

**Reproduced from
Studio Sound, September 1990**

Fixing It Outside The Mix

There are times when there is no alternative to mastering from a faulty stereo recording. Michael Gerzon describes a collection of practical techniques with available equipment that improve your chances of success

It is surprising how often defective recordings need to be released commercially. Perhaps it is a live stereo recording on cassette, or a straight-to-stereo digital recording with intermittent popping sounds, or the only surviving safety duplicate Japanese copy master of a long-lost classic recording that has developed a drop-out or flaky oxide on the right channel.

Maybe that classic once-in-a-lifetime performance only exists in a recording done on a personal stereo from the audience in a semi-illegal bootleg that the band was offered after the gig. Or an anonymous Dixon's cheapest cassette, in a misaligned domestic machine, taped from the PA desk as an afterthought. Or the technically beautiful stereo digital recording incorporating a technically perfect conversation on one side of the stereo image by a member of the audience about the outrageous price of beer at the venue during the climax of the gig.

Behind the industry facade of 48-track digital technical competence and solving all problems by expertise and wizardry, sad tales of sow's-ear recordings like this are commonplace. Often, such recordings are the ones that, for reasons of historic importance, artistic excellence or uniqueness, need to be released on CD. Ask any mastering house how often such recordings are brought in needing a degree of 'fixing' to be acceptable for release and how often releases prove to be impossible. Fixing it in a mix is fine if one has a multitrack tape to mix from – but what about premixed stereo recordings with no multitrack fall-back in which the faults are inextricably already in the mix?

This article is about some techniques of fixing it outside the mix. Few mastering consoles are equipped with the facilities to do this properly – likewise most studios. Equipment designed for equalisation, processing and tinkering with multiple mono tracks is generally totally unsuitable for zapping faults buried in the middle of a complete stereo mix.

It is not my aim here to describe the use of well-known studio processes. Rather, we look at how to target faults that will respond to specific specialised treatments tailored to those faults. Nearly all these treatments involve subtractive or cancellation methods that cancel out faulty information from one part of a stereo mix by information taken from another part of the mix.

Dodgy channels

One of the most common types of faults is a stereo recording in which one channel is OK and one is, intermittently or partially, faulty. There is a simple, if drastic, fix for such recordings – use the 'good' channel, and feed it to both left and right outputs. The result will be a technically good recording with no stereo effect and losing sounds predominantly in the faulty channel.

If the dodgy channel were sufficiently bad the whole time, then maybe this is all that can be done but, in most cases, there is still useful information from the dodgy channel that can be used to rescue some degree of genuine stereo effect.

Basically, we want to bring down the level of any fault in the dodgy channel to the point where the results are acceptable, while still preserving both the centring of the stereo and as much of the original stereoism as we can rescue. Also, if a fault is present or audible only some of the time, we want to have the full genuine stereo the rest of the time, with the ability to 'fade in' and 'fade out' the fault-removing processing only when needed.

The basic philosophy is 'if it ain't broke, don't fix it'. Any processing applied should be only to parts of the sound that need fixing, everything else should be left alone.

Suppose the left (L) channel is OK but the right (R) channel has a part-time fault. Then we want to crossfade between two situations: one in which the output stereo L_{out}, R_{out} is the same as the input stereo,

$$L_{out} = L$$

$$R_{out} = R$$

and the other in which the left input L is presented in mono at both outputs when R is totally unusable, ie

$$L_{out} = L$$

$$R_{out} = L$$

In order to crossfade between these two situations we put $L_{out} = L$ the whole time, and

$$R_{out} = R + k(L - R)$$

where k is a gain that can be varied from 0 (for the original stereo) to 1 (for mono from the L channel).

Alternatively, we could put

$$R_{out} = L + k'(R - L)$$

where the gain $k' = 1 - k$ varies from 1 to 0.

The basic processing setup required to crossfade between mono from the left channel input and true stereo is shown in Fig 1. Essentially, the left input L is fed to a channel of a stereo mixer and panned to left, the right input R is fed to a second channel and panned to right, the difference signal L-R is derived by a subtraction network and fed to a third channel of the mixer, which is also panned to the right. The equalisers and faders on all three mixer channels are set to the same settings and levels. When the L-R fader is at the same level as the other two faders, the right output becomes

$$R + (L - R) = L$$

and mono from the left input appears at both outputs. As the L-R fader is pulled down, the output becomes the original true stereo.

Thus, varying the L-R channel fader between the full level and zero allows a continuous crossfade between mono from the left input and full stereo. Because $L - R = 0$ for mono central images at the input, operation of this fader has no effect whatsoever on the level or position of central mono images. Since most stereo programmes have the most important sound sources panned near centre, this means that operating the L-R fader gives minimal disturbances of the overall sound when crossfading is needed.

You may or may not realise it but your studio is likely to be full of subtraction networks that can be used in the arrangement of Fig 1. Every electronically balanced input to a piece of studio gear is effectively a subtraction network for the unbalanced signals going to the 'hot' and 'cold' sides of the balanced input. Thus, for example, the left input could be fed to pin 2 or the hot side of such an input, and the right input could be fed to pin 3 or the cold side of such an input (or *vice-versa*, depending on the signal polarity of the equipment) with the processing of the piece of equipment switched off or bypassed. I use a spare channel of Symetrix 544 with the gate/expander switched out as a subtraction device (see Fig 2).

One may have to wire-up special leads or adaptors to derive the unbalanced feeds and to feed them to the two sides of an electronically balanced input. Also, one must use equipment that has unity gain and little phase shift in the audio band when its processing is switched out so the cancellation of L-R with R works well across the audio band. Also note that one should beware of 6dB

gain differences that can be caused in Fig 1 if some signals are fed to the mixer in unbalanced form and some in balanced form. The acid test is to ensure, with the L-R fader at the same level as the other two, that a right-only input is cancelled out at the right output. Otherwise, one has to use a '0dB reference level' for the L-R fader that differs by 6dB from the other two.

If a fault on the right input is intermittent, one can fade the L-R fader up shortly before the fault starts and down again after it ends to preserve stereo for as much of the time as possible. Also, if the fault is at a lowish level, by fading up the L-R fader to a slightly lower level, one can preserve some of the stereoism while reducing the level of the fault significantly. Since the fading action is smoothly continuous, one can control the degree of processing in real time, rather than having to edit between two dubs. It helps to have the master being processed in a form that has a precision time read-out so that the crossfades can be rehearsed.

A rather more sophisticated version of this processing is possible if the faults on the right channel only affect some frequencies. For example, a recording might have excessive hiss or rumble only in the right channel, or the interference or unwanted audience noises might be concentrated over a fairly well-defined frequency range, or the drop-outs might only affect the higher treble frequencies with little effect on bass frequencies. These are all quite typical situations encountered with faulty material. All one does is insert an equaliser (eg a graphic equaliser) in the L-R signal path of Fig 1, giving the schematic in Fig 3.

The equaliser used should be such that, when set out flat, it should have virtually no effect on polarity, amplitude or in-band phase response on the audio signal passing through it. (Digital equalisers invariably incorporate large time delays and so are totally unsuitable here. Also beware of polarity-inverting equalisers and of equalisers incorporating non-switchable lowpass and highpass filters, since these spoil the phase response at the flat setting.)

When set to flat, the modification of the right channel works just as before but if one cuts some frequencies on the graphic equaliser, the processed output reverts to stereo at those frequencies. For example, if one wishes to 'monoise' faulty treble frequencies only, then the bass frequencies on the equaliser should be cut. In practice, graphic equalisers are particularly suited for 'tuning' quite precisely what frequencies are monoised from the left input and what are allowed to remain in true stereo.

In general, only the most audibly disturbing frequencies should be highly processed. For example, if the problem is excessive hiss in the right channel, it may be enough to use the graphic to allow through a narrowish frequency band around 8kHz, leaving both lower and higher frequencies (which contribute less to the audibility of hiss) in true stereo. Also, in passages where the fault is absent or masked by the wanted information, the L-R fader can still be pulled down to restore the original stereo at all frequencies.

Fig 4 shows a version of the processing that can be used to 'tune out' faults in the left or the right channel, using an R-L signal path panned to the left to provide a similar cancellation when the faults are in the left channel. Although, in general, faults in both channels at the same time are much more difficult to deal with, Fig 4 does allow processing of faults in one frequency band on the left channel while other faults in a non-overlapping frequency band occur on the right channel. This does not happen very often but Fig 4 allows processing of recordings where faults occur in both channels at different times without the need for any switching or rewiring.

More advanced stuff

The 'cancellation' processing for faults in one channel only as described above, where faults in say R are cancelled out by mixing it with variable amounts of L-R, can be extended to more sophisticated processing. For example, suppose that the fault is audible only at low levels and is masked by the wanted signal at higher levels. In this case, one wishes to switch out the processing at high levels. For example, one might have a low-level buzz, crackle, hum, rumble or hiss in the right channel that is masked when the music gets loud.

In this case the procedure is to insert a limiter in the L-R signal path of Fig 1 or Fig 3, with the usual provisos about the limiter being non-inverting with little phase shift in the audio band and having unity gain below the limiting threshold. Then, above the preset threshold level, the L-R signal is, in effect, attenuated, leaving a substantially full stereo effect without any cancellation. Only at low levels when L-R is below the threshold level will the partial (if a graphic is used) or total (if no graphic is used) cancellation of the R channel occur to a full extent.

This way, for example, one can tune the graphic equaliser to remove only the most audibly disturbing part of a buzz on the R channel, and then set the limiter threshold to remove the processing at higher music

levels, leaving a minimally processed signal in which the buzz on the R channel is audibly suppressed. If the limiter is placed after the graphic equaliser, as in Fig 5, it will respond predominantly to music energy in the same frequency band as the frequencies of the buzz that is being suppressed, which are those components that will tend to mask the buzz best. An additional sidechain equaliser for the limiter, as shown in Fig 5, can emphasise the effect of those specific frequency components further – I have found that connecting the two channels of a consumer stereo graphic equaliser in series (with an attenuator between the two to prevent overload) is necessary to achieve a sufficient boost of some frequencies and attenuation of others in such a sidechain equaliser, where emphasis of the desired band by up to 50dB can be needed for best results.

The arrangement of Fig 5 can also be used to perform the reverse trick of getting rid of high-level occasional thumps or bangs in the right channel by replacing the limiter by a noise gate. In this case, one wishes to switch off the L-R signal except when the thump occurs. This is done by using a noise gate with a rapid attack time in the L-R path with the threshold set so the gate is permanently in the 'off' state except when the high level thump signal occurs. Since any central mono information is absent in L-R, the music-signal level it contains will often be quite low, making it easier to detect the presence of high-level thumps or bangs.

When such thumps or bangs do not have a much higher level than the music signal this processing becomes more difficult but two things can help increase the chances of it working. One is to use a sidechain equaliser for the noise gate that is flat at most frequencies (so that the start of the transient of the bang or thump is responded to fast) but with the highest-energy music signal components, *ie* those below 1kHz, attenuated. The other thing one can do is use an expander rather than a gate, which gives a smoother transition between stereo and mono from the left channel at high signal levels, at the expense of a tendency for the signal to monoise during high level music. However, this problem can be minimised by manually adjusting the expander threshold immediately before and after each thump or bang.

The above examples should by now have given a general idea of how faults in only one channel can be processed away and only at frequencies and signal levels at which they are troublesome. Obviously, I can't cover all possible faults of this kind that might be encountered, but the general principles of processing *only* the L-R

signal to cancel out faults in the R channel should allow many other cases to be dealt with without any disturbance of the central part of the stereo stage. It does not seem generally realised by recording and mastering engineers that processing of the difference signal L-R or R-L is the way to deal with faults on only one of the two stereo channels – but any processing that is to avoid disturbing the centre of the stereo stage must process only that signal that is always zero for such central stage images. The trick in all this processing is in how the processed L-R signal is mixed back into the stereo.

The processing described so far is certainly not the end of the story – just the start. Before looking at still more advanced tinkering, let's have a look at what can go wrong with what I have described so far.

Problems

The central images remain unaffected by the above processing because it only affects L-R, which equals zero for such images. However, this only applies if the two channels have substantially the same amplitude and phase response. With poorly adjusted analogue media, or with digital media without CTC correction, or with microphone techniques that are not precisely coincident or with mismatched microphones, this is not going to be the case. If there is a time delay between left and right, eg due to an azimuth error on tape or cassette, then L-R is no longer zero for central sounds and there will be some comb-filter colouration in the right output channel when the L-R gain k lies between 0 and 1.

One moral here is the need to adjust tape azimuth (and, in the case of copy tapes that may already incorporate recorded interchannel errors, the interchannel time delay) carefully before using the above processing, so as to maximise the cancellation of central sounds in the L-R signal. It should also go without saying that it is important to make sure that the two channels are not out-of-phase before processing. It helps to be able to monitor the L-R and L+R signals separately, or to examine the stereo signal on an XY oscilloscope display (or a similar display such as *The Box* or a phase meter).

Even when there are no interchannel gain or phase errors, a second problem is that the above processing causes a loss of stereo width, since whenever a fault occurs in the (say) right channel, the result tends towards mono derived from the left channel. In the extreme case when the right channel is being totally suppressed, such a reversion to mono is inevitable but there is no reason why we should accept the degree of

loss of width that occurs when the right channel is only somewhat reduced rather than suppressed altogether.

Consider, for example, the case where a right channel defect is reduced by 6dB rather than totally suppressed. This is the case for $k=0.5$, for which

$$L_{out} = L$$

$$R_{out} = R + 0.5(L - R) = 0.5L + 0.5R$$

Here, any wanted sound on the right at the input still comes from the right (at a level of -6dB) at the output, and any central sound $L = R = M$ still comes from the middle (with unaltered level) at the output. However, a sound coming from the left at the input now emerges (at a level of about +1dB) from the output panned only about a third of the way over to the left, since $L_{out} = L$ and $R_{out} = 0.5L$.

The subjective effect of this quite severe narrowing of the left side of the stereo image is a definite marked loss of stereo width. This loss of width can be compensated for by following the processing by a width control to increase the width again. Width control will be discussed more later but there is a simple trick to achieve an increase of width in the above situation without the use of any additional signal processing.

The trick is to pan the position of the L-R signal in Fig 1 to a position between right and centre, rather than to hard right as shown. This has the effect of introducing some antiphase crosstalk of the right input onto the left output, which has the effect of increasing the width at the expense of preventing a complete cancellation of the R signal at the full L-R fader gain. The operation of the panning of the L-R signal to half-right is not altogether obvious without looking at a detailed worked example.

By way of example, suppose the panpot has a constant power law, and the L-R signal is panned to half-right. Then the L-R signal is fed to the left output with gain $0.38k$, and with gain $0.92k$ to the right output, where, as before, k is the gain of the L-R fader. This gives

$$L_{out} = L + 0.38k(L - R) = (1 + 0.38k)L - 0.38kR$$

$$R_{out} = R + 0.92k(L - R) = 0.92kL + (1 - 0.92k)R$$

For example, if the L-R fader gain is -5dB (ie $k=0.562$), then

$$L_{out} = 1.22L - 0.22R$$

$$R_{out} = 0.52L + 0.48R$$

which cuts the R gain by -5.6dB and boosts the L gain by 2.4dB. This also positions the L input rather further over to the left than before, and positions the R input beyond the right loudspeaker at the output due to the antiphase crosstalk of R onto the left output. This results in a wider overall image than before while still cutting

the level of the R signal in the output.

Thus, although panning the L-R signal towards the centre prevents a complete cancellation of the R signal, it does allow an increased width when one only needs to reduce the R level to tame a fault rather than to eliminate it altogether. In general, when the L-R panpot is set to a position partway to the centre, the best subjective cancellation will be at a fader setting k less than the full gain $k = 1$. For the half-right case considered above, setting k to a maximum of about -5dB gives the best cancellation.

One can still use all the techniques described – varying the L-R fader setting and using a graphic equaliser or dynamic processor to fade the reduction of the R-channel fault in and out. While this will no longer permit a complete cancellation of any R channel fault, it does permit a reduction of the fault still without affecting central images (since $L - R = 0$ for such images) and with a reduced effect on the stereo width as compared with the processing described earlier. One now sets the L-R panpot position to achieve the desired maximum degree of R cancellation before switching in the other processing.

There is a bonus in panning the L-R signal partway to the centre. Earlier, I mentioned that timing or phase errors between the two channels, which are inevitable with some microphone techniques, caused comb filter colourations when k lies between 0 and 1. However, the effect of putting some antiphase crosstalk of R onto L is to cause the comb filter colourations on the two channels to be complementary, *ie* frequencies cut on one channel are boosted on the other and *vice-versa*, thereby making them much less audible. In particular, panning the L-R signal to half-right, and changing the L-R fader gain anywhere between zero ($-\infty\text{dB}$) and -5dB causes little colouration even when applied to stereo made with microphones having a significant spacing.

Besides its use in reducing faults on one channel, this kind of processing is equally applicable to altering the level balance between the left and right sides of a stereo image without affecting central images either in level or position. This kind of 'asymmetry control' is subjectively a better way of altering left/right level balance than ordinary balance controls when the problem is not caused by channel gain errors but by prior misjudged mixing decisions or by a natural imbalance of sound levels arriving at a well-adjusted stereo microphone. In particular, asymmetry control does not cause the 'lop-

sided' ambience reproduction of natural acoustics when an ordinary balance control is used with a stereo microphone.

One application of asymmetry control is to reduce unwanted live sounds on one side of the stereo stage while still giving a reasonable portrayal of ambience.

To summarise the above, asymmetry control to reduce sounds from the right stereo channel can be achieved by using the arrangement of Fig 1 but setting the L-R pan control to half-right position and the L-R gain between zero ($-\infty\text{dB}$) and -5dB of the level of the left and right faders. Selective asymmetry control of only some frequencies can be achieved by inserting a (non-inverting) equaliser in the L-R signal path. Faults on the left channel can be similarly controlled by using an R-L signal panned to half-left.

Crossfade techniques

The above techniques of reprocessing faults in one stereo channel are special cases of more general 'crossfade' reprocessing techniques that have many other uses. To consider the general case, suppose that a mono or stereo signal is to be reprocessed in two possible ways, which I shall refer to as A and B. A might consist of no processing whatsoever, whereas B might, say, consist of equalisation or dynamic processing, or something more complicated such as feeding the left input to both stereo outputs. Then often we wish to be able to crossfade for results between those of processes A and B. One way of doing this, shown in Fig 6a, is to have two sets of faders, one for A and the other for B, carefully matched and ganged so that their combined gains add up to unity at all settings, and to mix the resulting outputs together. Especially for stereo signal processing, this involves a largish number of ganged faders.

The same result can be achieved rather more simply as shown in Fig 6b, using the fact that $B = A + (B - A)$. Thus, as the fader gain k varies between 0 and 1 in Fig 6b, the output becomes

$$A + k(B - A) = (1 - k)A + kB$$

which smoothly crossfades between process A and process B.

The technique of Fig 6b is widely useful for adjusting the degree of processing smoothly but does depend on processes A and B being generally matched in level polarity and phase so that the crossfading does not produce unwanted cancellations or comb-filter colourations. There are many different applications of

the technique in Fig 6b that are useful in reprocessing and remastering applications, hopefully examples given will evoke other applications.

For example, if a signal requires a complicated equalisation for some of the time but should be flat at other times, then making A a straight-wire connection and B an equaliser (with the usual no inversion and no time-delay requirements satisfied) set to the required complicated equalisation allows a crossfade between flat (fader at zero) and full equalisation (fader at gain 1) without having to adjust 10 or 20 separate equalisation controls at the same time. This is very useful when a problem requiring corrective equalisation (eg hiss or an occasionally dominant badly-equalised instrument) varies in level from moment to moment, since one can continually vary the adjustment with one finger.

Also, the level adjustment need not be manual but can be via a dynamic signal processor. For example, in Fig 7, A is a graphic equaliser, B is a straight-wire connection, and the 'fader' is an expander with unit gain above its threshold. At low signal levels, when the expander acts as an 'off' switch, the output is the input passed through the graphic equaliser A, whereas at high signal levels, the gain k of the expander now equals one, so that the output equals the input passed through B, ie equals the unmodified input signal. Thus the circuit of Fig 7 acts as a completely adjustable dynamic filter for low level noise, tunable to any specific frequency band of noise. As earlier, the sidechain equaliser helps the chosen frequency band to more selectively operate the dynamic processing in order to aid masking.

This circuit works particularly well for dynamically filtering rumble noises, by cutting the graphic equaliser only in the bass and boosting the sidechain equaliser by up to 25dB at the most audible rumble frequencies and cutting it by 25dB at other frequencies.

The subtraction network in Fig 7 may well already exist as the electronically balanced input of the expander. Thus, apart from some adaptors or Y splitter leads, Fig 7 requires nothing not already present in most studios. Manufacturers of expanders should consider incorporating all this circuitry, apart from the equalisers, into their expanders with an extra insert point for an external graphic equaliser A. This way their products could be turned into infinitely adjustable dynamic filters by the addition of such an external equaliser.

Putting the graphic equaliser at B rather than A converts the dynamic filter into one that is flat at low levels and

equalised at high levels above the threshold. Such a dynamic filter might be used, for example, to 'brighten up' high level sound without bringing up low level hiss. Thus, in this configuration, the dynamic filter can be used to add 'excitement' if required, although it could equally be used to strengthen high level bass without bringing up low level rumble. Incorporating a switch allowing the equaliser to be inserted either in the A or the B path allows the same equipment to be used for both low level and high level dynamic filtering. Equalising high level signals also permits, if required, high level distortion to be filtered without affecting the signal the rest of the time.

This kind of corrective dynamic noise filtering requires intelligent setting of the equaliser to minimise the tonal alteration of the wanted signal while taming any faults. The philosophy of 'if it ain't broke, don't fix it' applies here very strongly. For example, to tune out hiss, the naive reaction of most engineers is to filter out the treble above a given frequency but this causes excessive damage to sound quality. It is usually better to cut a fairly narrow band of frequencies at which the hiss is subjectively most audible (often centred around 6 to 10kHz, being higher for cassette than for reel-to-reel tape or FM tuner noise) and to boost neighbouring frequency bands slightly to minimise the overall tonal alteration away from the band that is cut.

With a graphic equaliser, this usually means boosting adjacent band sliders by about one quarter the amount that the rejected band is cut, depending on the equaliser design. With parametric equalisers, one can use a fairly high Q cut and a series low Q (say 20% of the high Q) boost by a small amount (say 0.25 of the cut) centred at the same frequency. The aim is to achieve a frequency response somewhat as in Fig 8a. In such responses, if the cut is deep (say over 12dB), then the measured boost adjacent to the cut band might be around 1.5dB in order to preserve tonal quality, with a smaller boost adjacent to a shallower cut.

This kind of filtering strategy also reduces bass noises well with minimal tonal alteration. Even if it is necessary to cut the most extreme treble or bass frequencies, the general strategy of bringing the response up again beyond the maximally cut frequency band, as in Fig 8b for a hiss filter, still gives a better tonal quality than simply cutting out all the extreme frequencies. When used with the dynamic noise filter of Fig 7, these strategies achieve a reduction of the noise with a minimal audible effect on the wanted signal. Done carefully, the main effect of the dynamic noise filter is

then on the noise and care in adjustment of thresholds, expansion ratios and sidechain equalisation is required to minimise the audible 'pumping' of noise levels.

The arrangement of Fig 7 can also be used for more specialised processing tasks. For example, in a live recording between musical numbers, a performer might speak 'off microphone'. Providing that the recording is on a low-noise medium (*eg* digital), such announcements can be brought up in level automatically by giving the graphic equaliser A a gain greater than one, possibly with a degree of hiss filtering to tame any resulting noise increase. This converts the expander to a device that brings up signals below the threshold level by a predetermined amount

Bandsplit methods

The schematic of Fig 6 can also be applied to processing on bandsplit signals, where the two frequency bands are to be processed differently. If A is a graphic equaliser and B a straight-wire connection, then the outputs of A and of the subtraction network between them form a fully adjustable bandsplitting filter whose outputs sum back to the input signal.

For example, in the cancellation processing described earlier for reducing faults in (say) the right stereo channel, one can apply different degrees of processing in the two bands by putting an adjustable bandsplit in the L-R signal path, as shown in Fig 9. Different degrees of manual or dynamic control of faults can be done in the two frequency ranges, perhaps using a different setting of the L-R panpot controls in the two bands to obtain different compromises between maximum R rejection and width.

There is no need for either of the 'bands' in Fig 9 to be simple low- or high-pass filters. For example, band 1 might well consist of all frequencies except those around the presence region of 2 to 4kHz, whereas band 2 might consist just of the presence region. It is often best if the in-band sliders at the edges of the desired frequency band are slightly boosted (typically by 2 to 3dB) above unity gain to compensate for the in-band effect of the 'cut' sliders just outside the band.

One should also be aware that a non-flat equaliser will not have a flat phase response so the signal going through different processing in the two bands will have varying phase shifts, especially around the crossover frequencies between the two bands. When using bandsplitting to process stereo signals, this can cause some degradation of the resulting stereo effect due to

the 'phasiness' of the results in the crossover region. More advanced 'phase compensated' bandsplitting filters, such as used in the Audio Design *Filmex* processor, or a digital linear-phase bandsplitting filter, can avoid this problem, at the expense of increased cost, reduced flexibility and possibly undesirable phase distortions.

Once one starts using processing with this degree of sophistication, it is obviously essential that the engineer has a very good theoretical understanding of what he/she is doing and a fair degree of practical experience of the precise subjective effects of the processing.

Noise reduction

Another application of the technique of Fig 6b is controlling the side effects of sophisticated digital single-ended processing and noise reduction systems such as *CEDAR* or *NoNoise*. These systems use expansion of the signal in a large number of narrow frequency bands and, like every dynamic filter expansion system, risk a degree of audible side effects. Although the user of such a system will make some of the subjective decisions about the trade-offs to be made at the time of processing, it is worth noting a technique of changing these trade-offs after the digital processing. One arranges that the processed and the unprocessed signals are recorded in precise time synchronism on parallel digital tracks. This can easily be done for mono source material using a stereo mastering medium but stereo source material requires the use of a 4-channel mastering medium.

One can then crossfade between the two signals either dynamically or in a frequency dependent way such as shown in Fig 10. This lets through the unprocessed signal when the graphic equaliser has unity gain with no phase shift but lets it through at frequencies where the equaliser has maximum cut. One starts by equalising the unprocessed signal until the noise is acceptably reduced (ignoring the effect on the wanted signal) and then inserts the same equaliser in the network of Fig 10.

The effect here is that only those frequencies that need the full processing have processing artefacts. Thus the degree of processing and the trade-offs between audible faults and subtle side effects can be re-decided for different release markets, without having to re-do the digital processing each time.

This will work best if fully phase-compensated equalisers are used, and if the processing is used simply to reduce the annoyance value of noise, rather than to eliminate it altogether. I haven't been able to try out this proposal practically but it does suggest strongly that if processing

of a type such as *NoNoise* is to be used for archiving historic mono material, it is wise to put the unprocessed source material in exact time synchronism on the other channel of a stereo mastering medium.

Stereo manipulation

Another aspect of remastering substandard material is the use of linear stereo-to-stereo processing, which manipulates stereoism. The best known process is width control. This can be implemented in a large number of ways, the best known (due to Alan Blumlein in 1931) using sum-and-difference processing as shown in Fig 11. Here the stereo is converted into sum-and-difference (or mid/side) form

$$M = 0.707(L + R)$$

$$S = 0.707(L - R)$$

where the gain 0.707 = -3dB is chosen for convenience, the gain of the difference or 'side' signal S is adjusted by a gain *w* (width) between 0 for mono and about 2.5 or 3 for superwide stereo, being 1 for ordinary stereo, and the modified stereo L_{out} and R_{out} is then derived by a second sum-and-difference operation

$$L_{out} = 0.707(M + wS)$$

$$R_{out} = 0.707(M - wS)$$

Such width controls often help to provide a more satisfactory width for recordings, widening (for $w > 1$) over-narrow recordings or narrowing (for $w < 1$) over-wide ones. They can be used to alter the level-balance between the middle and the edges of a stereo recording. Since the difference or 'side' signal S has no energy from central sounds, increasing its gain, and hence the width, increases the reproduced stereo level of edge sounds relative to central sounds. Thus the relative level of an over-dominant central sound can be reduced by increasing the width.

If, as in Fig 12, a graphic equaliser is inserted into both the sum M and difference S signal paths, this can allow a useful degree of adjustment of both width and of centre-to-edge level-balance at different frequencies, although care is needed to prevent the resulting sound from being excessively 'monoish', when the difference gain is too small, or excessively 'phasey', when the difference gain is too large.

One can reduce noises concentrated in the difference channel, such as rumble from vinyl record playback or wind noise from some stereo microphone systems, or hiss and interference on FM multiplex broadcast reception, by using a dynamic noise filter as described earlier with reference to Fig 7 in the difference channel S, without too much audible effect on the stereoism.

Although it should be used with extreme care, a dynamic filter in the sum channel can sometimes help tame 'common-mode' sum noises in some stereo recordings, it should, however, be realised that excessive reduction of the sum gain can cause stereo that is largely out-of-phase at levels and frequencies at which such a filter is active.

Width control has the defect that, when used with non-coincident stereo microphone techniques or with tapes with significant azimuth errors, it can cause a lot of comb filter colouration. In such a case, it may be expedient to confine the action of the width control to lower frequencies only (say below 700Hz), where this problem is likely to be less.

Besides width control, there are many other useful forms of stereo-to-stereo processing. I have earlier described the use of asymmetry control to adjust left/right level-balance without altering central images. An opposite type of control is rotation control, which alters the centring of a stereo image without any alteration of the level-balance. Such a control was developed and marketed by Telefunken in the 1950s but is now virtually unknown outside German stereo broadcast mixing desks. It is particularly useful for re-centring off-centre live recordings made with stereo microphones without any unacceptable effect on the ambience pick-up and can be used to rescue live recordings that have a 'lop-sided' direct sound pick-up due to performers being at unexpected locations relative to the microphone position and direction.

It would lengthen this article too much to go into the full practical implementation of a rotation control but is described by the mathematical formula

$$L_{out} = L\cos\theta + R\sin\theta$$

$$R_{out} = R\cos\theta - L\sin\theta$$

where $-45^\circ \leq \theta \leq 45^\circ$ describes the degree of 'rotation' of the stereo sound image. For $\theta = 0^\circ$, we have normal stereo, for θ greater than zero, the stereo image is rotated to the left, and for θ less than 0, the stereo image is rotated to the right, all without any alteration of any sound level whatsoever.

In a practical stereo manipulation system, I have found it convenient to incorporate an adjustable bandsplitting of the input stereo as described earlier, so that the stereo-to-stereo processing can be separately adjusted for different frequency components of the sound. This enormously increases the ability to fine tune any corrective processing to different aspects of the input programme content.

Ergonomics and system design

The range of remastering techniques possible with what has been described earlier is enormous, and the techniques described really are very useful for remastering work. Unfortunately, using all these possibilities can be a bit of a nightmare, because equipment has to be constantly repatched and replugged in all kinds of non-standard ways. Although I hope that the above techniques will prove practically useful problem-solvers for remastering engineers, I can't pretend that using them together in different combinations is always easy.

Because of the non-standard wiring involved, one is very liable to get a rat's nest of connections that make no immediately obvious sense, and the controls are liable to be equally non-obvious without detailed crib sheets to remind one of their functions. My own reprocessing system, which I have built up over two and a half years, is still a rat's nest but I have found ways of organising the processing to make it much more ergonomic and easy to use.

Different kinds of processing interact with each other in different ways depending on which comes first. This necessitates a logical system design for organising the systems architecture of how the processing is to be used. Such a system design is in itself a major task. Suffice it to say that conventional studio signal processing architectures are specifically designed for the separate processing of multiple mono signals, with relatively little thought given either to stereo-to-stereo processing or to cancellation or crossfade processing techniques.

No-one has yet designed a really appropriate stereo mastering console architecture to do this kind of reprocessing. I believe that one really does need a new architecture to cope with a wide range of problem recordings. In particular, one needs intuitive methods of setting up alternative signal flow architectures. This system design problem in some ways resembles that for early analogue synthesisers, where systems of patch cords or pin matrices were used to set up different configurations. In principle, digital mastering consoles could be reconfigurable under software control but designing such software is not the easiest of tasks.

The remastering problem is not confined to material in the form of mono or stereo masters, since even multitrack material often incorporates mono or stereo submixes that individually could be usefully reprocessed. A rethink of processing architectures could thus prove

useful even in multitrack mixdown applications.

A good design of mastering console architecture requires a fundamental rethink to ensure an intuitive ease of control of several complex signal-processing chains that interact with one another. Although I have given quite a lot of thought to such systems design over the last 3 years, I see little prospect of any commercial product being developed unless and until a manufacturer or people in the industry perceive a need for a new approach to remastering work. By looking at some of the signal processing involved in remastering, at least this article may have helped to clarify some of the reasons why a new approach might be needed.

The methods described can be improved in the future in two main respects: first, improved control ergonomics so altering the adjustment of complex processing is easier; second, improved subjective performance. The areas of technical improvement include improved dynamic signal processing laws for specific applications and the avoidance of the undesirable subjective side effects of equaliser phase shifts in the above processing algorithms.

In conclusion, good reprocessing involves making sure that the source material is the best available (and it is worth seeking out the earliest-generation copies where possible, no matter how poor the basic recording quality) and is played back as carefully as possible, with attention paid to things like azimuth, speed, equalisation and noise reduction tracking – even if the source is a poor quality cassette recorded on a maladjusted domestic machine. No amount of clever reprocessing technology can properly compensate for carelessness in the original transcription of the source.

Ultimately, the results of any reprocessing involve the artistic judgement and 'golden ears' of the reprocessing engineer. All reprocessing ultimately involves a careful tailoring of the reprocessing to the original faults – and no magic 'automatic clean up' processing is ever likely to be developed – at least until machines develop intelligence and artistic judgement.

Reference

1) Michael Gerzon, 'Stereo Shuffling: New Approach – Old Technique', *Studio Sound*, July 1986, pages 122-130.

The Gerzon Archive
www.audiosignal.co.uk

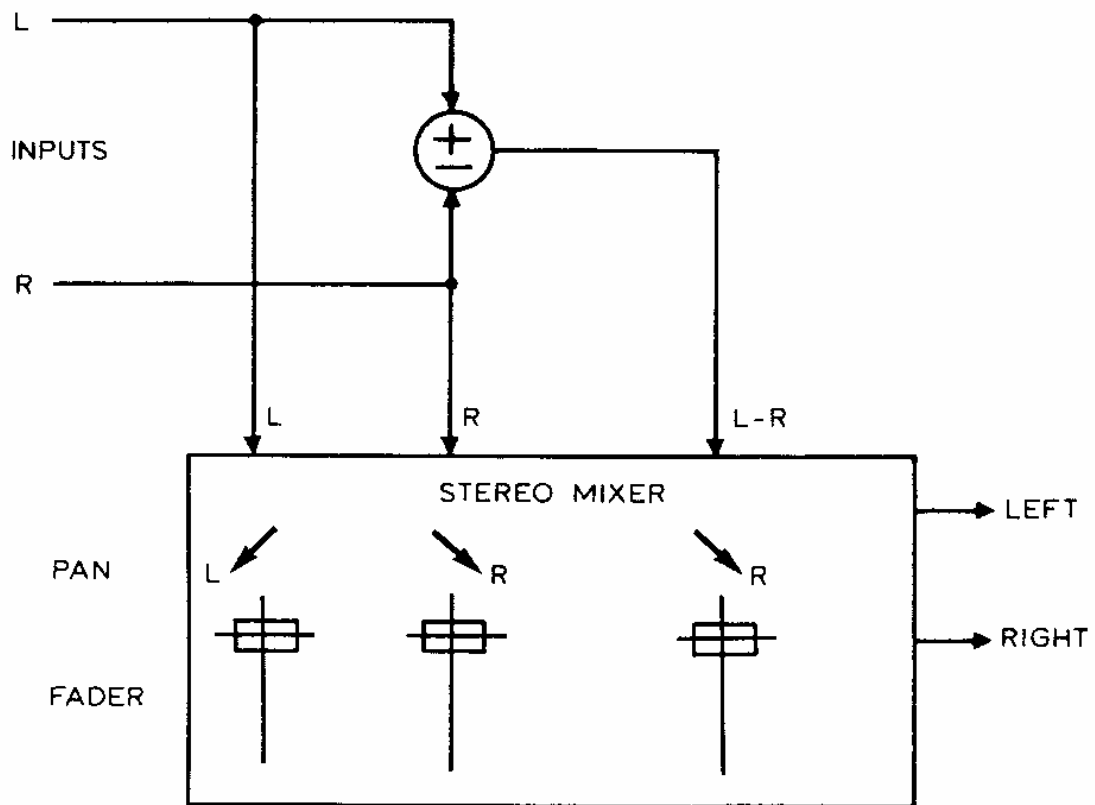


Fig 1: Processing a faulty right channel

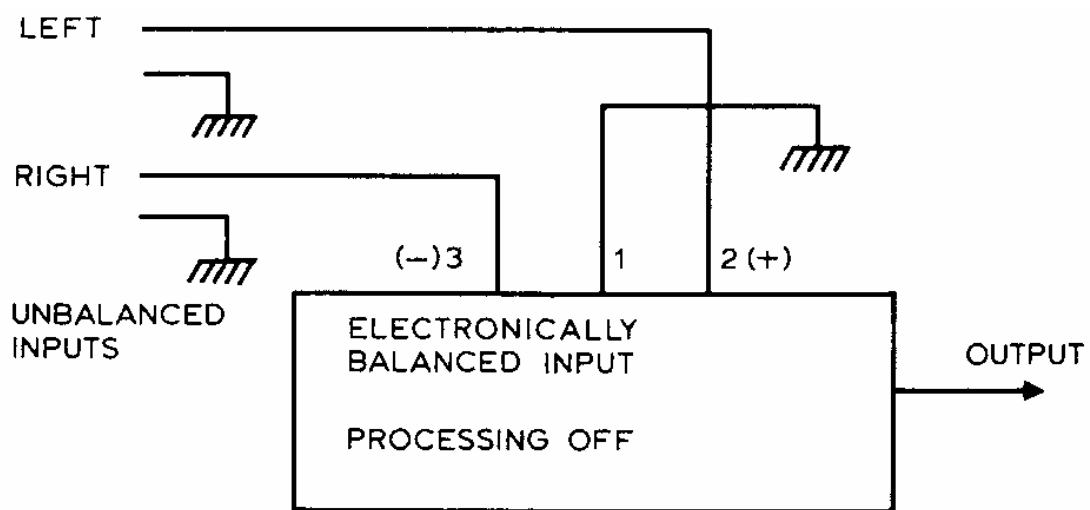


Fig 2: Use of equipment with an electronically balanced input as a subtraction network

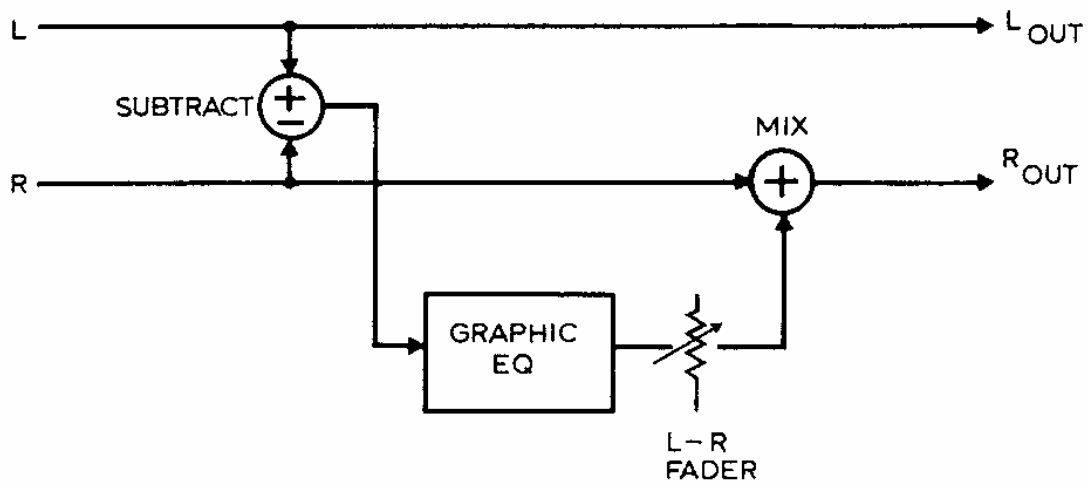


Fig 3: Inserting a graphic equaliser into the L-R path of Fig 1

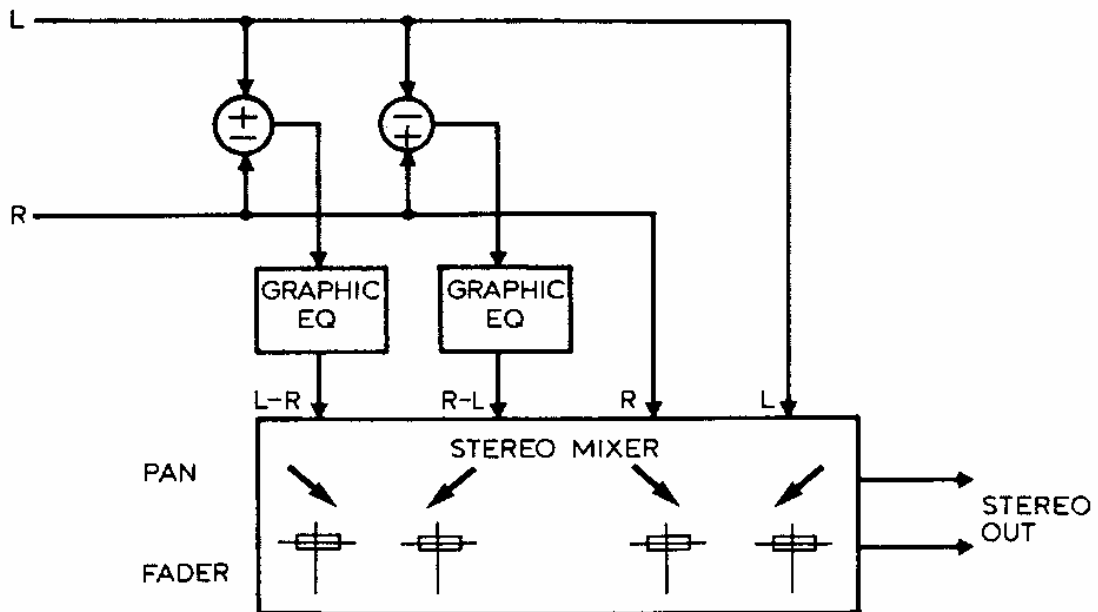


Fig 4: Processing for cancelling faults in the left or right channels




Fig 5: Removing R channel faults with dynamic processing of the L-R signal




Fig 6: Two methods of crossfading between two different forms A and B of signal processing

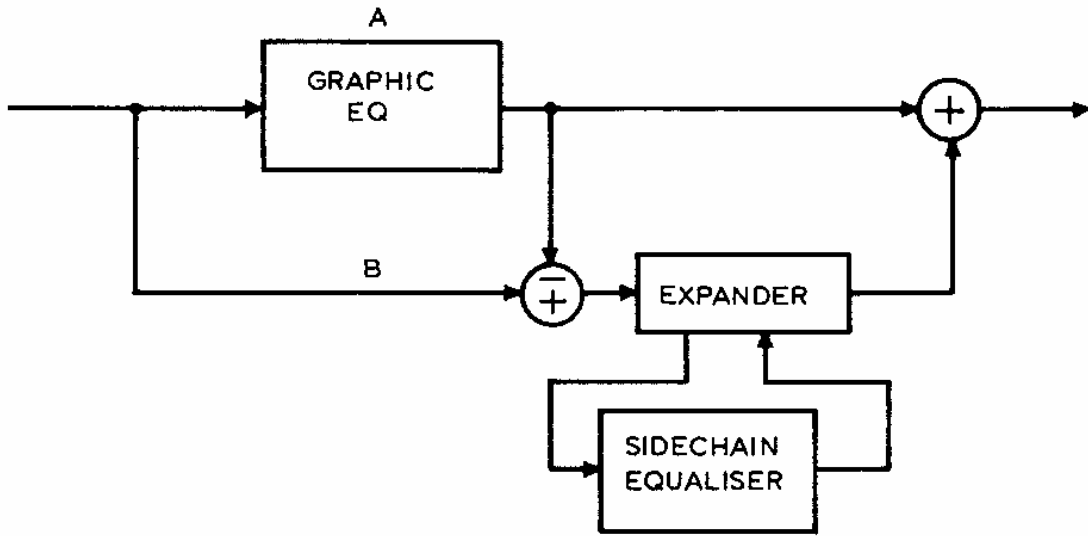


Fig 7: Fully adjustable dynamic filter based on Fig 6b




Fig 8: Typical preferred responses of noise filters

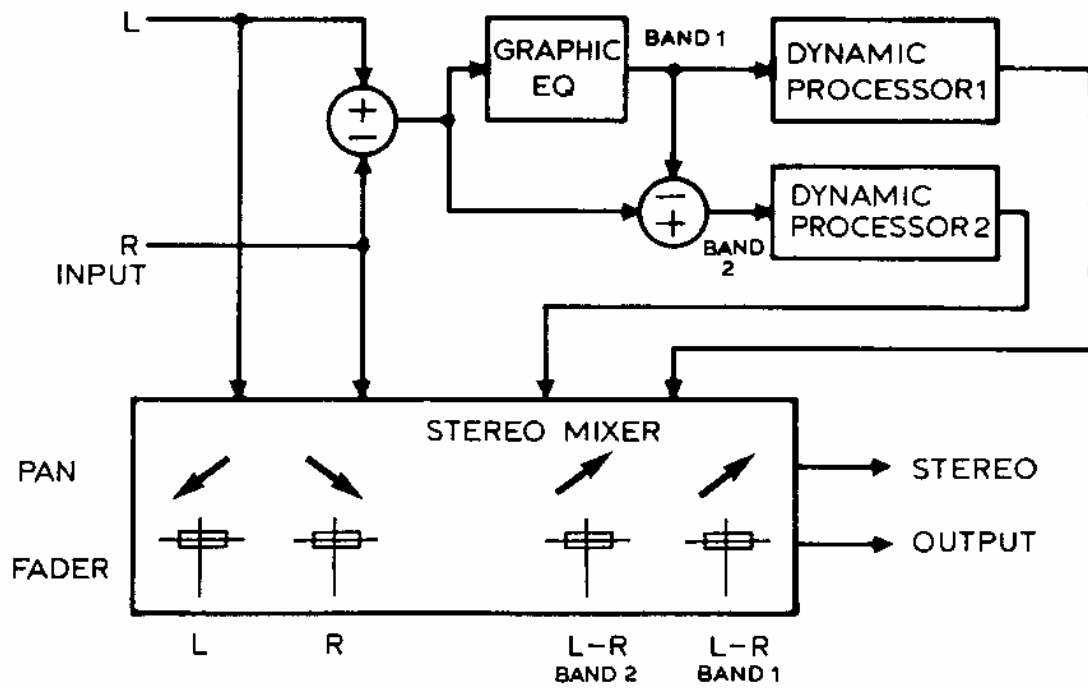


Fig 9: Example of use of bandsplitting for control of right channel faults

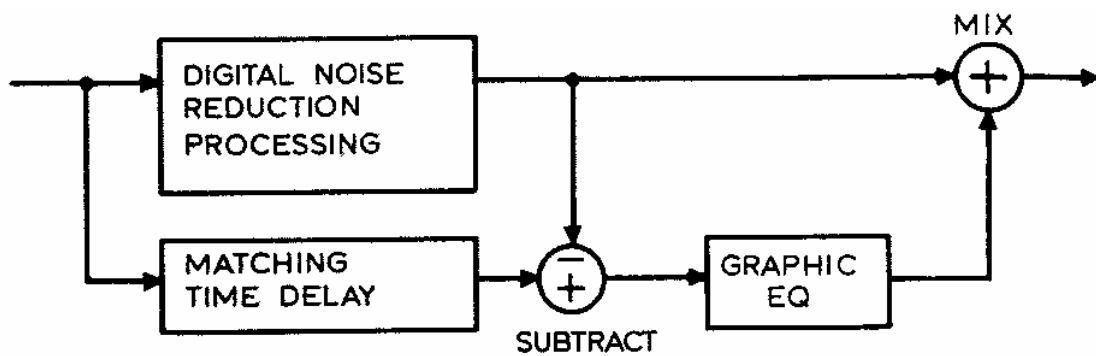


Fig 10: Control of digital noise reduction side effects




Fig 11: Schematic of width control




Fig 12: Frequency-dependent width control and edge-to-centre rebalancing system using graphic equalisers in the sum and difference channels