Cutting the noise

Quantisation noise from digital musical instruments annoys recording engineers. Richard Brice outlines an electronic noise gate design which puts the dynamic range back into the recording process.

Using modern digital audio recording equipment with a dynamic range of more than 90dB, the noise sources up-stream of the recorder become a source of worry and annoyance. It becomes a constant fight to reduce noise from interference, air-borne and structure-borne noise picked up in the microphone channels and electrical noise in the analogue mixing and processing electronics.

A further source of noise in the modern audio studio is the digital synthesisers. Some of this is due to Johnson and 1/F noise in the analogue output stages of this equipment, the rest derives from quantisation noise on the audio samples.

Digital mixing processes also add their own noise. This builds up in the mixing and processing stages of these systems, the residue of truncated digital multiplications and additions.

The audible character of digital noise is coarse and unpleasant. Having chased noise contributions from microphone pre-amps and the desk electronics, it is very galling to have to put up with noise on a high level signal which enters the mixer at mix-bus voltage levels[1].

A noise gate tackles this problem. The single ended noise reduction system outlined here provides a beneficial level of noise reduction without introducing the "breathing" or "pumping" effects that so often plague noise reduction.

The design employs an unusual semiconductor device, the photon controlled bilateral analogue fet. The part combines the advantages of the photoresistor and optoisolator with the speed of a fet.

But it doesn't suffer from non-linearity of these other devices nor the complexity of circuits employing a "Blackmer" voltage controlled amplifier[2].

Theory of Operation

The noise reducer uses an expansion technique where low level audio signals are amplified less than high level signals. A threshold is set so that the residual system noise has insufficient energy to cause the variable gain amplifier to change to its higher gain regime. The design must ensure that the presence of useful audio signal raises the gain of the amplifier sufficiently quickly without destroying the transient start of the signal. The period of time the amplifier remains in the high gain regime once a high level signal has ceased must be carefully controlled. Too long and the amplifier will

Fig. 1. Stereo noise gate circuit, working as if the treble control of a pre-amp were quickly varied. The effect is of a single ended noise reduction system which cuts the treble response during quiet passages.

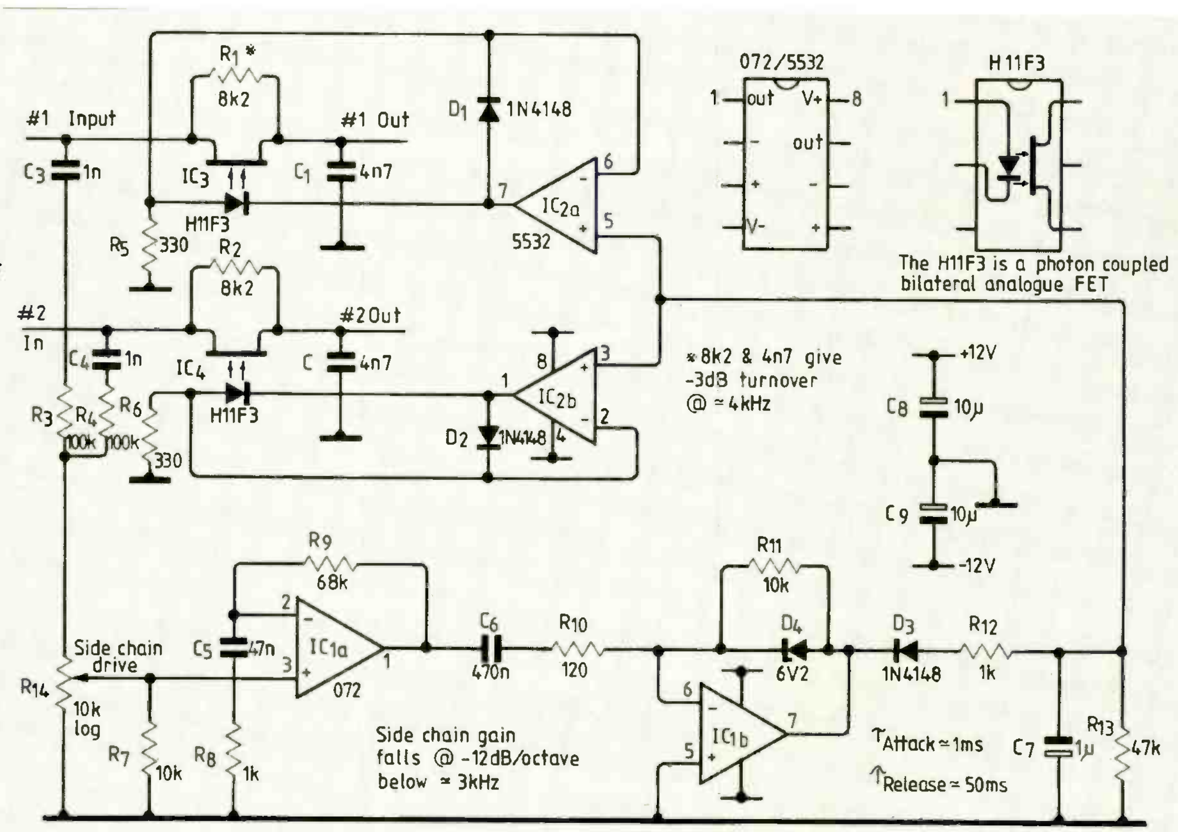

be unable to return to the low gain state in the gaps between wanted signal. Too short and the expander will mutilate the reverberant part of the audio signal.

Expander systems produce a characteristic noise signature. This comes from envelope modulation of constant input noise behind the wanted signal. When this is amplified by the changing gain of the expander, the noise takes on a varying quality which has been described as sounding like breathing or swishing. All forms of noise reduction using signal dependent amplification suffer to a greater or lesser extent from the phenomenon and it places an upper limit on the maximum expansion ratio employed before noise modulation becomes noticeable.

The noise-reducer also takes advantage of two psychoacoustic phenomena; masking and the non-linear frequency response of the ear. As a result of nature's signal processing, the presence a single pure tone at 60dB above the threshold of perception will cause the desensitisation of the ear by as much as 20dB in the octave and a fifth above the original tone. The threshold shift may be as much as 40dB near the tone frequency. Music has very many pure-tones simultaneously present but, for the majority of the time, its masking effect only operates at low to middle frequencies. This is because system noise, whether generated by thermal agitation or from quantisation errors, has a very flat energy versus frequency characteristic.

It sounds hissy to the human ear because the frequency response which, at low levels at least, is about 28dB more sensitive at 3.5kHz than at 100Hz and then falls away slowly to around -10dB at 10kHz reference 3.5kHz. Music, in order not to sound shrill and thin to our non-linear ears, must complement the ear's rising frequency response and has an average energy versus frequency characteristic which falls with frequency.

A signal which has an uneven energy versus frequency characteristic like music will sometimes fail to mask one which has: noise. This is especially significant since the music signal fails to mask the noise in the high frequency portion of the audible range where our ears are most sensitive to detecting the noise part of the signal.

The circuit shown in **Fig. 1** operates by automatically controlling the frequency response of the forward signal path of the noise reducer using controlling information derived from the high frequency content of the programme. The principle, which is used in some commercial noise reduction systems[3] works rather as if the treble control of a pre-amp is constantly and quickly varied. Thus in the absence of a music signal, the control attenuates the system noise in the HF part of the audio spectrum where it is most subjectively annoying. The control is only turned up when the wanted signal, containing high frequencies, occurs.

The circuit has this effect. If a bass guitar is the only instrument playing, the high frequency gain is not raised since a bass instrument does not require a wide frequency response for faithful reproduction, and has insufficient upper partials to mask the system noise. The subjective effect is to lose the hiss. If, on the other hand, the signal is a snare drum hit in a reverberant acoustic, control must operate very quickly to allow the high frequencies of the snare drum to pass through the system. Gradual reduction must follow... not so slow that the noise resurfaces after the masking effect of the signal has ceased yet not so fast that the reverberant tail of the signal is destroyed.

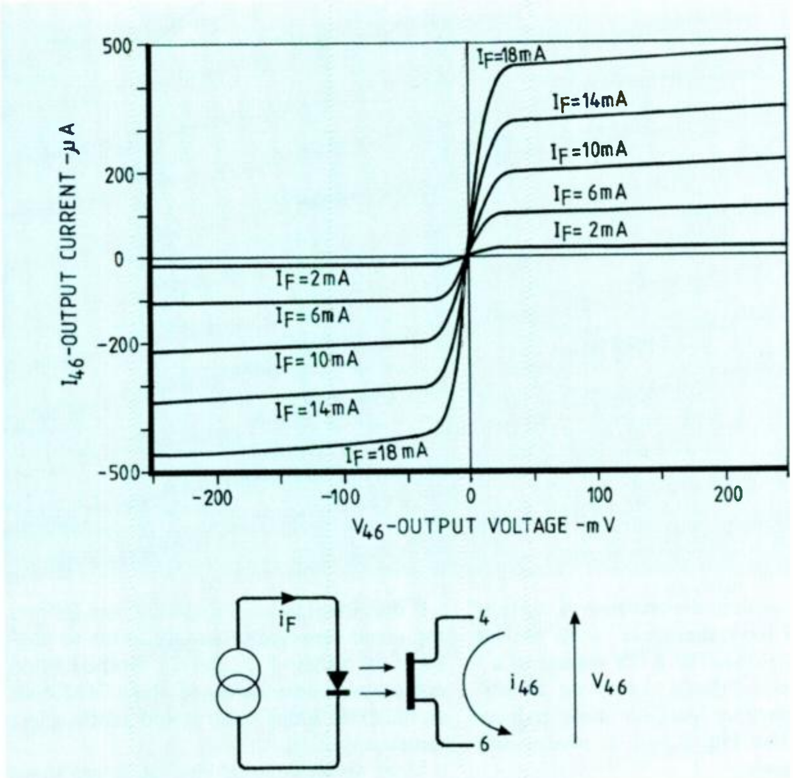

Fig. 2. Output characteristics of the H11F3 — photon coupled analogue bilateral fet.

The active device

The gain controlling element in the noise-reducer is the H11F3 made by Harris Semiconductor. The part rejoices in the description photon coupled bilateral analogue fet. It integrates a gallium arsenide led with a symmetrical bilateral silicon photodetector. The detector is electrically isolated from the control input and performs like an ideal isolated fet. The part is designed for control of low-level AC and DC analogue signals. It responds in less than 15µs.

Figure 2, device output characteristics, demonstrates that it looks like a fet operated below pinch-off and controlled, not by a gate potential, but by the current through the led. Because the element is not a triode, and controlling circuit need not share a terminal with controlled circuit, the fet may be used bi-directionally, with current flowing from source to drain or from drain to source. Practically, within certain operating limits, one may think of the device as a variable resistance element. The huge advantage that these parts have over normal fets is that control can be effected by a completely isolated circuit. It is also much simpler than the traditional variable transconductance multiplier arrangement.

Low-Pass Circuit

The device resistance is controlled by the current flowing through the integral diode. The turnover frequency, in the absence of signal, is defined by R_l and C_l since, when no current flows in the led, the fet has a very high resistance indeed. In this design, the turnover frequency is set on the high side: about 4kHz.

There is a complex compromise to be made between signal level handling capacity, the effectiveness of the noise reduction and the minimising of modulation noise.

The prototype unit was designed to operate directly at a line level of 0VU = -10dB(V). I settled on the 4kHz turnover empirically.

Refinements

Operation of the fet directly at -10dB(V) line level is a little optimistic. The device is specified in terms of the maximum allowable RMS signal current and maximum allowable RMS signal voltage it can handle linearly at specified resistances. Consider two cases: 1) a single tone input 1kHz at -10dB(V) and 2) a single tone input 4kHz at -10dB(V).

Case 1: I have assumed that the side-chain sensitivity has been set so as to start widen-

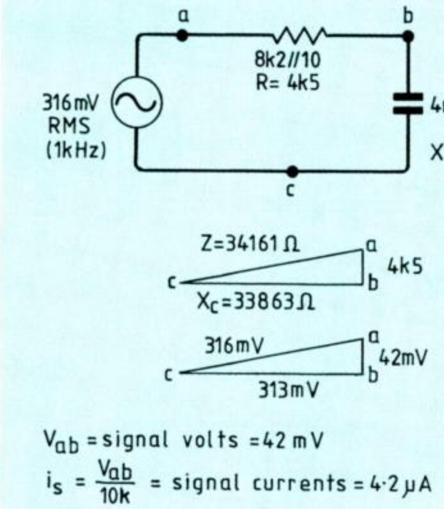

Fig. 3. Noise reducer stage — assuming an infinite non-capacitive load — and its phasor diagram associated with the conditions of Case 1.

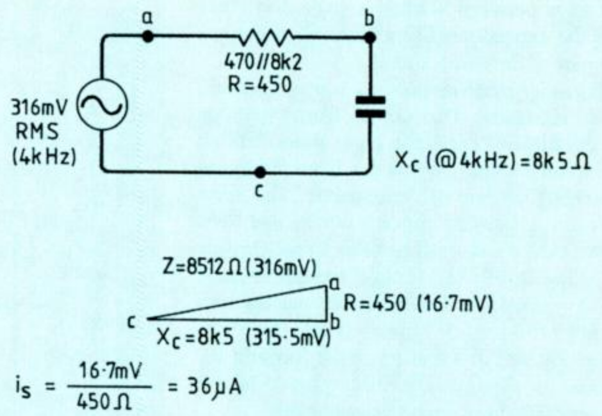

Fig. 4. Noise reducer stage and phasor diagram for Case 2. -10dB(V) tone input of 4kHz is assumed to open the bandwidth. Fet resistance will fall to an assumed 470Ω.

ing the bandwidth in the presence of 1kHz at -10dB(V). I have therefore taken the fet resistance as 10kΩ. The RMS voltage of a -10db(V) tone is 316mV. Assuming an infinite, non-capacitive load, the noise reducer stage looks like **Fig. 3** and its phasor diagram is as shown.

The maximum signal levels are therefore, 4.2μA and 42mV. The specification of the device states the maximum allowable RMS current and voltage in the linear region as 6μA and 50mV respectively.

Case 2: the -10dB(V) tone input of 4kHz is assumed to open the bandwidth right up and for that the fet resistance will fall to an assumed 470Ω. The diagrams now look like **Fig. 4**.

The maximum signal levels in case 2) are therefore 36μA and 17mV. The maximum permissable figures are 35μA and 15mV — right on the specification limit of the device and, since one must assume peak signal some 10dB higher than the steady state 0VU signal, some distortion from the circuit seems inevitable. In fact measurements failed to confirm such pessimism, with distortion being less that 0.1% in the high or low gain regime at levels of -10dB(V) at all audible frequencies. The onset of waveform distortion of several percent happened abruptly at about twice the limits suggested in the data sheet. When they appear, the distortion artifacts are predominantly 3rd harmonic.

I have been unable to detect any audible distortion from the prototype unit at all when used in the studio. I suspect this is principally due to the phenomenon noted earlier that music does not have much high-frequency energy and that such distortion as there may be is of a fairly innocuous kind. When the fet moves beyond its linear region (see Fig. 2), the element operates like a constant current source with an essentially capacitive load. The unit will slew-rate limit rather than amplitude limit and slew-rate limiting is a far less audible form of distortion than clipping.

If the circuit is left unbuffered, any following cable capacitance simply forms to augment the value of C_1 and C_2. Normal audio cable has a capacitance of about 30pF/foot so the extra 300pF or so would hardly affect circuit operation.

More serious is that the following input resistance must be high, at least 47kΩ, otherwise overall operation will be adversely affected. A buffer based on an AC coupled NE5532 makes connection far more flexible and, if the following op-amp stage is arranged to give some gain, the signal may be attenuated before entering the noise-reduction stage to improve headroom.

Side chain

Only one side chain is necessary for a stereo noise reducer. In fact, unless the control is common to both channels, the stereo image can appear to wander. The side chain comprises the mixing resistors R_3 and R_4 which take an equal proportion of the left and right signals to the sensitivity control. R_7 pads VR_1's control law to be more manageable. IC_{1A} is straight-forward HF gain stage. IC_{1B} gives a very large gain until the output swings minus 6.2V when the Z_1 goes into zener conduction and the gain reduces very quickly. The output can only swing positive by about 0.6V when the zener diode goes into forward conduction. C_6 and R_{10} roll-off the voltage gain below 2.8kHz. The IC_{1B} stage virtually provides the rectification necessary to derive a control signal. D_3 ensures no positive signals are applied to the control-voltage storing C_7. The two time constants formed by R_{12} and R_{13} with C_7 determine the attack and release times of the noise reducer.

Strictly speaking, the side-chain should full-wave rectify the audio signal since the magnitude of negative and positive peak signals can differ by up to 8dB. A purist might install another dual op-amp, the first part acting as an inverter and the second part as an identical twin to the IC_{1B} stage in Fig.1. This stage could feed R_{12} via another diode so that the most negative signal always predominated.

The control voltage stored on C_7 drives the dual op-amp IC_2. The infrared emitting diodes are connected within the feedback loop of the op-amps; the diode current is controlled linearly by the voltage on the non-inverting inputs of IC_2 and the value of R_5 and R_6. I found that no adjustment of R_5 and R_6 was necessary to achieve channel matching, so well matched were the H11F3's devices.

D_1 and D_2 are included so that IC_{2A} and IC_{2B} do not saturate in the unlikely event of their receiving a positive voltage input.

The design of the side chain ensures that nothing bottoms — good audio design requires as little spikey current about the place as possible. The led offset voltage is continually compensated for by the action of the feedback loop of IC_{2A} and IC_{2B}. This guarantees the dynamic response of the circuit is dominated by the time constants formed by R_{12}, R_{13} and C_7. Part of R_{13} could be replaced with a potentiometer so allowing control of decay action; the best value for this is not always the same for different types of music.

The circuit requires no setting up and it is very easy to operate: One simply, gently increases the sensitivity of the side chain, using VR_1, until all the high frequencies in the untreated signal seem present in the noise reduced signal. While setting the side-chain drive, I use the bypass route to compare the untreated and treated sound and to check the noise reduction operation. ■

References

1. Brice R. Audio Mixer Design. *EW + WW* July 1990. Vol 96 number 1653. pp. 575-580.

2. Blackmer D. U.S. Patent 3,714,462 Multiplier Circuits

3. LM1894 Dynamic Noise Reduction System DNR *National Semiconductor Databook: Special Purpose Linear Devices 1989 Section 1* pp. 165-172.