

NAME

level version 2.71 – process an audio input file in WAV format to normalise the signal level

SYNOPSIS

level [*options ...*] [*input file(s)* or *]

DESCRIPTION

level processes WAV format audio file(s) into WAV format output file(s) with the audio signal level either normalised (fixed gain) or peak-limited (gain is continuously adjusted automatically).

File names may be specified individually, or an asterisk (*) may be used as a wild-card to process all WAV files in the directory. The .WAV file suffix is not required with file names. **level** processes each specified file into an output file of the same name but with a "c_" prepended. If "*" is used, WAV files with names starting with "c_" are skipped for input.

Limiting mode

Unless the `-n` (normalise only) flag is used, *level* operates in peak-limiting (infinite audio-compression) mode. Under these conditions, *level* acts as a mono-band peak-limiting amplifier.

Peak-limiters operate as normal amplifiers for signals below the defined limiting threshold. For signals above the limiting threshold, the instantaneous gain is reduced in exact proportion to the increased input, which prevents the output signal from rising further. This gain reduction is achieved with no audible distortion by maintaining its value at the new level for a period of time.

Limiting slope

With analogue peak-limiting amplifiers, finite-gain feedback loops theoretically restrict the performance so that, above limiting threshold, an increase in the input signal will still cause a finite increase in the output signal. The ratio of *the rise in output* to *the rise in input* (for an input which is above the limiting threshold) is called the **limiting slope**. This slope must be at least 10:1 for an amplifier to be useful as a peak-limiter. Modern designs achieve limiting slopes in excess of 100:1. With a limiting slope of 100:1, an input that increases by 25db will result in an output increase of 0.25db.

level has an infinite limiting slope.

Attack time

Another key parameter for specifying the performance of a peak-limiting amplifier is its **attack time**. This determines how quickly the amplifier can reduce its gain in response to a signal increase. Ideally this time interval would be zero, but analogue peak-limiting amplifiers have feedback loops with a finite (non-zero) response time. Early valve designs could take as long as 5 milliseconds to respond to a signal peak. Current designs measure their response times in microseconds. During this *attack time* period, the output momentarily overshoots the nominal maximum value.

level has a response time which is instantaneous in the sense that the gain reduction will always occur fully along the leading edge of the peak of the half-cycle which is attempting to exceed maximum output. The output remains at exactly the maximum value during this time, ie: there is no overshoot.

Gain recovery

For an input signal which is (on average) always above the limiting threshold, each time a peak-limiting amplifier reduces its gain to maintain correct output, it must subsequently *restore* the gain when the signal falls below the specified maximum output. The *rate and manner* of gain restoration must be carefully controlled or waveform distortion (both harmonic and intermodulation) resulting from the gain restoration process may become audible.

The normal approach taken with peak-limiting amplifiers is to employ *multiple* (2 or more) coupled time-constants in the gain-recovery block.

The main high-speed time constant is used to permit a controlled amount of rapid gain-recovery from short isolated transients. In **level**, the *amount* of this high-speed recovery (in db) is selected via `-r`, the fast/slow ratio (default=1), and its *rate* in db per second via `-f` (default=12).

A *sample-and-hold* circuit operates in parallel with this high-speed time constant to delay recovery for a fixed period of 40 milliseconds. This minimises audible distortion on sustained low frequency sinusoids. The sample-and-hold timer is reset each time full output is reached, so for a steady-state sinusoid which is already at full output, it will defeat the recovery time-constants indefinitely.

A longer time constant operating *in parallel* with these provides a much slower gain recovery for sustained

program peaks or longer term drops in level. The slow time-constant rate (in db/sec) is controlled by the `-s` value (default=1).

The combination of these recovery controls provides arbitrarily fast recovery across the top few db of the signal in combination with a slow and less obtrusive recovery for the tracking of sustained drops in level. This reduces *gain chasing* effects, such as bass or percussive instruments audibly modulating the volume of other signals.

Normalise mode

If normalise mode is used (via `-n`), dynamic compression (peak-limiting) is **disabled**. In normalise mode, the input file must be read twice. The first scan locates the highest sample value (+ve or -ve). The file is then read again with the gain preset to the appropriate value for normalised output, and the result written to the selected output file.

The `-l` (output level) control is still available in normalise mode.

Options

- `-d` Debug. Not recommended at all. Generates vast amounts of data on standard output, and will multiply processing time by several orders of magnitude. For emergencies only – absolutely useless to anyone except the author.
- `-e` Print some *examples* of parameter usage (see also `-h`)
- `-f N` Fast recovery time-constant (db/sec). Following a gain reduction, `-f` defines the speed at which gain will recover until either (a) the slow time-constant takes over, or (b) maximum output is reached – whichever occurs first. Default = 12db per second. (See also `-r` and `-t`)
- `-g N` Initial gain value (+db). This is the equivalent of the input attenuator in a real (hardware) peak-limiting amplifier. Default = +6db. If you set this to more than the `-m` value, it is silently reduced to the `-m` value.
- `-h` Print some basic help re parameter usage
- `-l N` Line level output value (-db). Acts as an output attenuator, with a default setting of 0db, corresponding to full output. If a value is specified, it will reduce the output level by that amount. For example, `"-l6"` will set the output level to -6db, ie: to 50%.
- `-m N` Maximum gain (+db). Specifies the gain ceiling. Default value = +30 db. Maximum allowed value = +120 db. Note: the *initial* gain value is set via `-g`.
- `-n` Perform normalise only (dynamic compression is disabled). The gain is adjusted to ensure that the highest peak in the input WAV file reaches maximum level (or the output level as selected via `"-l"`).
- `-p` Before processing the file, display the main file parameters (sample rate, sample size, file size). Then display progress via a dotted line across the screen.
- `-r N` Ratio of fast gain recovery to slow gain recovery (db). On sustained program material, `-r` db of fast recovery (`-f`) will occur before switching to the slow recovery (`-s`) rate. Default = 1db.
- `-s N` Slow recovery time-constant (db/sec). Following a gain reduction, a specified amount of fast gain recovery occurs (as set by `-f`), and this is then followed by slow gain recovery (as set by `-s`) until full output is reached. Default = 1db per second. (See also `-r` and `-t`)
- `-t N` Transient (tracking) parameter (db/sec). Probably the most difficult to explain in humble English. Controls the speed at which the slow recovery time-constant responds to changes in program level.

The slow time-constant recovery block is supposed to come into operation when the program level falls by `-r` db. However, when isolated short transients (such as a "click" on a vinyl recording or a brief transient speech sibilant) force a substantial gain reduction, one would intuitively expect a rapid recovery back to the average program level. This is achieved by the introduction of

a "tracking" time-constant for the slow time-constant block.

This may be best explained via an example. The default value for `-t` is 26db per second. This means that for any change in program level, the slow time-constant *reference value* will track any change in program level at a rate of 26 db per second.

So, for example, if an *isolated 10 mS transient* of +20db occurs, the reference point for the slow time-constant will shift at a rate of 26 db/sec for 10 milliseconds. In total, the slow time-constant *reference point* will move up by +0.26db.

Because this shift is fairly negligible, once our 10mS transient has disappeared, fast recovery will now occur for the entire 20db gain-loss that we just suffered.

The difficulty in "tuning" `-t` relates to the type of program material being processed, of course. The default (26 db/sec) is fine for speech and pop music, but it may need to be increased for "serious" (classical) music, or decreased if transient "clicks and pops" are causing problems.

Usage examples

level myfile

This just levels the input file `myfile.wav` to `c_myfile.wav`, accepting all the built-in defaults, ie: `r = 1db`, `f = 12db/sec`, `s = 1db/sec`, `g = 6db`, `m = 30db`.

level -s 0.02 -r.5 -g3 -m9 file

Processes 'file.wav' into 'c_file.wav' to maximise the average level somewhat less obtrusively. The initial (fast) gain-recovery for sustained signal peaks still = 12db/sec (the `-f` default), but it then slows to 0.02db/sec (`-s 0.02`). Initial gain (`-g`) = 3db. The maximum gain during quiet passages will be +9db. The fast/slow ratio (`-r`) has also been halved to 0.5 db.

level -f 30 -s 15 -r6

Quite dramatic; will sound like some of the more punchy rock/metallica radio stations (heavy volume pumping). Fairly insane for music. May be useful for live speech that has widely varying levels.

level -f 20 -s 2 -r 1

More gentle, but still levels mercilessly. Good for barbeques and listening in the car.

Usage hints

The term 'limiting threshold' is relatively easy to visualise in the case of analogue (real) peak-limiting amplifiers, because they always have various real-time *indicators* for displaying input and output level and "db of limiting". Hence, one can simply feed in program material and make adjustments until the degree of limiting looks (and sounds) appropriate. The *limiting threshold* is easy to understand – it simply corresponds to the point at which the "db of limiting" indicator first begins to kick away from zero.

Level version 2.71 is command-line driven only and has no "db of limiting" indicator as such, so best results require some care and with some music may require a number of runs before the most satisfactory result is achieved.

One simple test that *can* be carried out in advance is to type "level -pn ..." to do a dummy normalisation run (with parameter display). This will tell you how much gain was required for *normalisation*. That figure may then be used as the `-g` (initial gain) parameter for your compression processing run. If you were to set the `-m` value to a value 6db higher than this, then the maximum compression will be 6db.

Acknowledgement

As a teenager in 1968, I sent a letter of enquiry to RCA America and was totally amazed to receive an original, full manual of the **RCA Broadcast Equipment Automatic Audio-Signal Processing Systems** range,

complete with full schematics and descriptive information! This led me into a good deal of private research, construction, and fun during the ensuing years. I have no idea who it was at RCA America who took the decision to send this manual to me, but I am forever in their debt. (I had sent similar enquiries to various electronics companies in the UK and drawn a complete blank.)

Caveats

Although **level** may be used as an active *audio compressor* or *AGC* system to maintain constant audio level, the infinite compression ratio can sound less than satisfactory with some types of material. In such cases, one may need to set the various control parameters rather critically and do several processing runs before the result is satisfactory.

Copyright

Copyright (C) 2001-2004 Bluehaze Solutions (Australia).

See web page at <http://www.bluehaze.com.au/unix/level.html> for latest usage and pricing details.

NAME

level version 2.71 – process an audio input file in WAV format to normalise the signal level

SYNOPSIS

level [*options ...*] [*input file(s)* or *]

DESCRIPTION

level processes WAV format audio file(s) into WAV format output file(s) with the audio signal level either normalised (fixed gain) or peak-limited (gain is continuously adjusted automatically).

File names may be specified individually, or an asterisk (*) may be used as a wild-card to process all WAV files in the directory. The .WAV file suffix is not required with file names. **level** processes each specified file into an output file of the same name but with a "c_" prepended. If "*" is used, WAV files with names starting with "c_" are skipped for input.

Limiting mode

Unless the `-n` (normalise only) flag is used, *level* operates in peak-limiting (infinite audio-compression) mode. Under these conditions, level acts as a mono-band peak-limiting amplifier.

Peak-limiters operate as normal amplifiers for signals below the defined limiting threshold. For signals above the limiting threshold, the instantaneous gain is reduced in exact proportion to the increased input, which prevents the output signal from rising further. This gain reduction is achieved with no audible distortion by maintaining its value at the new level for a period of time.

Limiting slope

With analogue peak-limiting amplifiers, finite-gain feedback loops theoretically restrict the performance so that, above limiting threshold, an increase in the input signal will still cause a finite increase in the output signal. The ratio of *the rise in output* to *the rise in input* (for an input which is above the limiting threshold) is called the **limiting slope**. This slope must be at least 10:1 for an amplifier to be useful as a peak-limiter. Modern designs achieve limiting slopes in excess of 100:1. With a limiting slope of 100:1, an input that increases by 25db will result in an output increase of 0.25db.

level has an infinite limiting slope.

Attack time

Another key parameter for specifying the performance of a peak-limiting amplifier is its **attack time**. This determines how quickly the amplifier can reduce its gain in response to a signal increase. Ideally this time interval would be zero, but analogue peak-limiting amplifiers have feedback loops with a finite (non-zero) response time. Early valve designs could take as long as 5 milliseconds to respond to a signal peak. Current designs measure their response times in microseconds. During this *attack time* period, the output momentarily overshoots the nominal maximum value.

level has a response time which is instantaneous in the sense that the gain reduction will always occur fully along the leading edge of the peak of the half-cycle which is attempting to exceed maximum output. The output remains at exactly the maximum value during this time, ie: there is no overshoot.

Gain recovery

For an input signal which is (on average) always above the limiting threshold, each time a peak-limiting amplifier reduces its gain to maintain correct output, it must subsequently *restore* the gain when the signal falls below the specified maximum output. The *rate and manner* of gain restoration must be carefully controlled or waveform distortion (both harmonic and intermodulation) resulting from the gain restoration process may become audible.

The normal approach taken with peak-limiting amplifiers is to employ *multiple* (2 or more) coupled time-constants in the gain-recovery block.

The main high-speed time constant is used to permit a controlled amount of rapid gain-recovery from short isolated transients. In **level**, the *amount* of this high-speed recovery (in db) is selected via `-r`, the fast/slow ratio (default=1), and its *rate* in db per second via `-f` (default=12).

A *sample-and-hold* circuit operates in parallel with this high-speed time constant to delay recovery for a fixed period of 40 milliseconds. This minimises audible distortion on sustained low frequency sinusoids.

The sample-and-hold timer is reset each time full output is reached, so for a steady-state sinusoid which is already at full output, it will defeat the recovery time-constants indefinitely.

A longer time constant operating *in parallel* with these provides a much slower gain recovery for sustained program peaks or longer term drops in level. The slow time-constant rate (in db/sec) is controlled by the -s value (default=1).

The combination of these recovery controls provides arbitrarily fast recovery across the top few db of the signal in combination with a slow and less obtrusive recovery for the tracking of sustained drops in level. This reduces *gain chasing* effects, such as bass or percussive instruments audibly modulating the volume of other signals.

Normalise mode

If normalise mode is used (via -n), dynamic compression (peak-limiting) is **disabled**. In normalise mode, the input file must be read twice. The first scan locates the highest sample value (+ve or -ve). The file is then read again with the gain preset to the appropriate value for normalised output, and the result written to the selected output file.

The -l (output level) control is still available in normalise mode.

Options

- d Debug. Not recommended at all. Generates vast amounts of data on standard output, and will multiply processing time by several orders of magnitude. For emergencies only – absolutely useless to anyone except the author.
- e Print some *examples* of parameter usage (see also -h)
- f N Fast recovery time-constant (db/sec). Following a gain reduction, -f defines the speed at which gain will recover until either (a) the slow time-constant takes over, or (b) maximum output is reached – whichever occurs first. Default = 12db per second. (See also -r and -t)
- g N Initial gain value (+db). This is the equivalent of the input attenuator in a real (hardware) peak-limiting amplifier. Default = +6db. If you set this to more than the -m value, it is silently reduced to the -m value.
- h Print some basic help re parameter usage
- l N Line level output value (-db). Acts as an output attenuator, with a default setting of 0db, corresponding to full output. If a value is specified, it will reduce the output level by that amount. For example, "-l6" will set the output level to -6db, ie: to 50%.
- m N Maximum gain (+db). Specifies the gain ceiling. Default value = +30 db. Maximum allowed value = +120 db. Note: the *initial* gain value is set via -g.
- n Perform normalise only (dynamic compression is disabled). The gain is adjusted to ensure that the highest peak in the input WAV file reaches maximum level (or the output level as selected via "-l").
- p Before processing the file, display the main file parameters (sample rate, sample size, file size). Then display progress via a dotted line across the screen.
- r N Ratio of fast gain recovery to slow gain recovery (db). On sustained program material, -r db of fast recovery (-f) will occur before switching to the slow recovery (-s) rate. Default = 1db.
- s N Slow recovery time-constant (db/sec). Following a gain reduction, a specified amount of fast gain recovery occurs (as set by -f), and this is then followed by slow gain recovery (as set by -s) until full output is reached. Default = 1db per second. (See also -r and -t)
- t N Transient (tracking) parameter (db/sec). Probably the most difficult to explain in humble English. Controls the speed at which the slow recovery time-constant responds to changes in program level.

The slow time-constant recovery block is supposed to come into operation when the program level falls by $-r$ db. However, when isolated short transients (such as a "click" on a vinyl recording or a brief transient speech sibilant) force a substantial gain reduction, one would intuitively expect a rapid recovery back to the average program level. This is achieved by the introduction of a "tracking" time-constant for the slow time-constant block.

This may be best explained via an example. The default value for $-t$ is 26db per second. This means that for any change in program level, the slow time-constant *reference value* will track any change in program level at a rate of 26 db per second.

So, for example, if an *isolated 10 mS transient* of +20db occurs, the reference point for the slow time-constant will shift at a rate of 26 db/sec for 10 milliseconds. In total, the slow time-constant *reference point* will move up by +0.26db.

Because this shift is fairly negligible, once our 10mS transient has disappeared, fast recovery will now occur for the entire 20db gain-loss that we just suffered.

The difficulty in "tuning" $-t$ relates to the type of program material being processed, of course. The default (26 db/sec) is fine for speech and pop music, but it may need to be increased for "serious" (classical) music, or decreased if transient "clicks and pops" are causing problems.

Usage examples

level myfile

This just levels the input file `myfile.wav` to `c_myfile.wav`, accepting all the built-in defaults, ie: $r = 1$ db, $f = 12$ db/sec, $s = 1$ db/sec, $g = 6$ db, $m = 30$ db.

level -s 0.02 -r.5 -g3 -m9 file

Processes 'file.wav' into 'c_file.wav' to maximise the average level somewhat less obtrusively. The initial (fast) gain-recovery for sustained signal peaks still = 12db/sec (the $-f$ default), but it then slows to 0.02db/sec ($-s 0.02$). Initial gain ($-g$) = 3db. The maximum gain during quiet passages will be +9db. The fast/slow ratio ($-r$) has also been halved to 0.5 db.

level -f 30 -s 15 -r6

Quite dramatic; will sound like some of the more punchy rock/metallica radio stations (heavy volume pumping). Fairly insane for music. May be useful for live speech that has widely varying levels.

level -f 20 -s 2 -r 1

More gentle, but still levels mercilessly. Good for barbeques and listening in the car.

Usage hints

The term 'limiting threshold' is relatively easy to visualise in the case of analogue (real) peak-limiting amplifiers, because they always have various real-time *indicators* for displaying input and output level and "db of limiting". Hence, one can simply feed in program material and make adjustments until the degree of limiting looks (and sounds) appropriate. The *limiting threshold* is easy to understand – it simply corresponds to the point at which the "db of limiting" indicator first begins to kick away from zero.

Level version 2.71 is command-line driven only and has no "db of limiting" indicator as such, so best results require some care and with some music may require a number of runs before the most satisfactory result is achieved.

One simple test that *can* be carried out in advance is to type "level -pn ..." to do a dummy normalisation run (with parameter display). This will tell you how much gain was required for *normalisation*. That figure may then be used as the $-g$ (initial gain) parameter for your compression processing run. If you were to set the $-m$ value to a value 6db higher than this, then the maximum compression will be 6db.

Acknowledgement

As a teenager in 1968, I sent a letter of enquiry to RCA America and was totally amazed to receive an original, full manual of the **RCA Broadcast Equipment Automatic Audio–Signal Processing Systems** range, *complete with full schematics and descriptive information!* This led me into a good deal of private research, construction, and fun during the ensuing years. I have no idea who it was at RCA America who took the decision to send this manual to me, but I am forever in their debt. (I had sent similar enquiries to various electronics companies in the UK and drawn a complete blank.)

Caveats

Although **level** may be used as an active *audio compressor* or *AGC* system to maintain constant audio level, the infinite compression ratio can sound less than satisfactory with some types of material. In such cases, one may need to set the various control parameters rather critically and do several processing runs before the result is satisfactory.

Copyright

Copyright (C) 2001-2004 Bluehaze Solutions (Australia).

See web page at <http://www.bluehaze.com.au/unix/level.html> for latest usage and pricing details.