What The Hell Is... Gating?

Gating is an essential technique to understand if you want to make your studio life much easier, cleaner, sharper and snappier.

Call me an old plagiarist but I believe one of the most exciting things about pop and dance music is its ability to creatively exploit almost everything around it. Other people's music, different cultures, new bits of technology...

it's all grist to the mill for the musician on the look-out for new ideas. In fact, there's something very healthy about picking up whatever happens to be lying around and using it to produce music. And where better for the musician to look than the recording studio, literally stuffed with gadgets that can potentially produce interesting noises, even if it does piss the engineers off.

God knows how many years the noisegate was left languishing around the studio, muting the odd hiss, cutting the occasional hum, before someone decided it could be used for much more interesting things. But once its potential was realised, the creative musician found himself with yet another potent weapon in his armoury.

The start of noise
The origins of the noisegate aren't altogether clear, but it seems likely it was first developed to overcome the problem of hiss and background noise on the soundtracks of early films. This could be a major distraction for cinema audiences, at times threatening to obscure the dialogue altogether.

But noise is a fact of life in audio systems generally. Though the development of low-noise chips and noise-reduction circuitry has reduced overall levels considerably over the years, for musicians using tape machines, PA systems and monitoring equipment - not to mention instruments of dubious design - noise remains a major irritant. In situations where noise is already present in an audio signal, gating is the most effective way of limiting its apparent effects, and thanks to some clever design work, modern processors are easy to set up and surprisingly effective in use.

Why do I speak of gating or limiting the apparent effects of noise, rather than reducing it or eradicating it altogether? The fact is, once noise has found its way on to an audio signal, it's virtually impossible to remove it without affecting the signal. Even with today's digital waveform editing techniques, noise superimposed over a complex signal like music can only be removed by cutting out part of the waveform and, inevitably, that means losing part of the music too.

The best we can do is either filter the noise out and hope it doesn't affect the audio signal too much, or rely on a very effective natural phenomenon known as 'acoustic masking' to convince our ears that the noise is no longer there. Let me explain...

Though our ears are extremely good at picking up sound across a broad frequency range and over a wide dynamic range, when two signals of similar frequency are heard, our brains tend to filter out the quieter of the two, or the one which doesn't appear to be part of what we want to listen to. Understanding how it does this requires an understanding of that most esoteric of sciences: psychoacoustics, and that's just a little too complex to encompass in a four-page article. All we really need to know is that it happens and that harnessing it provides us with a very effective weapon in the battle against noise.

Big breadths
In complex audio signals such as those present in music, we can normally rely on there being a sufficient breadth of frequencies to mask out the most prominent components of noise. As a result it tends to be more noticeable during the parts in the music when the level drops.

If you're old enough to remember the days of vinyl albums you'll almost certainly recall being far more aware of those annoying pops and crackles between tracks (or when the music was very quiet) than you were when the music was playing. This wasn't because the surface noise of the vinyl decreased during the louder parts, it's

simply that the music was able to mask its effects, unless a scratch produced a pop loud enough to be heard over the level of the music... which always seemed to happen after you'd lent them to friends.

Anyone who has done any home recording will have probably replicated the action of a noisegate by turning down the record level controls until just before the music starts, so as to minimise the noise going to tape. The problem, of course, is knowing exactly when the music is going to begin and being able to respond quickly enough with the movement of the controls. However, human reactions simply aren't fast enough to pull back exactly in time.

The alternative is to use an electronic circuit which can react much more quickly to an incoming signal in order to restore level virtually instantaneously. Correctly designed and set up, it is possible for the circuit to react in milli- or even micro-seconds to an incoming signal, rather than the one or two seconds it would take for a human to complete the task. It's also possible for an electronic device to act quickly to lower the overall level once signals have died away. Indeed, because audio signals usually 'die away' rather than cut off abruptly, the device can track the level of the input signals and lower output accordingly to provide a smooth fade down to silence.

Up and down like a bride's nightie

The job of the noisegate is to automatically raise and lower overall levels in response to input signals, maintaining apparent signal-to-noise ratio throughout a piece of music and creating virtual silence immediately before and after it. Early designs where effectively simple electronic switches which were either in an on or off state. This made gating a pretty uncompromising process with quiet sections often chopped off prematurely and some sounds muted altogether. As you might have guessed, this is why they were called gates. Signals were either high enough to get over the gate or low enough to be shut out by it.

Since those early devices, noisegates have grown considerably in sophistication and now offer a range of control parameters which help them cope with virtually any type of speech or music signal. Even so, the principle of gating remains the same; adjust the overall volume level in proportion to the input signal level and you lower the perceived noise.

One of the first refinements to the design of the gate was the introduction of a control which allowed you to set a threshold, above which no gating took place. This helps in reducing the overall range within which the gate operates and makes setting up the rest of the controls easier. To provide you with a visual indication of the threshold control an LED is almost always included; you may find two or more on more sophisticated models.

The problem of dealing with audio signals that don't start and end instantaneously (and of course, that's most of them) is addressed by the inclusion of attack and release controls to alter the speed at which the gate responds to the start of the signal and the rate at which it 'releases' the signal after it has fallen below the threshold level. Obviously, settings vary considerably with different types of material, but even with complex signals which change rapidly it's usually possible to select useful average levels which will allow the noisegate to track them accurately.

Gating a percussive sound such as a snare drum, for example, would demand very fast attack and release rate settings if the noisegate is to capture the initial impact of the drum and also cut off quickly enough to silence any noise which might follow. By contrast, a string or pad sound on a keyboard would require a much slower attack setting and a long release time, reflecting its slow build up and decay. Somewhere between the two, a plucked guitar string would probably demand a fast attack but a slow release as the note dies away.

On signals which fluctuate constantly in level such as heavily modulated synth sounds, there is often a problem of the gate 'chattering' as it attempts to keep track of the input. This can produce a very unnatural effect that's very annoying to listen to. Consequently, a more sophisticated noisegate will often include an additional hold control to prevent it entering its release phase until a predetermined time after the input falls below the threshold level.

Unfortunately, because the hold level remains fixed, irrespective of the waveform being gated, this doesn't provide a wholly successful solution to the problem of chattering. So manufacturers came up with a refinement in the design of the noisegate so it exhibited what's known as a 'hysteresis' effect. As soon as they have to deal with a difficult signal, they start kicking and scream... sorry, that was a rather predictable joke.

Don't get hysterical

Hysteresis is a method of ensuring that whatever your threshold level for opening the gate, the threshold level to which the signal must drop to close the gate will always be lower. It's a sort of built-in automatic difference selector to ensure your two threshold levels are never the same. How does this prevent the dreaded chattering of our noisegate? Because any fluctuation in signal level would have to be greater than the difference between the two threshold levels for it to occur. Don't worry too much if you don't fully understand this. Hysteresis is built into the design of the gate and is invariably fixed, so there are no extra controls to worry about.

A further inclusion on more up-market designs is a sidechain circuit, much the same as that included with a

compressor. This features a 'key' input which can be used to control the level of the main signal. This is often employed in the creation of 'ducking' effects, where one signal automatically reduces the level of another. The most common application of this is in creating voice-overs where the signal from a microphone automatically reduces the level of the music you're talking over, allowing it back up once you've finished speaking.

With an equaliser inserted into the sidechain the gate effectively becomes 'frequency conscious', responding to some frequencies but not others. This can be extremely useful where normal gating might be too obtrusive. Often the noise you're trying to get rid of only occurs within a limited band of frequencies so tailoring the response of the gate to act only within this range makes it a much less heavy-handed approach.

Tighter drums n' bass

One of the most surprising things about gates is the fact that they can be used to achieve effects quite similar to those of compression which is really the theoretical opposite, in terms of its effect on dynamic range. For example, noisegates can be used to give one instrument dynamic control over another: perhaps a drum pattern lending something of its rhythmic feel to a bassline. To set up this kind of effect, the bass would be fed to the noise gate as the main signal and the drum track used to provide a series of triggers, via the sidechain. By adjusting threshold, attack and release levels, it's possible to 'tailor' the bass so it locks with the drums and takes on something of their rhythmic feel.

Obviously, we're only controlling the dynamics of the bassline, but the time-related controls of the noisegate help create the impression of a tighter bond between the two instruments. If the effect becomes too exaggerated, you could try limiting the drums to just bass and snare or perhaps the hi-hat.

Applying a noisegate directly to a drum loop can also produce some interesting effects. A 'busy' loop can be thinned out using gating to produce a much tighter, more focused pattern. It should even be possible to lose a particular instrument, providing its dynamic range is lower than the rest and remains pretty constant. The obvious example would be a hi hat ticking away in the background, which can often be annoying.

Speaking of drums, one of the least attractive aspects of the gated reverb effect that became so popular a decade or so ago was that it was always associated with Phil Collins. In fact, so popular did it become that most manufacturers still include gated reverb among their presets on effects processors. But don't let this ubiquity put you off. There are still untold ways in which reverb (and other effects) can be 'shaped' by a noisegate to produce some fascinating effects.

In fact, there's a lot you can do with those parts of the sound you'd normally use the noisegate to remove: for example, a gated snare drum's reverb 'footprint' with the snare sound itself removed. Similarly, those long, cyclic splurges of noise produced by flanging after the signal has died away can be cleaned up and 'packaged' by the noisegate and used as sounds in their own right. This has a lot of potential when putting together industrial rhythm loops using non-instrumental sounds.

Another promising line of enquiry is treating samples. A gate can be used to recreate a more natural envelope for a sample which has had its attack or decay cut off prematurely. In the days of expensive RAM, samples were often chopped to the very bone in the interests of saving a few bytes. A few minutes with your gate could make a world of difference, particularly if your sampler (or direct-to-disk system) lacks the necessary editing facilities. Needless to say, you could also set about removing the noise from them. Nice, grainy, low-resolution samples are acceptable enough - especially if you're a lo-fi, grunge nut - but hiss is another matter entirely.

A little time spent with a noisegate could work wonders for your sample library. And of course, there's nothing to stop you changing sample envelopes much more radically, just as you would using the VCA on an analogue synth. In these days of pre-packed samples on CD, anything you can do to make a sample your own has to be worth the effort.

Crossing talk

To return to the slightly more utilitarian: one of the things that noisegates are indispensable for is dealing with crosstalk. It may be the sound of cymbals bleeding through on to drum tracks, backing vocals finding their way on to the lead vocal track or even the sound from a pair of monitoring headphones being picked up by a microphone. Whatever form it takes, crosstalk can be a real nuisance. Even when it's not threatening to completely spoil a performance or a recording, it only serves to 'muddy up' a mix. In that sense, you could regard crosstalk as noise and of course, that makes it fair game for the noisegate.

The difference in level between crosstalk and your main signal is usually enough to make noisegating very straightforward and effective though, of course, you need to be careful setting up the controls on something like a vocal if notes are going to be cut off. In terms of entire mixes, the noisegate is, admittedly, rather limited; the very complex signals present in a full musical arrangement make it difficult for the noisegate to respond in any

meaningful way. That said, with the right settings, it can be used to remove an annoying equipment hum at the beginning of a track or to prevent noise becoming apparent as it fades out.

In all its applications as a means of removing noise, the gate is actually much more effective than most people would believe possible. Simply allowing noise to be 'hidden' behind the music doesn't sound like any kind of adequate solution, but it actually is. And gates are now so inexpensive there really is no justification for not using them in both recording and live applications. I've heard tons of bands introduce their next number over a wash of equipment hum, chorus and flanging whooshes and general background noise, all for the sake of an inexpensive noisegate.

For recording too, you should think about having at least one noisegate to hand, preferably a stereo model which can be used with any instrument to keep noise to an absolute minimum. Once that's out of the way, you can roll up your sleeves and get creative.

Nigel Lord