

# Noise Reduction Tools & Techniques

Tips & Tricks By [Paul White](#) - Sound on Sound Magazine - Published [January 2012](#)

Modern recording systems are clean and quiet, so electronic noise in the recording chain is rarely a problem. However, what was once an acceptable amount of noise in a mix (because it could be masked by tape hiss, for example) will now stick out like a sore thumb. Many of us also face challenges from noises such as the hum and hiss of guitar amps, mysterious digital clicks, camera whine, air-conditioning noise and so on. Audio-restoration software is quite affordable now, but you still need to know how to get the best from it, so in this article I'll discuss the pros, cons and applications of various noise-removal tools.

## Gating & Expanding



*CEDAR Audio's ground-breaking noise-reduction software has proved so useful that they've been able to release dedicated hardware noise-reduction processors like the DNS1000 pictured here, to speed up 'cleaning' tasks in professional broadcast, restoration and mastering work.*

The humble noise gate (or simply 'gate'), relies on modest levels of noise being masked when the wanted audio is present, and being audible only during pauses. The gate opens as the signal rises above a threshold, and closes when it falls below it. If the threshold is set just above the noise floor, the gate will close during pauses and open as soon as the wanted audio is detected. To avoid clicks due to the gate opening very quickly, and to prevent slow fades being cut off abruptly, most gates have adjustable attack and release times.

Masking requires the wanted sound to be louder than the noise in the same frequency range, especially in the 500Hz to 5kHz region to which human hearing is most sensitive. So if the wanted audio is a low-frequency sound and the noise is a high-pitched whine, the noise will remain audible. Although you might improve the situation using a low-pass filter, this usually also dulls high-end details you want to keep.

This all means that gates are only really effective when (a) the level of unwanted noise is fairly low, and (b) the wanted audio masks the noise when the gate is open. The greater the level of background

noise or spill, the less natural the result will sound as the gate opens and closes. You can generally get away with higher levels of noise when gating short, percussive sounds than when working on delicate sources, but ultimately you must let your ears decide.

In most applications, the complete removal of noise is counter-productive, as it draws attention to the processing and becomes a distraction. In practice, it's usually far more effective and less obvious to use more subtle processing, even if the degree of noise reduction is more modest. Reducing the attenuation range of a gate to, say, 12dB (so it doesn't completely close, but just attenuates by 12dB) usually sounds subjectively better than a hard gate action.

A more sophisticated approach is to split the audio into multiple frequency bands, to create a multi-band expander or gate. Dolby did this with their CAT43, a four-band device based on the Dolby-A noise-reduction system operating in a modified replay mode. It's since been recreated as a plug-in by Waves. Four sliders set the expander thresholds for each band, and a fifth acts as a master. This doesn't work well for broadband noise such as hiss, but it can reduce the level of consistent mechanical sounds such as camera noise and lighting whistles. Heavy processing will affect the wanted audio and create unnatural tone-shifting artifacts, but where noise contamination is low or moderate it can work well, and it has proven very popular in cleaning location dialogue. Being optimised for dialogue, this approach may improve a dirty vocal but is unlikely to be the best tool for musical instruments.

## Digital Solutions

Many of the processes I've described were originally implemented using analogue circuitry, but high-power digital processing has made possible some serious advances in de-noising technology. The earliest digital methods were similar to the CAT43 process, but they split the audio into a far greater number of bands, sometimes several hundred. Many more affordable plug-ins, including the broadband section of BIAS Sound Soap Pro, work in this way. CEDAR's DNS series of dialogue noise-reduction systems, which are very popular in broadcast and film production, also use multi-band dynamic noise-removal techniques.

It would be impossible to set hundreds of expander thresholds by hand, so most systems allow you to take a 'fingerprint' of the offending noise by playing a short 'noise-only' section. This is analyzed and used to set the expander thresholds automatically. This fingerprint sample only needs to be a second or so long. Tape noise can usually be taken from just before or just after the recorded performance, as the noise is continuous, but for noisy DAW tracks where you don't have this luxury, you'll need to find an exposed section of noise. Where no solo noise sample is available, some systems offer a selection of typical noise profiles (tape, circuit hiss and so forth), which can be adjusted up or down using a single threshold fader. Plug-ins using this system usually have a graphic display to show both the noise-curve threshold and the spectrum of the wanted audio. There may also be tools to adjust the shape of the noise profile curve, to help optimize the settings for minimum side-effects.

The above approach only really works well when the noise level and noise spectrum are reasonably constant. Otherwise, the fingerprint quickly becomes invalid, and either some noise will rise above the threshold and become audible, or, if the noise level falls, some of the wanted audio may be attenuated.

Nevertheless, when used appropriately, such processors can bring about a significant reduction in noise, and because of their multi-band nature they can work across all frequencies: if no audio is present in a specific part of the frequency band, the corresponding filter bands will close down. The masking process is thus very effective, as when audio is present it only needs to mask noise residing in the same part of the audio spectrum.

Common side-effects when tackling higher levels of noise include a dulling of top-end detail (again!) and a watery 'tinkling' or 'bubbling' sound, as the various bands being muted affect the wanted audio in an unnatural way. The latter effect can be minimized by careful algorithm design, using intelligent linking between frequency bands to smooth out their effect, but it's also related to the amount of noise reduction being applied. In most cases, adjusting for a noise attenuation of around 10dB or less produces a good compromise between unwanted side effects and a useful level of noise reduction. I also find that processing in a couple of gentle passes often produces better results than a single, heavier application.

## **Variable Noise Floor**

A changing noise floor poses a greater challenge, and may be an issue when treating a complete mix in which noisy sources, such as guitar amps or old synths, are turned up and down. Companies at the cutting edge of audio restoration, such as CEDAR Audio and Sonic Solutions, have devised algorithms that can recognize noise in the presence of a wanted signal and create a dynamic noise profile to follow them.

Similarly, the faster the noise spectrum and noise level changes, the less well software will be able to track the noise. This is certainly true for many of the more affordable packages that offer an 'Auto' mode: if there's an abrupt change in noise character, the noise may become audible for a few seconds while the system responds to it, and changes its own noise-profile settings. Most noise-reduction software aimed at the home user tends to dull the wanted audio more in its automatic mode than if a fingerprint is taken, but further manual adjustments can often be made to find the best compromise settings. Popular software capable of this type of processing includes products from Izotope, TC Electronic, Waves and BIAS.

To avoid problems at the start of a piece of audio, where the system is still trying to figure out the noise profile, try temporarily pasting a piece of the audio before the start, just to give the algorithm something to work from. Once the audio has been denoised, this section can be discarded. The same strategy can be used if the noise floor changes abruptly and significantly: break the audio into sections, process as just described, then rejoin them.

## **Hum, Buzz, Clicks & Crackle**

Some noises are, by their nature, more problematic. For example: hum, buzz, clicks, vinyl crackle, and single outside noises such as a slamming door or a car horn. None of these respond well to the processing methods I've described, so specific processes have been devised for each.

Removing hum is fairly straightforward, as it generally relates to the mains (AC) supply frequency and its immediate harmonics. For example, In the US, the mains frequency is 60Hz, so the harmonics will

occur at 60Hz intervals. Narrow-band notch filters tuned to the fundamental and its harmonics will attenuate the hum while having minimal effect on the wanted audio. The narrower the filter bands, the less the wanted audio will be affected, but if they're too narrow, they may not be wide enough to accommodate small fluctuations in mains frequency (or recording speed, or sample-rate variations). It's also possible for very narrow-cut filters to introduce resonance-like artifacts.

Buzz can be more serious. The fundamental is still often related to the mains power supply, but another factor introduces waveform distortion and a more extended series of harmonics. Lighting dimmers are a common cause of buzz, as they reshape the mains waveform, leaving square edges with a lot of harmonics. Buzz can be addressed by a more extended series of filters at multiples of the mains frequency, but in my experience this is usually only partially successful, as really bad buzz has enharmonic components that escape the filters.

Clicks. An individual click can often be edited out in your DAW's waveform editor by 'drawing in' a new point, of course. However, some software can automatically differentiate between the fast rise time of a digital click and an intentional percussive sound. It can then replace the sample(s) causing the click with values interpolated from the samples either side of the click.



*De-clicking and decrackling plug-ins, such as the Sonnox De-clicker shown here, have become quite sophisticated.*

Vinyl crackles and clicks are a little harder to deal with, as they can extend over many samples and occur far more frequently. Algorithms have been created, though, which can usually discriminate between these artifacts and intentional percussive transients. As the events are longer than a single sample, the process of filling the gaps in such a way as to hide the edit becomes more complicated, but when done well it leaves a pretty seamless join. CEDAR pioneered this field with their Declick and Decrackle products, but many other companies provide similar tools. Sonnox have come up with a refinement that prevents their restoration software misinterpreting the sound of some harsh brass instruments as crackles: a display shows the detected crackles on a level/frequency graph so when a brass instrument plays, any apparent crackles appear as momentary bubbles in a specific area of the screen. It's then possible to draw a box around this area to exclude it from processing.

In the case of de-crackling software, the user has to make adjustments to find the best compromise between side effects and noise reduction. The main side effect is that where the processing is intensive, so much reconstruction takes place that the cleaned audio suffers noticeable distortion. To aid setting up, this type of software often has an audition mode that lets you hear only the audio that will be removed by the processing. If you can hear music as well as crackles, you're probably over-processing!


Clicks are dealt with using a different algorithm, as they tend to be less frequent than crackles and are usually much higher in amplitude, but, again, samples are generated to bridge the gap as seamlessly as possible. When cleaning up vinyl, it's common to treat major clicks and pops first, then address surface crackle. Sometimes a third stage of broadband noise reduction can be beneficial to reduce more general surface noise. There are many programs of varying sophistication aimed at those transferring their vinyl collections to a digital format, but you really do get what you pay for. The expensive, high-end systems are usually capable of better results, as their reconstruction algorithms produce less distortion.

## **Things That Go Bump...**

Until recently, the one type of noise that couldn't be dealt with by signal processing was the cough or door slam in the middle of, for example, a string-quartet recording. (Adobe Audition and )CEDAR tackled the impossible with displays showing the audio as spectral content against time: the different frequencies are shown vertically, and the signal amplitude as different colors. Often, the offending noise can be seen on the display as a noticeable discontinuity in part of the spectrum, and, once it's identified, the operator can draw a box around it, and either attenuate it to lie below the noise floor or replace it with material interpolated from the region around it. The idea is that only the frequencies present in the intruding sound are processed, and then replaced by what I can best describe as 'spectral fog' to match what's going on either side of the problem section.

(The Healing Brush and ) CEDAR's Retouch has been used to restore and improve countless recordings (including much of the Beatles catalogue), as well as to eliminate individual instruments from mixes to facilitate a Rock Band game! I've used the system several times to erase squeaky chairs and floorboards from classical concert recordings, to remove squealing car brakes from an ambient soundtrack, and to get rid of organ ciphers. It can also be used to remove just enough speech elements to render a profanity unintelligible in a song for radio play, without destroying the flow of the song, and to remove... well, almost anything that shouldn't be there, without damaging the audio in any audible way.

Sounds covering only a limited part of the spectrum are easier to remove without audible side effects than ones with a wide spectral spread, and short-duration sounds are easier to fix than long ones. From an operational viewpoint, every rogue sound usually has to be located and treated manually, so it can be labour intensive, but if the sound repeats regularly, you can select the offending region, and then select repeat regions at regular intervals. Even if you end up having to do your selections manually, though, you're achieving the impossible, so who's complaining?



*Spectral editors like (Audition and) CEDAR's Retouch and the one in Steinberg's Wavelab (shown here) allow you to perform incredibly precise edits.*

Several companies now offer spectral editing software clearly 'inspired by' the CEDAR model (Steinberg's Wavelab includes similar tools, and the latest version of Magix Samplitude offers spectral editing on every track), but they all seem to have their own methods of synthesizing the signal that's used to fill the part of the spectrum where the rogue sound used to be.

## Clean Up Your Act!

I'll leave you with a few guiding principles about how best to put the tools to use. First, remember that although there are ways to tackle most forms of unwanted noise, most of the tools aren't perfect, and some require a bit of skill to get the best from them — so try to get things right when recording! Second, noise is only a problem if you notice it. A little continuous tape hiss, for example, might not be a problem, and its presence might mask a few other low-level 'evils'. So use your ears to decide what's needed, rather than trying to clean every track.

If you have identified problem noises, though, all of the aforementioned processes can be used more effectively when you understand how they work and appreciate their limitations (3). Often multiple passes with a process, removing less noise with each pass, will help you keep side effects to a minimum (4). Similarly, you'll often find that some noise problems, such as guitar-amp noise, respond best to a combination of different treatments (probably de-hum, de-buzz and broad-band de-noising, in this case) (5). There's some pretty powerful software now, ranging from free plug-ins to pricey professional tools, so if your recordings suffer from noise, there's no reason not to give at least some of them a try!

## Noise Reduction In Practice: Combining Tools



*A series of separate processors was deployed on a bass track to clean it up for use in a mix. Note that the compressor is not part of the noise-reduction process. It can be worth auditioning typical processing after the noise-reduction chain, just to see how usable the resulting file is going to be!*

Here's an example of noise reduction in practice. I've placed some illustrative audio clips on the SOS web site at [/sos/jan12/articles/noisereductionmedia.htm](http://sos/jan12/articles/noisereductionmedia.htm). The example is a bass-guitar recording, recorded by an artist for a demo, with a lovely lazy character and a nice laid-back, lilting style.

However, the recording clearly suffers badly from buzz, hum and other noise, and becomes worse when adding typical mix processing such as compression. The file had already been edited, so I didn't have access to any suitable sections of audio to provide a noise 'fingerprint'.

The first step was to tackle the all-too-audible 50Hz mains hum [50 Hz in England, 60 Hz in USA], and for this I turned to Tone Boosters' excellent TB HumRemover, which allows you to apply a notch filter at a user-defined frequency, and creates notches to remove the harmonics. Achieving the best result is a matter of balancing what you remove with what you wish to retain, by altering the number of notches and their bandwidth (Q) settings, and setting the overall amount of attenuation.

I was pleased with the result, as the bass retained its character but was much more audible, but there was still audible noise at the top end. My first instinct was to see if I could use EQ to roll off the top end. This removed some annoyances, but unless I brought the filter down far enough that it interfered with the bass sound, I could still hear gremlins. To tackle the residual problems, then, I inserted an instance of TC Electronic's De-noise plug-in. This works most efficiently with a noise fingerprint, but as none was available, I bypassed the low-pass filter (so I could hear what I was doing) and waded through the presets. The 'Strong Noise' preset got me closest, and juggling the threshold and reduction level controls brought further improvement. I then switched on both the 'draw' and 'audition' modes, which allowed me to hear only the noise that I wanted to remove and to customize the curve so that the processor wasn't removing any of the wanted signal. This got me some way towards the result I wanted, and reinstating the low-pass filter after the De-noise plug-in yielded a reasonable result.

Finally, to check that the result would be usable, I jammed a compressor-limiter on the end and gave the signal a good squeeze. Although there were some audible issues, they were not show-stoppers, and I was happy that they'd be masked in a mix. Matt Houghton