data (in seconds) to trigger stretching/squeezing/filling or trimming of the data to make it match the timestamps. The default is that stretching/squeezing/filling and trimming is disabled (`min_comp = FLT_MAX`).

min_hard_comp

For swr only, set the minimum difference between timestamps and audio data (in seconds) to trigger adding/dropping samples to make it match the timestamps. This option effectively is a threshold to select between hard (trim/fill) and soft (squeeze/stretch) compensation. Note that all compensation is by default disabled through `min_comp`. The default is 0.1.

comp_duration

For swr only, set duration (in seconds) over which data is stretched/squeezed to make it match the timestamps. Must be a non-negative double float value, default value is 1.0.

max_soft_comp

For swr only, set maximum factor by which data is stretched/squeezed to make it match the timestamps. Must be a non-negative double float value, default value is 0.

matrix_encoding

Select matrixed stereo encoding.

It accepts the following values:

`'none'`

select none

`'dolby'`

select Dolby

`'dplii'`

select Dolby Pro Logic II

Default value is none.

`filter_type`

For swr only, select resampling filter type. This only affects resampling operations.

It accepts the following values:

`'cubic'`

select cubic

`'blackman_nuttall'`

select Blackman Nuttall Windowed Sinc

`'kaiser'`

select Kaiser Windowed Sinc

`kaiser_beta`

For swr only, set Kaiser Window Beta value. Must be an integer in the interval [2,16], default value is 9.

`output_sample_bits`

For swr only, set number of used output sample bits for dithering. Must be an integer in the interval [0,64], default value is 0, which means it's not used.

3 See Also# TOC

ffmpeg, ffplay, ffprobe, ffserver, libswresample

4 Authors# TOC

The FFmpeg developers.

For details about the authorship, see the Git history of the project ([git://source.ffmpeg.org/ffmpeg](https://source.ffmpeg.org/ffmpeg)), e.g. by typing the command `git log` in the FFmpeg source directory, or browsing the online repository at <http://source.ffmpeg.org>.

Maintainers for the specific components are listed in the file `MAINTAINERS` in the source code tree.

This document was generated using *makeinfo*.

FFmpeg Resampler Documentation

Table of Contents

- 1 Description
- 2 Resampler Options
- 3 See Also
- 4 Authors

1 Description# TOC

The FFmpeg resampler provides a high-level interface to the libswresample library audio resampling utilities. In particular it allows one to perform audio resampling, audio channel layout rematrixing, and convert audio format and packing layout.

2 Resampler Options# TOC

The audio resampler supports the following named options.

Options may be set by specifying *-option value* in the FFmpeg tools, *option=value* for the aresample filter, by setting the value explicitly in the `SwrContext` options or using the `libavutil/opt.h` API for programmatic use.

`ich, in_channel_count`

Set the number of input channels. Default value is 0. Setting this value is not mandatory if the corresponding channel layout `in_channel_layout` is set.

`och, out_channel_count`

Set the number of output channels. Default value is 0. Setting this value is not mandatory if the corresponding channel layout `out_channel_layout` is set.

`uch, used_channel_count`

Set the number of used input channels. Default value is 0. This option is only used for special remapping.

`isr, in_sample_rate`

Set the input sample rate. Default value is 0.

`osr, out_sample_rate`

Set the output sample rate. Default value is 0.

`isf, in_sample_fmt`

Specify the input sample format. It is set by default to none.

`osf, out_sample_fmt`

Specify the output sample format. It is set by default to none.

`tsf, internal_sample_fmt`

Set the internal sample format. Default value is none. This will automatically be chosen when it is not explicitly set.

`icl, in_channel_layout`

`ocl, out_channel_layout`

Set the input/output channel layout.

See (ffmpeg-utils)the Channel Layout section in the ffmpeg-utils(1) manual for the required syntax.

`clev, center_mix_level`

Set the center mix level. It is a value expressed in deciBel, and must be in the interval [-32,32].

`slev, surround_mix_level`

Set the surround mix level. It is a value expressed in deciBel, and must be in the interval [-32,32].

`lfe_mix_level`

Set LFE mix into non LFE level. It is used when there is a LFE input but no LFE output. It is a value expressed in deciBel, and must be in the interval [-32,32].

`rmvol, rematrix_volume`

Set rematrix volume. Default value is 1.0.

`rematrix_maxval`

Set maximum output value for rematrixing. This can be used to prevent clipping vs. preventing volume reduction A value of 1.0 prevents clipping.

`flags, swr_flags`

Set flags used by the converter. Default value is 0.

It supports the following individual flags:

`res`

force resampling, this flag forces resampling to be used even when the input and output sample rates match.

`dither_scale`

Set the dither scale. Default value is 1.

`dither_method`

Set dither method. Default value is 0.

Supported values:

`'rectangular'`

select rectangular dither

`'triangular'`

select triangular dither

`'triangular_hp'`

select triangular dither with high pass

`'lipshitz'`

select lipshitz noise shaping dither

`'shibata'`

select shibata noise shaping dither

`'low_shibata'`

select low shibata noise shaping dither

`'high_shibata'`

select high shibata noise shaping dither

`'f_weighted'`

select f-weighted noise shaping dither

`'modified_e_weighted'`

select modified-e-weighted noise shaping dither

`'improved_e_weighted'`

select improved-e-weighted noise shaping dither

`resampler`

Set resampling engine. Default value is `swr`.

Supported values:

`'swr'`

select the native SW Resampler; filter options `precision` and `cheby` are not applicable in this case.

`'soxr'`

select the SoX Resampler (where available); `compensation`, and filter options `filter_size`, `phase_shift`, `filter_type` & `kaiser_beta`, are not applicable in this case.

`filter_size`

For `swr` only, set resampling filter size, default value is 32.

`phase_shift`

For `swr` only, set resampling phase shift, default value is 10, and must be in the interval [0,30].

`linear_interp`

Use Linear Interpolation if set to 1, default value is 0.

`cutoff`

Set cutoff frequency (`swr`: 6dB point; `soxr`: 0dB point) ratio; must be a float value between 0 and 1. Default value is 0.97 with `swr`, and 0.91 with `soxr` (which, with a sample-rate of 44100, preserves the entire audio band to 20kHz).

`precision`

For `soxr` only, the precision in bits to which the resampled signal will be calculated. The default value of 20 (which, with suitable dithering, is appropriate for a destination bit-depth of 16) gives SoX's 'High Quality'; a value of 28 gives SoX's 'Very High Quality'.

cheby

For soxr only, selects passband rolloff none (Chebyshev) & higher-precision approximation for 'irrational' ratios. Default value is 0.

async

For swr only, simple 1 parameter audio sync to timestamps using stretching, squeezing, filling and trimming. Setting this to 1 will enable filling and trimming, larger values represent the maximum amount in samples that the data may be stretched or squeezed for each second. Default value is 0, thus no compensation is applied to make the samples match the audio timestamps.

first_pts

For swr only, assume the first pts should be this value. The time unit is 1 / sample rate. This allows for padding/trimming at the start of stream. By default, no assumption is made about the first frame's expected pts, so no padding or trimming is done. For example, this could be set to 0 to pad the beginning with silence if an audio stream starts after the video stream or to trim any samples with a negative pts due to encoder delay.

min_comp

For swr only, set the minimum difference between timestamps and audio data (in seconds) to trigger stretching/squeezing/filling or trimming of the data to make it match the timestamps. The default is that stretching/squeezing/filling and trimming is disabled (`min_comp = FLT_MAX`).

min_hard_comp

For swr only, set the minimum difference between timestamps and audio data (in seconds) to trigger adding/dropping samples to make it match the timestamps. This option effectively is a threshold to select between hard (trim/fill) and soft (squeeze/stretch) compensation. Note that all compensation is by default disabled through `min_comp`. The default is 0.1.

comp_duration

For swr only, set duration (in seconds) over which data is stretched/squeezed to make it match the timestamps. Must be a non-negative double float value, default value is 1.0.

max_soft_comp

For swr only, set maximum factor by which data is stretched/squeezed to make it match the timestamps. Must be a non-negative double float value, default value is 0.

matrix_encoding

Select matrixed stereo encoding.

It accepts the following values:

`'none'`

select none

`'dolby'`

select Dolby

`'dplii'`

select Dolby Pro Logic II

Default value is none.

`filter_type`

For swr only, select resampling filter type. This only affects resampling operations.

It accepts the following values:

`'cubic'`

select cubic

`'blackman_nuttall'`

select Blackman Nuttall Windowed Sinc

`'kaiser'`

select Kaiser Windowed Sinc

`kaiser_beta`

For swr only, set Kaiser Window Beta value. Must be an integer in the interval [2,16], default value is 9.

`output_sample_bits`

For swr only, set number of used output sample bits for dithering. Must be an integer in the interval [0,64], default value is 0, which means it's not used.

3 See Also# TOC

ffmpeg, ffplay, ffprobe, ffserver, libswresample

4 Authors# TOC

The FFmpeg developers.

For details about the authorship, see the Git history of the project ([git://source.ffmpeg.org/ffmpeg](https://source.ffmpeg.org/ffmpeg)), e.g. by typing the command `git log` in the FFmpeg source directory, or browsing the online repository at <http://source.ffmpeg.org>.

Maintainers for the specific components are listed in the file `MAINTAINERS` in the source code tree.

This document was generated using *makeinfo*.