HI-FI DYNAMIC RANGE COMPRESSOR

Several designs for dynamic range compressors have previously been published in Elektor, but this is the first design that provides truly high-fidelity performance. It can accept nominal signal levels from 600 μV to 2.2V and is thus suitable for use with microphones as well as higher output circuits such as audio preamps. The harmonic distortion is below 0.3% and the decay of the compressor can be adjusted to suit different types of programme material.

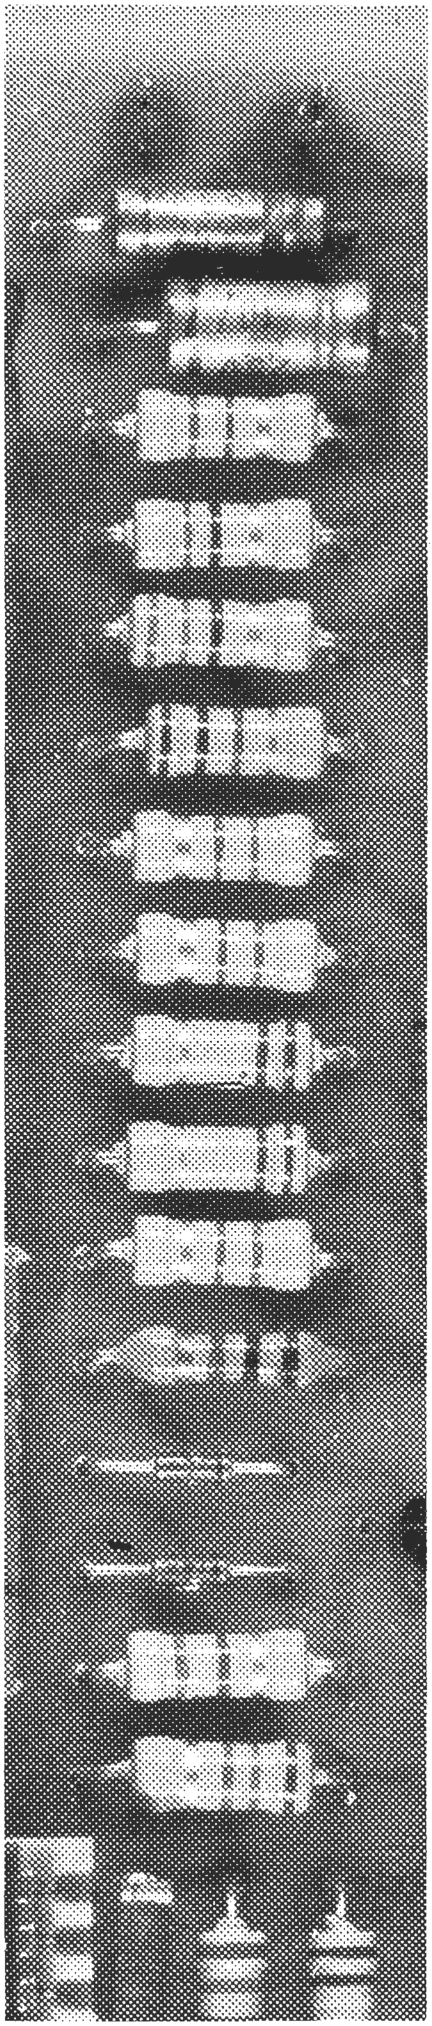

Dynamic range compressors find many applications. The recording industry would be unable to make a single disc without some method of compressing the dynamic range of the programme. The reasons for this are fairly obvious. The dynamic range of live music, from the softest ppp of the piccolo to the loudest fff of the bass drum can be in excess of 80 dB. However, even the best recording media have a usable dynamic range of only 60 dB or so. The smallest signal level that can be recorded on tape is limited by tape noise, while the largest signal level is limited by tape saturation. In the case of disc the lowest signal level is limited by surface noise and the highest signal level is limited by the tracking ability of the cartridge.

At the other end of the recording chain compressors are very useful in discotheques to compress the dynamic range of the recording still further and so maintain a fairly high average signal that can be heard above the background noise (of people conversing in shouts) without the danger of the equipment overloading on peaks. Compressors can also be used in home-tape recording, public-address systems and by radio amateurs, to name but a few applications.

Principles of compressors

Although most compressors operate according to the same broad principles the compression characteristics may vary widely depending upon the intended application. For example, the 'limiters' found in some tape recorders do very little at all to the signal until the signal level approaches tape saturation, but they then operate in a very heavy-handed manner to ensure that the signal does not exceed tape saturation level. Compressors intended for P.A. systems, on the other hand, begin to operate at a very low signal level to try to maintain a fairly constant signal level for maximum intelligibility. The only real difference between these two types is the threshold level at which they begin to operate.

Compressors with pretensions to high-fidelity, on the other hand, apply compression over the entire dynamic range of the signal, rather than simply clamping all signals which exceed a certain level.

However, these systems all have several things in common. Figure 1 shows a block diagram of a generalised compression system. The input signal passes through a voltage (or current) controlled attenuator followed by an am-

1

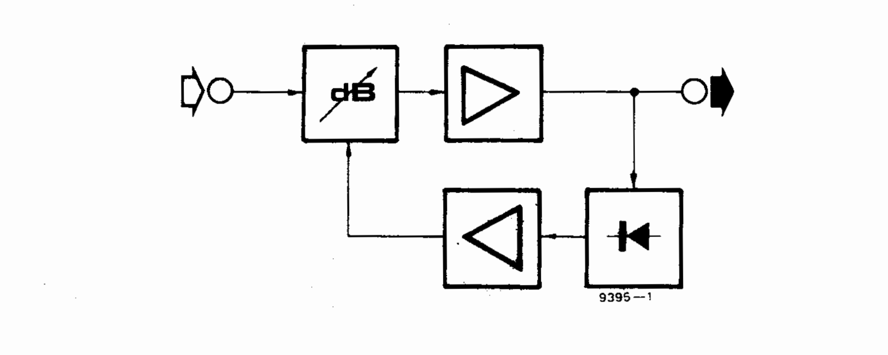

2a

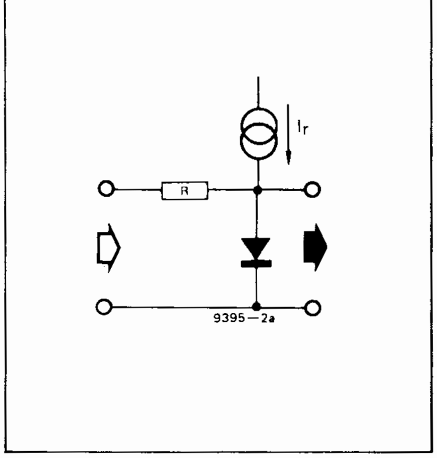

2b

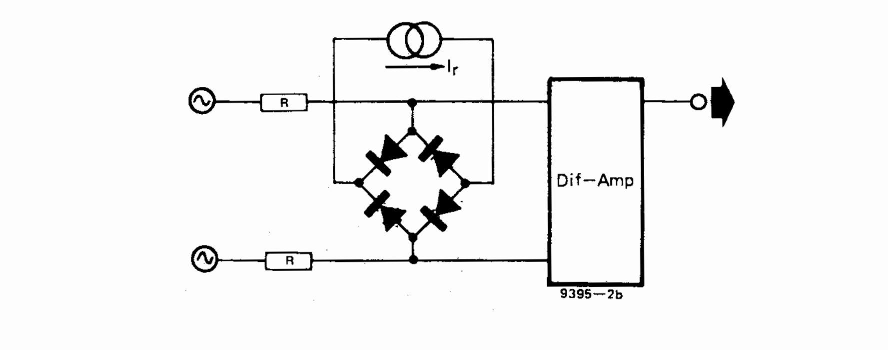

3a

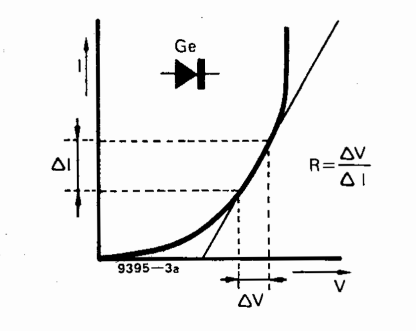

3b

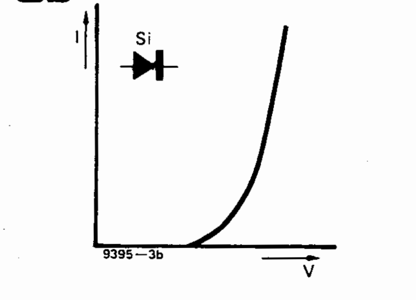

Figure 1. Block diagram of a dynamic range compressor.

Figure 2a. Basic current-controlled diode attenuator.

Figure 2b. Bridge type diode attenuator with differential signal input and output to reject the common-mode control voltage.

Figure 3a. Forward transfer characteristic of a germanium diode.

Figure 3b. Forward transfer characteristic of a silicon diode.

plifier. The amplified output is rectified, DC amplified and used to control the attenuator. Thus, as the input signal increases so will the control voltage applied to the attenuator and hence the degree of attenuation, so that, instead of the output increasing in a linear fashion proportional to the input, the increase in the output signal becomes smaller as the input signal increases. Of course, the control voltage applied to the attenuator must not be the 'raw' output of the rectifier, since the attenuation would then vary during each cycle of the signal waveform, leading to gross distortion. The output of the rectifier must be 'smoothed' to give a control voltage that will follow the envelope of the waveform, but not each individual cycle.

This inevitably leads to compromises in the design. If the control voltage has a long time constant then the compressor will be unable to respond to sudden changes in signal level, while if it is too short the signal will be distorted, especially at low frequencies. Fortunately an interesting psychoacoustic effect can be used to advantage. The ear is insensitive to distortion during sudden increases in signal level, so the attack time constant of the control voltage can be made fairly short. This can be important in preventing equipment overloads. The decay time constant however, must be relatively long to prevent the control voltage following the signal waveform, and this of course means that, should the signal level suddenly decrease, it will be some time before the attenuation is reduced, causing intervals when the signal may be much too quiet. In the present design this problem is reduced by providing three different decay time constants; a short (120 ms) time constant suitable for speech, where signal levels fluctuate rapidly and distortion is less important than intelligibility; a medium time constant (600 ms) suitable for mixed (speech and music) programmes, and a long (3 sec) time constant for music, where distortion is more important than the occasional quiet passage.

Controlled Attenuator

Many non-linear devices could be used as the control elements in the attenuator, for example field effect transistors, voltage dependent resistors or light dependent resistors, but one of the cheapest and most effective solutions is an attenuator using silicon or germanium diodes.

The forward conduction characteristics of a germanium and a silicon diode are shown in figures 3a and 3b respectively.

It will be seen that the dynamic resistance of a diode $\frac{\Delta V}{\Delta I}$, decreases as the current through the diode increases.

This can be put to use in an attenuator as shown in figure 2a. The diode forms the lower limb of a potential divider and is fed by a current source I_r. A current source is used since it has an infinite output impedance and cannot, of itself, attenuate the signal. If the control current through the diode is increased the diode resistance will fall and the attenuation of the signal will increase.

This simple circuit has several disadvantages. Firstly, the control current, as well as varying the dynamic resistance of the diode, also causes a voltage drop which is superimposed on the signal. Changes in this voltage as the control current varies can give rise to clicks and thumps. This problem can be overcome by arranging four diodes in a bridge configuration as shown in figure 2b. The signal is applied differentially and is amplified by the differential amplifier at the output of the attenuator, but the voltage produced by the control current appears in common mode at each of the differential amplifier inputs, and is thus rejected.

The second problem with the diode attenuator is that the signal voltage causes a current to flow through the diode, which varies the dynamic resist-

4

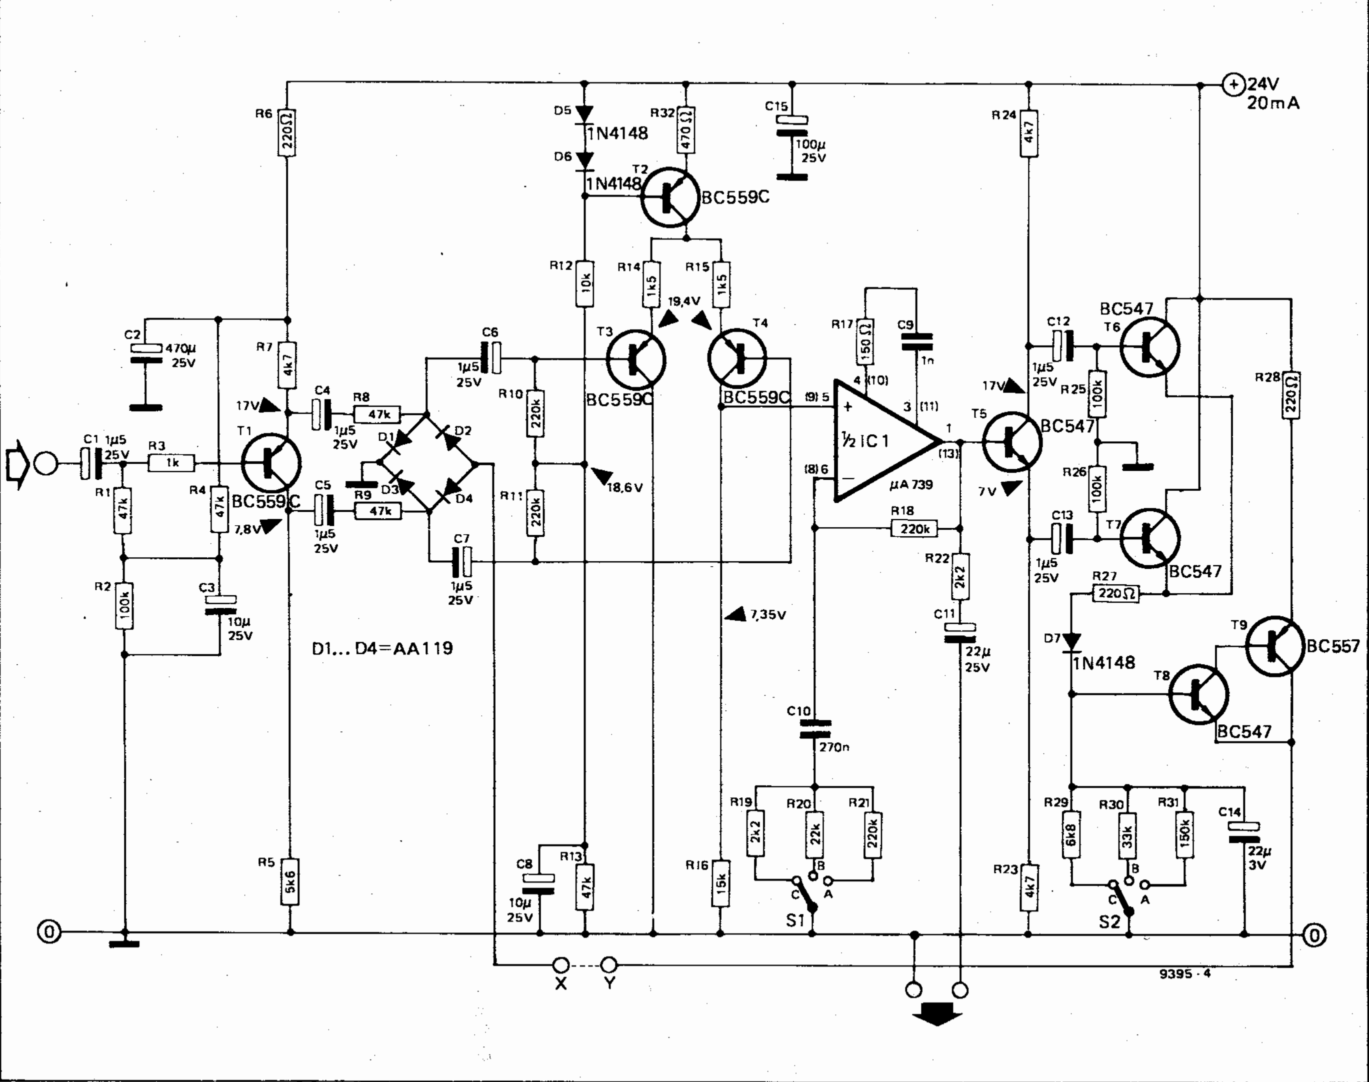

ance and hence the attenuation. This can lead to distortion. The solution is to make the signal voltage across the diode small compared to the total voltage drop across the diode, but here a compromise must be struck between distortion and signal-to-noise ratio, which is determined by the noise generated by the diodes. On the one hand, the signal voltage cannot be reduced without degrading the signal-to-noise ratio, while on the other hand, increasing the control current increases the diode noise, which also degrades the signal-to-noise ratio.

As a final note on the controlled attenuator, germanium diodes are to be preferred to silicon in this application. The only usable portion of the forward conduction curve is that part where it actually is curved. The initial portion of the curve where the diode does not conduct cannot be used, nor can the later portion where the curve becomes a straight line, since the dynamic resistance is then constant. It can be seen that a much larger portion of the curve can be used in the case of a germanium diode, which makes setting up of the attenuator much easier and also gives it a more favourable characteristic.

Practical circuit

The main parts of the circuit shown in the block diagram and discussed earlier can be seen fairly easily in the circuit of figure 4. Transistor T1 functions as a phase splitter and the antiphase signals from its emitter and collector are fed to the controlled attenuator comprising R8, R9 and the diode bridge D1 to D4. To ensure symmetrical operation of the attenuator and good control signal rejection diodes D1/D3 and D2/D4 should be matched pairs.

The output of the attenuator is taken from the cathodes of D2 and D4 and is fed to the differential amplifier consisting of T3 and T4. The control current from point X is fed to the anodes of D2 and D4 and since the voltages caused by it appear in equal amplitude and phase at the attenuator outputs they are not amplified by T3 and T4.

The signal at the collector of T4 is then fed into an amplifier (½ IC1) which has three switched gains of 2, 11 and 102 to suit different input signal levels. The compressed output is taken from the output of this amplifier via R22 and C11.

The rectifier that produces the control signal operates in a somewhat unusual manner. The output of IC1 is fed into yet another phase splitter T5. The antiphase outputs from the emitter and collector are fed into two emitter followers T6 and T7. These are operated with zero base bias and so only conduct on the positive half-cycle of the waveform fed to them, i.e. they both operate as half wave rectifiers, but since they are fed with antiphase signals a full-wave rectified version of the signal appears at the junction of their emitters. The output from T6 and T7 is used to charge C14 via R27 and D7. The attack time constant of the compressor is thus approximately R27xC14. Three decay time constants can be provided by switching in R29, R30 or R31. The voltage on C14 is used to control a current source T8/T9, which feeds a control current into the diode bridge via points X and Y. The connection between these points may be broken to open the feedback loop for test purposes. If continuous adjustment of gain and decay time is required it is quite permissible to replace S1, S2 and their associated resistors by potentiometers. To replace R19 to R21 a 220k potentiometer should be used in series with a 2k2 resistor to limit its control range. In place of R29 to R31 a 220k potentiometer should again be used, but a 390k resistor should be connected in parallel with it to limit the maximum resistance to 150k. A 6k8 resistor should be connected in series with this combination to limit the control range and

Figure 4. The complete circuit of one channel of the dynamic range compressor.

Figure 5. Transfer characterisitc of the compressor for the three different positions of S1.

also to avoid overloading the rectifier output.

Test Results

Figures 5a to 5c illustrate the transfer characteristic of the compressor with S1 in the low, medium and high gain positions respectively. The signal levels are plotted in dB relative to 0dB = 1 volt. It can be seen that the position of S2 has a very slight effect on the transfer characteristic. This is caused by the different loading of the rectifier.

Figures 6a to 6c show the frequency response of the compressor with S1 in its three different positions. Here the gain in dB at the input levels specified is plotted against frequency on a logarithmic scale. It will be noted that the gain with S1 in position C is rolled off below about 200 Hz by the increasing impedance of C10. This setting of S1 is intended principally for microphone inputs for speech use. Rolling off the gain at low frequencies does not impair intelligibility, but it does help to minimise hum pickup, flicker noise from T1, and thumps and rumbles caused by handling the microphone.

The final tests performed on the compressor were the measurement of decay with S2 in its three different positions. This is shown in figures 7a to 7c. The test method used was to feed a low level signal into the compressor, with bursts of a much higher amplitude superimposed upon it. In each case the input signal is the lower trace of the oscillogram.

In figure 7a it can be seen that the output signal level rises quickly at the start of the tone burst. This is due to the delay in the operation of the compressor caused by the attack time constant. As C14 charges, however, the amplitude is quickly controlled. At the end of the tone burst the amplitude of the output signal falls below what it was before the tone burst, then slowly recovers. In fact, comparing figure 7a with figures 7b and c it can be seen that it does not regain its original level before the next tone burst arrives. In figures 7b and 7c it can be seen that the signal recovers its original amplitude much more rapidly due to the shorter decay time constant. Figure 7c shows the whole sequence particularly well; the initial amplitude of the output before the tone burst, then the overshoot at the start of the toneburst, quickly controlled; the drop in amplitude at the end of the tone burst and finally recovery to the original output level.

5a

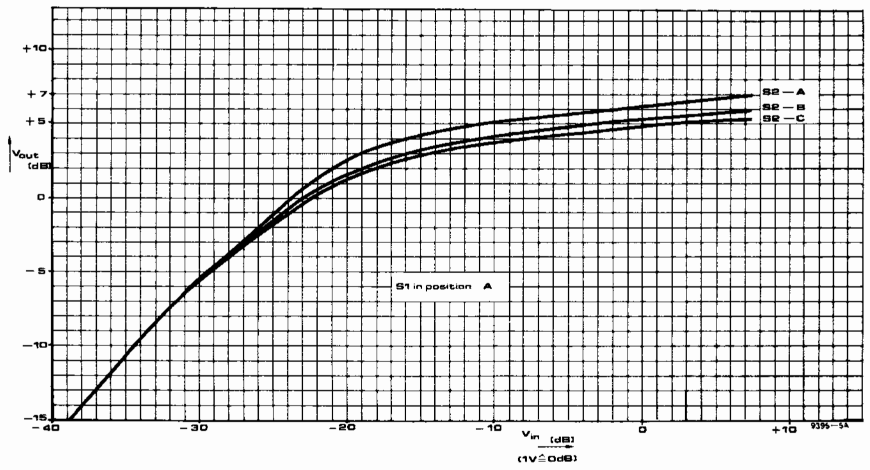

5b

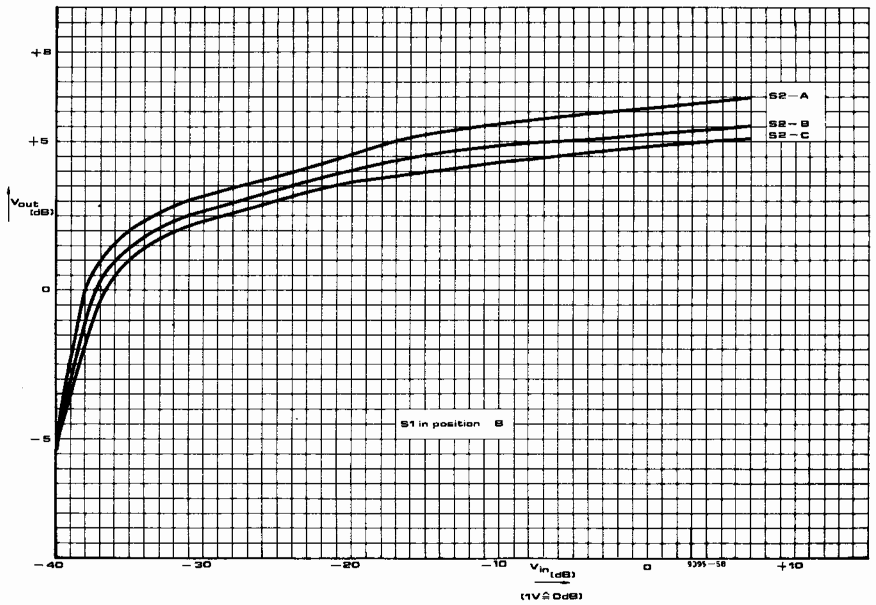

5c

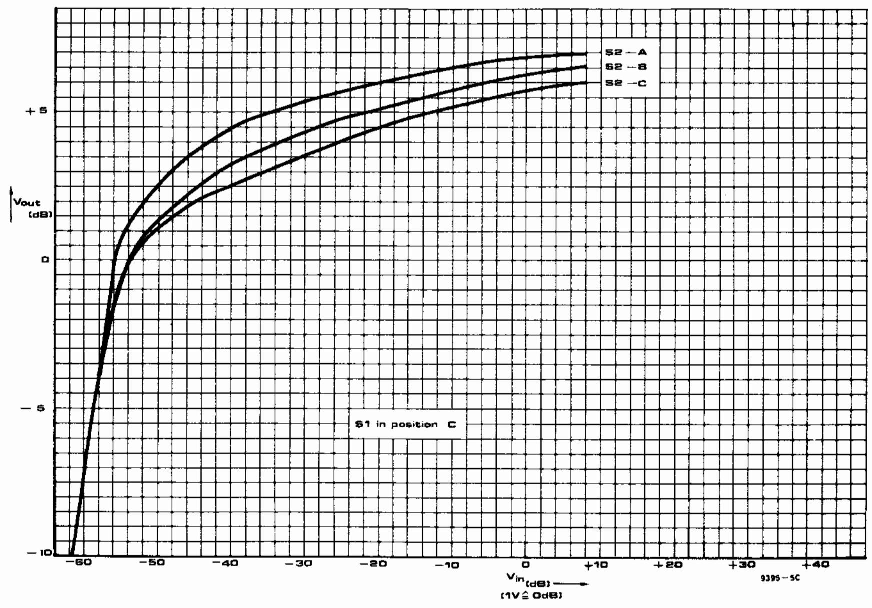

Distortion figures

Because of the many variables involved, such as decay time constant, input signal level and frequency, it is difficult to give comprehensive results for distortion. Obviously, higher distortion figures are to be expected at higher input levels due to interaction of the signal with the control current through the attenuator. Distortion might also be expected to be worse at lower frequencies and/or faster decay time constants.

However, to give a typical example, with S1 in position B and with the nominal input level of 50mV the harmonic distrotion was less than 0.3% over most of the audio spectrum. To give a worst case example, with S1 in position C (maximum gain) and an input level of 500 mV (40 dB greater than the nominal input level), the distortion was still less than 1%. This compares very favourably with commercially available compressors, which often introduce up to 10% distortion.

Construction

Since the compressor is intended for high-fidelity applications (which almost invariably means stereo) a two-channel board layout was designed. This also facilitates some interesting applications, which will be discussed later. The printed circuit board and component layout are given in figures 8 and 9. If the unit is to be used for straightforward stereo recording or reproduction then S1 and S2 should be double-pole three-way types. Otherwise refer to the section headed applications. Apart from this construction of the printed circuit board is extremely simple. Wiring of the unit into an audio system should conform to normal audio practice, with screened leads being used for inputs, outputs and connections to switches. Care should be taken to avoid earth loops, and the inputs of the compressor should be kept well away from mains transformers and/or the outputs of power amplifiers.

Applications

The provision of a link on the board between points X and Y, as well as being a break point in the feedback loop for test purposes, also makes possible some interesting applications. One example is a voice operated fader, shown in figure 10.

One channel of the compressor is fed from a microphone, the other from a music source such as disc or tape. The control output of the microphone channel is fed to the controlled attenuator input of the music channel, so that as soon as someone speaks into the microphone the level of the music is lowered – very useful for disc jockeys. For this application separate switches are required for S1 in each channel, since S1 will be set to maximum gain for the mic channel, but will be in either the medium or low gain position in the

6a

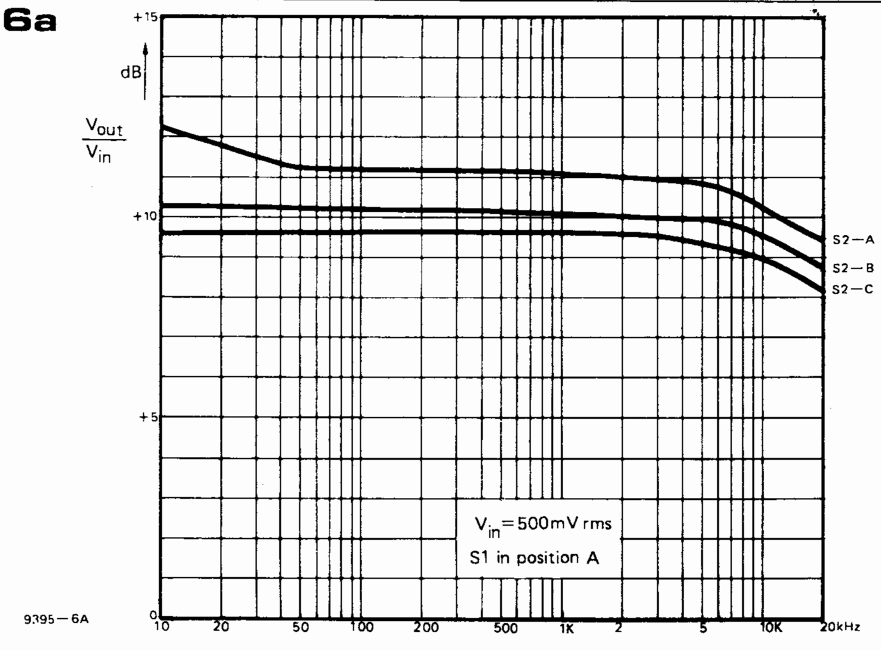

6b

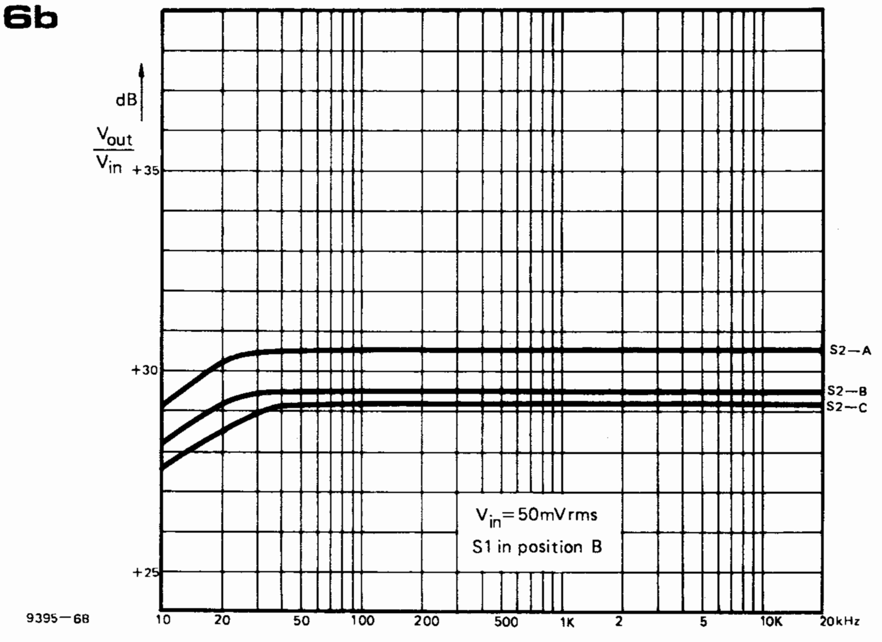

6c

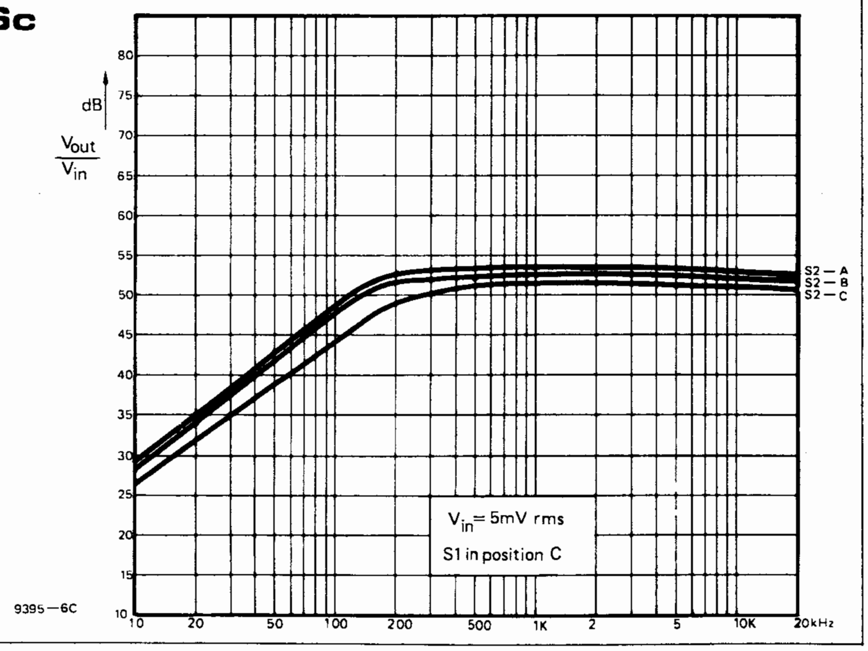

Parts list (For one channel)

Resistors:
R1, R4, R8, R9, R13 = 47 k
R2, R25, R26 = 100 k
R3 = 1 k
R5 = 5k6
R6, R27, R28 = 220 Ω
R7, R23, R24 = 4k7
R10, R11, R18, R21 = 220 k
R12 = 10 k
R14, R15 = 1k5
R16 = 15 k
R17 = 150 Ω
R19, R22 = 2k2
R20 = 22 k
R29 = 6k8
R30 = 33 k
R31 = 150 k
R32 = 470 Ω

Capacitors:
C1, C4, C5, C6, C7,
C12, C13 = 1µ5/25 V
C2 = 470 µ/25 V
C3, C8 = 10 µ/25 V
C9 = 1 n
C10 = 270 n
C11 = 22 µ/25 V
C14 = 22 µ/3 V
C15 = 100 µ/25 V

Semiconductors:
T1 . . . T4 = BC 559 C
T5 . . . T8 = BC 547 B, BC 107 B
T9 = BC 557 B, BC 177 B
D1/D3, D2/D4 = AA 119 matched pairs
D5, D6, D7 = 1N4148 (1N914)
IC1 = µA 739, SN 76131,
TBA 231, MC 1303

Miscellaneous:
S1, S2 Mono: single pole three way.
Stereo: double pole three way.

7a

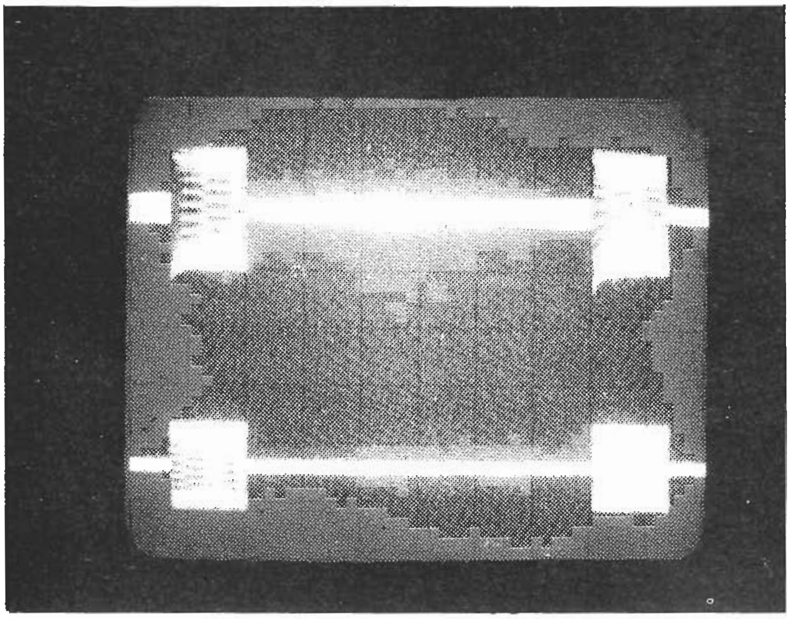

7b

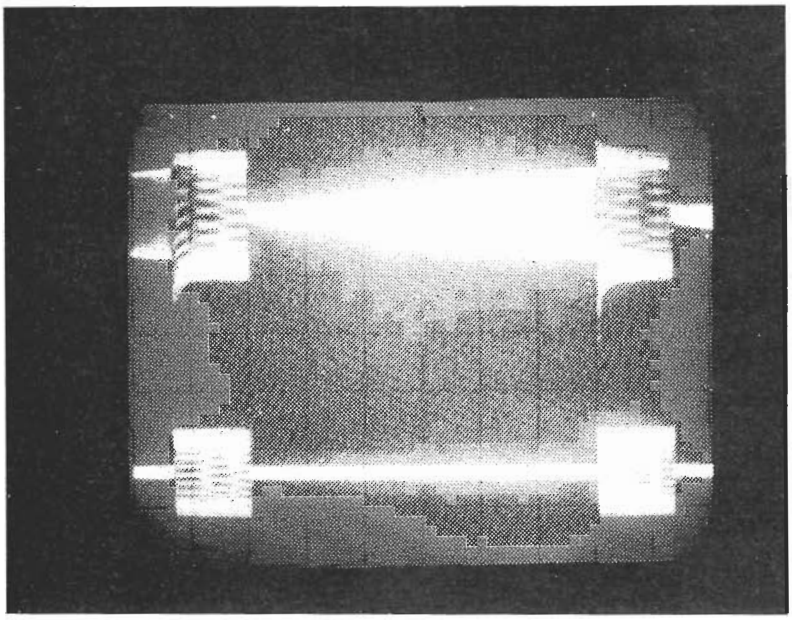

7c

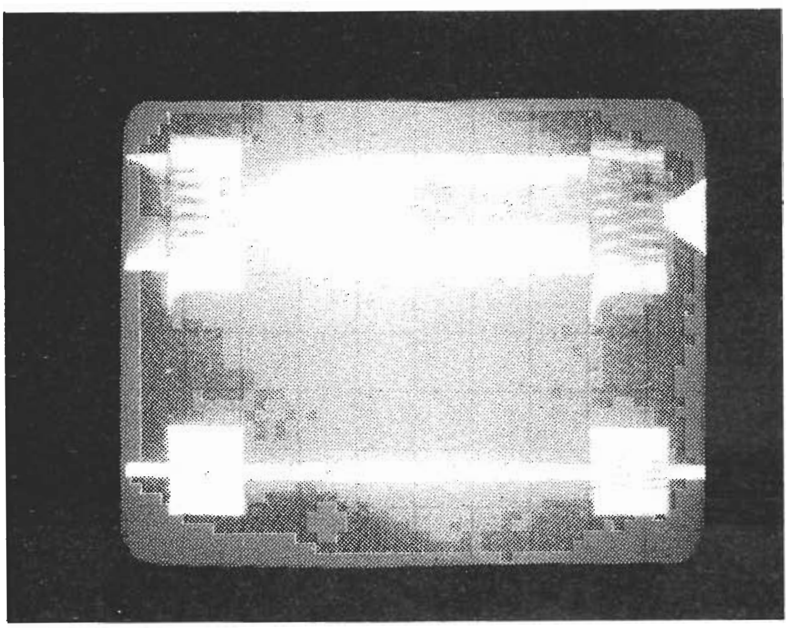

Figure 6. Gain/Frequency plot of the compressor at specified input levels for the three positions of S1.

Figure 7. Oscillograms showing the response to toneburst signals with different settings of S2.

8

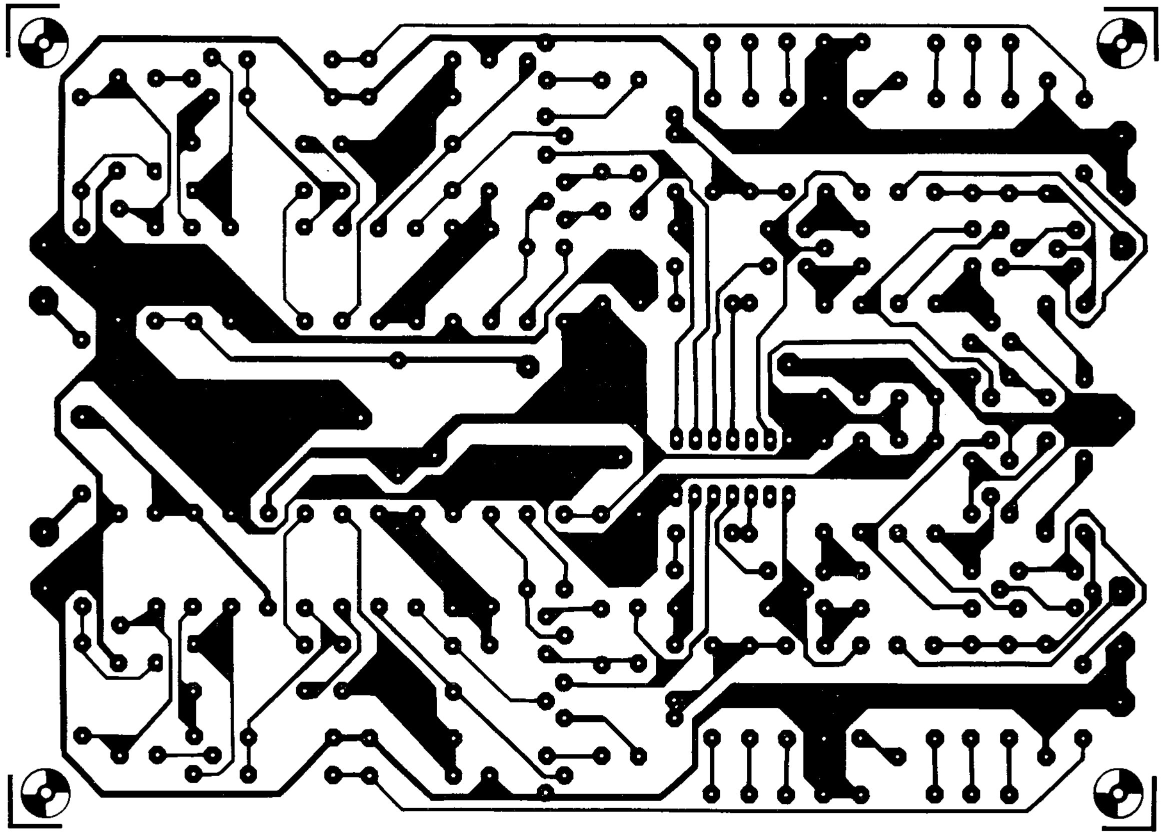

9

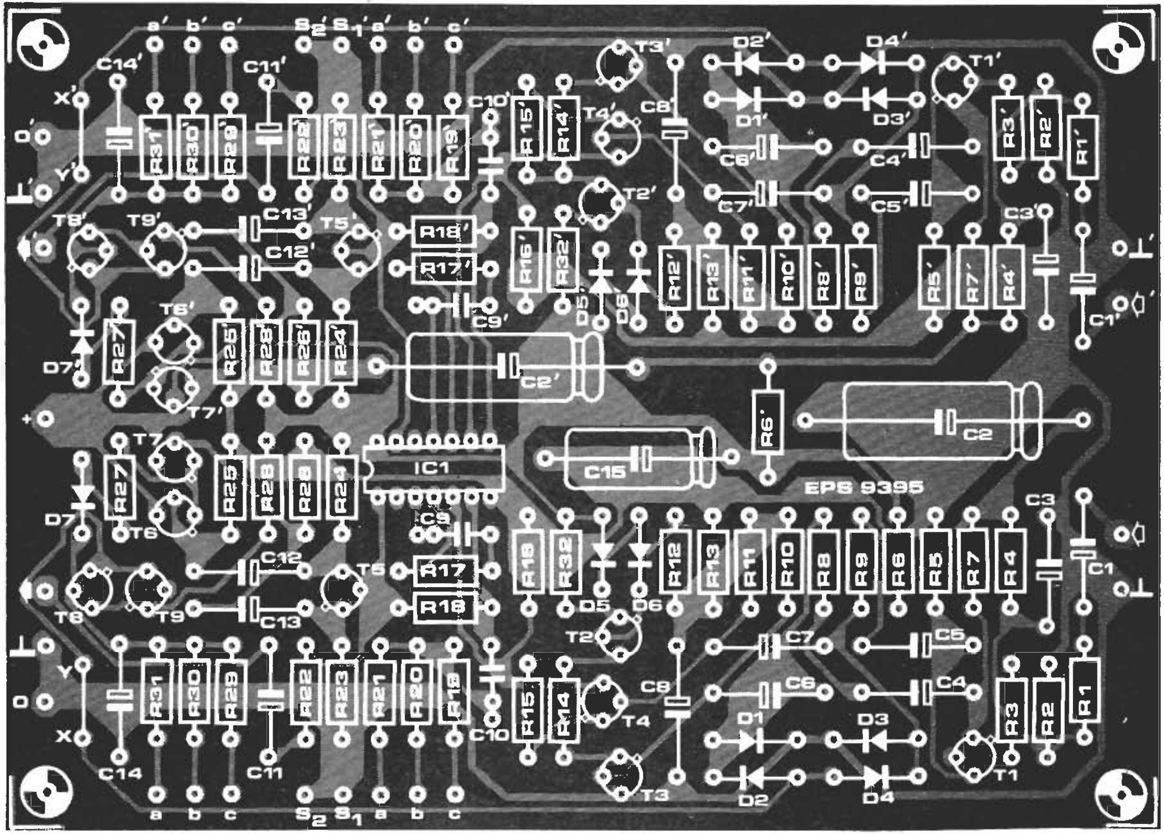
a' b' c' S2' S1' a' b' c'
C14' C11' C10' T3' D2' D4' T1'
X' 0' Y'
R31' R30' R29' R22' R23' R21' R20' R19' R15' R14' T4' C8' D1' D3' C6' C4' R3' R2' R1'
T8' T9' C13' T5' R18' T2' C7' C5' C3'
C12' R17' R16' R32' D5' D6' R12' R13' R11' R10' R8' R9' R5' R7' R4'
C9' C1'
D7' R27' T6' R25' R28' R26' R24' C2'
T7' R6' C2
T7 IC1 C15
EPS 9395
R27 R25 R28 R26 R24 C3
D7 T6 C9 R16 R32 R12 R13 R11 R10 R8 R9 R6 R5 R7 R4 C1
C12 T5 R17
T8 T9 C13 R18 D5 D6
Y T2 C7 C5
R31 R30 R29 R22 R23 R21 R20 R19 R15 R14 C8 C6 C4 R3 R2 R1
0 T4
X C10 D1 D3
C14 C11 T3 D2 D4 T1
a b c S2 S1 a b c

Figure 8 and 9. Printed circuit board and component layout for a stereo version of the compressor. (EPS 9395)

Figure 10. By linking the control output of one channel to the attenuator input of the other channel the compressor can be used as a voice operated fader.

Figure 11. For stereo operation the control circuits of both channels should be linked to avoid shifts of the stereo image.

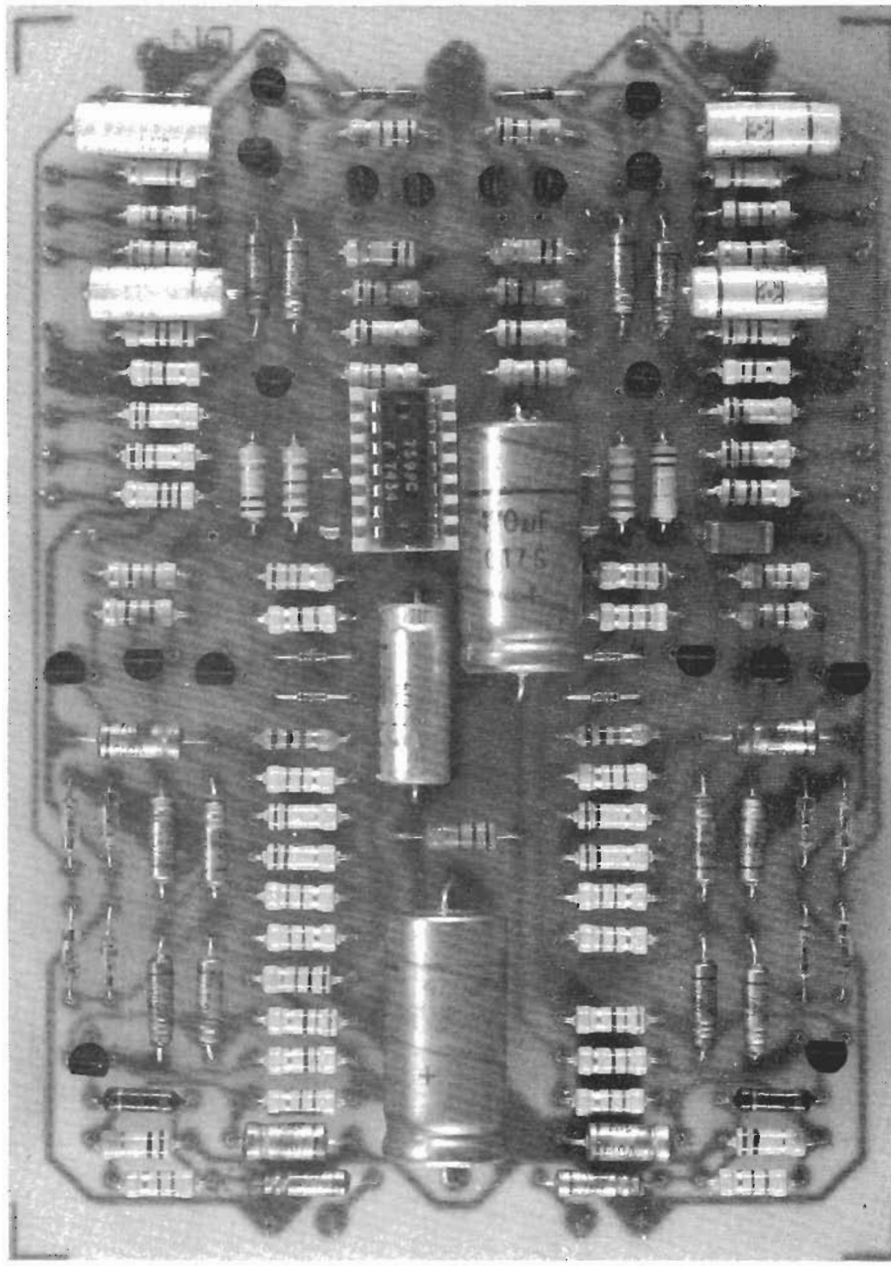

other channel. S2 will be required in the microphone channel, but will not be required in the music channel.

Many variations on this theme can be constructed, limited only by the ingenuity of the reader. For example, the system could be extended so that the microphone could fade several channels. The link between points X and Y in the microphone channel could be omitted so that this channel functioned simply as a microphone preamp, while the link between points X and Y in the music channel could be included, so that the compressor would be controlled either by the microphone or by the music signal.

Stereo Use

For stereo use it is important that points X and Y in both channels should be linked as shown in figure 11 so that the compressors operate 'in unison'. If the two channels are not linked in this manner some disturbing shifts of the stereo image may occur. For example, consider an orchestra arranged around a stage with a soloist centre.

While the left and right-hand sections of the orchestra were playing at about the same level no problem would occur. The solo would emanate equally from both loudspeakers and would thus appear central.

If, however, a crescendo occured in the left-hand section of the orchestra then the left-hand compressor would operate, compressing all the left channel signal including the soloist. The image of the soloist would thus shift to the right giving the impression that he or she was roller skating around the stage. With the control circuits linked a change in signal level in one channel will vary the attenuation of both channels, and the stereo image will be unaffected.

10

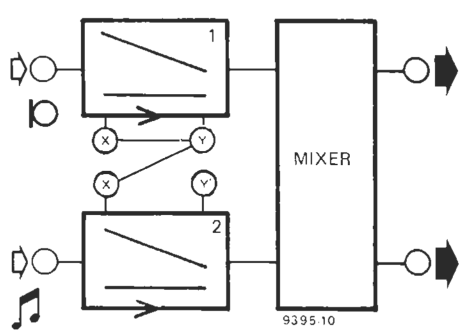

11