Dialogue Noise Suppression

TV walls, moving lights, and the fashion for hard reflective set designs has made capturing clean dialogue a serious challenge in modern television. Digital Noise Suppression offers a practical solution as Peter Best explains.

The CEDAR DSN10000 dialogue noise suppressor

The modern broadcast environment can be a source of many types of noise and unwanted sounds. Some of this, such as room ambience and noise introduced by the re-use of existing noisy material, is unavoidable. Others, such as those caused by undiagnosed electrical faults and poorly positioned equipment can – with care – be cured at source. But many noises cannot be eliminated by careful microphone positioning, intelligent placement of equipment and good maintenance, and it is here that broadcasters must turn to signal processing to suppress or, when possible, eliminate the offending sounds.

Evolution of Noise Reduction

Traditional methods of noise reduction included the use of low-pass filters to reduce high-frequency hiss and high-pass filters to eliminate rumble. Unfortunately, depending upon the signal spectrum, these affected the wanted signal as much as the noise and nowadays this may be unacceptable to a high quality broadcaster.

The use of dynamics processors such as noise gates and full-spectrum expanders was no better, resulting in unnatural break-up of the signal or unpleasant noise pumping.

In the 1960s, the multi-band expander was developed. An early version of this split the spectrum into four bands and, if the signal level in any band dropped below a user-defined threshold, the amplitude of that band was progressively suppressed according to a simple gain ratio. This technique worked on the assumption that signals with amplitudes above the threshold contained wanted audio, while those with amplitudes below the threshold were just noise. Conceptually simple, this was high technology for its time, and a product based on this idea – Dolby Lab's 'Cat 43' – proved to be surprisingly effective. Furthermore, it was simple to use, with a physical control surface that was both immediate and intuitive, so many of these units are still in use today.

The advent of large amounts of computing power in the 1980s made a new, more precise technology possible: spectral subtraction. This technique splits the spectrum into numerous bands (typically 512 or 1,024) and uses a noise fingerprint to determine whether the amplitude in each has dropped below the threshold below which gain reduction is applied. Although expansion and spectral subtraction are usually viewed as separate technologies, one is merely an extension of the other, albeit with different operational characteristics and different side effects. At one extreme, the single-band expander generates noise pumping; at the other, the 1,024-band spectral subtractive dehisser can generate artefacts usually referred to as twitter, gurgling, and even 'space monkeys,' depending upon the nature of the input and the processing applied.

Spectral subtractive systems can be very precise about what noise they remove and can produce remarkable results, but they are not always simple to set up optimally. Furthermore, they assume that the noise content is constant over a reasonable timescale, and this is not always the case, especially when microphones are being moved or when faders are being raised and lowered.

For broadcasters, there is an even more serious limitation to spectral subtraction. The algorithm uses mathematical techniques that require significant signal buffers, resulting in a few video frames of latency. Furthermore, with very few exceptions, spectral subtractive algorithms are implemented on computer platforms that add further delay in their I/O stages. While acceptable when working off-line, this latency often renders spectral subtraction inappropriate for use when the sound and picture have to remain synchronised, especially when broadcasting live to air.

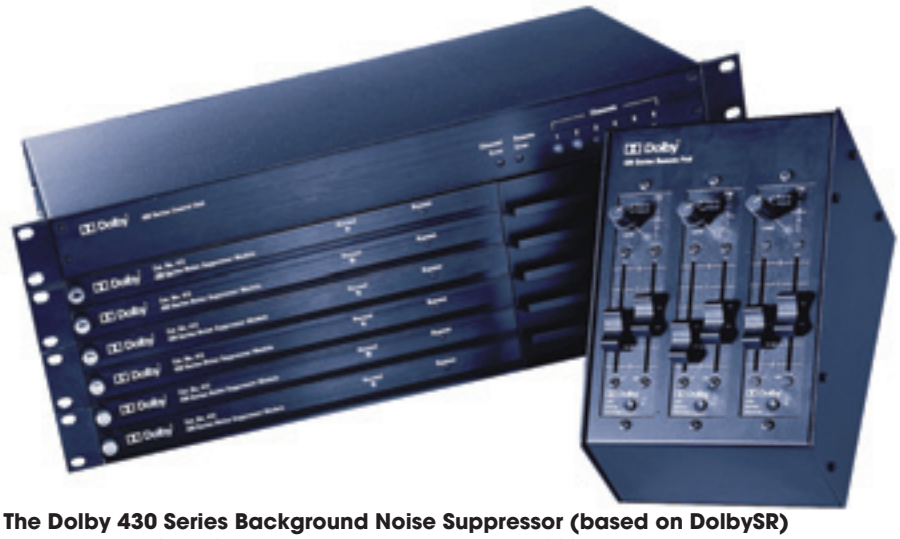

The Dolby 430 Series Background Noise Suppressor (based on DolbySR) replaced the Cat 43 unit (based on Dolby A) in 1992

Dialogue Noise Suppression

In the 1990s, many production houses and broadcasters had begun the transition from analogue to digital audio but, due to the limitations of computer-based techniques, they were (by and large) still using analogue expanders as their primary noise reduction systems. Clearly, there was a requirement for a digital solution that combined the simplicity of analogue units with a modern, digital algorithm that (as much as possible) avoided the twin evils of noise pumping and twittering in the de-noised signal. The solution also needed to exhibit near-zero latency, such that a number of units could be daisy-chained with no perceivable loss of lip-sync.

Several algorithms were proposed and at least one noise filtering product appeared in the mid-90s, the GML 9550 Digital Noise Filter. Meanwhile, in the UK, noise reduction specialists CEDAR Audio were experimenting with processes based on an intermediate number of bands. These were designed to suppress background noise in signals with a reasonable signal to noise ratio, and (as far as possible) to minimise unwanted side effects. The resulting algorithm, which they called Dialogue Noise Suppression, or simply 'DNS,' combined elements of existing analogue and digital processes, and had a negligible latency of just five samples (approximately 0.1ms at 48kHz, or 1/400th of a frame at 25fps).

DNS made its first public appearance in 2000 in the form of the CEDAR DNS1000 Dialogue Noise Suppressor. It was not the most powerful noise reduction system available – time spent on a spectral subtractive system could often yield superior results – but it was simple to use, benign in its generation of side-effects, and its near-zero latency meant that it could be used in real-time during broadcast. Furthermore, it was able to suppress a wide range of unwanted sounds including general background noise, babble, lighting buzz, and camera (shutter) noise.

"If there's a downside, it's that lighting directors now take advantage of noise suppression."

Big Brother is Listening

Oliver France of Oliver France Sound was an early adopter of the CEDAR system: "I was engaged for the first series of *Big Brother* in 2000 or thereabouts, and the directors were constantly seeking new 'looks', most of which generated large amounts of audio noise – such as big screens and lighting with built-in cooling fans. On many occasions Davina McCall was surrounded by screens and lights and the noise was challenging the programme audio to the extent that it was not fit for broadcast. The problem was to reduce the noise but still hear what we wanted to hear. I had heard of the CEDAR box, and this seemed an ideal opportunity to see what it could do. We brought in a DNS1000 and a Sonifex Redbox ADDA and the results were

amazing. We managed to filter out the majority of the noise without losing the quality of the source material – Davina's voice. Fortunately, once the soundstage was rearranged the noise was pretty constant, so we could set up the DNS1000 and leave it to do the job."

In 2006, Oliver moved from *Big Brother* to *Deal or No Deal,* where the problems were somewhat different. "The studio is built into a warehouse-type unit in Bristol, and unfortunately you can hear everything going on around it. There are also a lot of moving lights and a large screen, all generating noise. However, unlike *Big Brother*, it's possible to have up to 22 gooseneck mics open as well as Noel Edmonds' radio mic, so the noise content is constantly changing. At first, we used a DNS1000 in post-production, but after the first few shows we moved it to recording, placed on an insert of the talking group so that it did not process the music or studio applause. This was very effective, and there are now three DNS1000s installed in *Deal or No Deal.* Noise problems very often come down to the same old problems of moving lights and screens with fans in them. As long as you're careful, the results that you can get from the DNS1000 are almost too good to be true. If there's a downside, it's that lighting directors now take advantage of noise suppression, and as a consequence may be less careful with placement."

Dialogue noise suppression is equally important for outside broadcast and live event coverage. Steve Crisp, Senior Sound Supervisor for BSkyB explained: "We use the DNS1000 for numerous sports events including the Barclays Premier League. It's not necessary for the live commentary because the lip-mics we use provide great separation, and the crowd noise is not a huge issue, but studio coverage within sports grounds is a different matter. We are frequently positioned in places such as executive boxes rather than purpose-built studios and, because there is little if any acoustic treatment, these can suffer from all manner of background noises. The most common are crowd noise and air conditioning noise, but in one football stadium the studio is next to a huge outdoor screen with banks of fans that generate a peaky whine. We can't get rid of this acoustically, but we have identified the frequencies of the whine and fine-tuned the DNS1000 to remove them. In principle, we could tailor the audio using EQ, but the DNS removes the problem better than equalisation or expanding.

"In general, we apply the DNS1000 across the studio mic group so that it is cleaning the signals from all the microphones simultaneously, and the output is then compressed before being presented to the broadcast chain. However, the way we use it will depend upon the situation in which we find ourselves and, because the noise changes during the transmission, the sound supervisor will often tweak the settings on the fly. Also, we have found that if there is a lot of low frequency noise, we can attenuate this with EQ before presenting the signal to the DNS. The noise suppressor will then do an even better job with less risk of pumping."

Excess Reverberation

Another problem encountered with increasing regularity is that of excessive reverberation, as Nick Ashton, Sound Team Leader at ITN News, explains: "We produce the news broadcasts for ITV, Channel 4 and Setanta Sports News. Each has its own studio within the building, but the problems are much the same. For example, the Setanta studio is a converted office space with high-pressure air conditioning that generates an unacceptable amount of noise. Then there's the set itself; this has just fresh air above the cyc frame, which lets in ambient noise from the newsroom, and a Perspex wall and desks which creates a serious reverberation problem. To control this, we use our new DNS1500 placed across the mic group of the studio's Soundcraft desk. What it's doing is incredible, though especially challenging when we have two women presenting. To get any better results we would have to change the studio itself.

"The ITN News studio is also nearby, and that has a bank of four large back-projectors that generate a serious amount of noise. The Channel 4 newsroom is even worse, with isolation problems, a back-projector that causes us grief, and terrible reverb as a consequence of a glass wall, a partially untreated ceiling slab, and a set covered in Perspex. We use a DNS1000 in there, and that controls the noise and reverb very effectively.

"Finally, there's the atrium itself, which we sometimes use for live broadcasts. This is essentially a 10-storey hole and, in the General Election broadcast of 2001, the combination of a long reverb time and huge, fan-cooled screens could really have stuffed us up. To cure this, we used a DNS1000 in a novel way, with its left channel dedicated to the microphone of one of the main presenters, and its right channel dedicated to the other. The noise that each presenter was experiencing was quite different, but when set up in this way the DNS1000 coped admirably."

Noise suppression is required in post-production at least as often as in broadcast. In this environment, additional considerations become important. These include the ability to automate the noise reduction process, synchronisation to programme content, and integration with commonly used workstations. It may also be desirable to clean multiple channels simultaneously, either when running multiple channels of dialogue, music and effects to a final mix, or when the audio is already mixed for 5.1 or 7.1. This can be achieved using multiple DNS units, or by using a multi-channel software implementation on a workstation. Of course, the latter loses the benefit of near-zero latency, but that will be less of an issue in the post studio than elsewhere.

The Future

Digital signal processing is no longer in its infancy, but neither is it mature, and there are areas in which there is significant room for further improvement. Nowhere is this truer than for broadband noise reduction. Whereas clicks, crackle, buzz, hum and even clipping can usually be eliminated without any adverse effects on the wanted audio, there are almost always compromises to be made when removing background noise. Happily, today's tools far exceed the performance of those available just a decade ago.

jbs

Nick Ashton, Team leader ITN, in Sound Control One used for ITV News