The Spectral Recording Process*

RAY DOLBY

Dolby Laboratories Inc., San Francisco and London

A complementary audio signal encoding and decoding format, called spectral recording (SR), for use in professional magnetic recording and similar applications is described. The processing algorithm is highly responsive to the spectral properties of the signal. A further characteristic used during encoding deemphasizes high-level signal components in the frequency regions usually subject to channel overload. The process results in a significant reduction of audible noise and distortion arising in the channel.

0 INTRODUCTION

In 1980, some 14 years after the introduction of A-type noise reduction [1], the author began work on the development of the next generation system for general-purpose professional recording and transmission. A configuration that would employ the A-type characteristics as part of the new system, with switchable compatibility, was considered initially. However, this would not take full advantage of the new technology embodied in the C-type system [2], nor would it readily allow the incorporation of some further new concepts. Therefore, the particular parameters of the A-type system were abandoned as a starting point for the new development. However, the basic principles, which appear to be as valid as when they were first introduced, were retained: the use of a main signal path without any dynamic processing to pass high-level signals, coupled with a low-level side chain compressor to provide dynamic action.

The design goals of the new system were set high. The new technology, called spectral recording (SR), should provide master recordings of the very highest quality, especially with regard to audible signal purity. Yet the system should be practical and economical for routine applications, being suitable for easy and troublefree use in a wide variety of professional recording and transmission environments. Certain new techniques, to be described, provide the required signal quality and practicality but result in circuit complexity. Reliance has been placed on improved circuit implementation and manufacturing techniques to overcome the problems of complexity and to ensure economical production of the new system.

1 BRIEF OUTLINE

The goal of the spectral recording process is to modify the various components of the incoming signal in such a way as to load an imperfect recording or transmission medium in the most rational way. Generally, high-level signal components at both ends of the spectrum are attenuated, whereby a better match with the overload characteristic of the medium is provided. At the same time, low-level components of the signal are amplified substantially, in a highly frequency-selective way. These effects are reversed during reproduction, restoring the original signal. The result is a significant reduction of distortion and noise, in both the absence and the presence of signals.

The process has a number of layout and operating characteristics in common with the A-type ([1], B-type [3], [4], and C-type [2], noise reduction systems. The SR process takes these developments considerably further in the same general direction. An understanding of the new system requires reference to the technical papers on these earlier systems, the C-type paper [2] being particularly relevant.

Referring to Fig. 1, which will later be described in detail, a main signal path is primarily responsible for

* Presented at the 81st Convention of the Audio Engineering Society, Los Angeles, 1986 November 12–16.

conveying high-level signals. A side chain signal with the SR process characteristic is additively combined with the main signal in the encoding mode and subtractively in the decoding mode, whereby an overall complementary action is obtained.

The SR stage layout resembles that of the C-type system, except that three levels or stages of action staggering are used: high level, mid level, and low level (HLS, MLS, LLS). There are various advantages arising from the use of multilevel stages, including accuracy and reproducibility, low distortion, low overshoot, and action compounding for good spectral discrimination. For the high-level and mid-level stages both high-frequency and low-frequency circuits are used, with a crossover frequency of 800 Hz. The low-level stage is high frequency only, with an 800-Hz high-pass characteristic.

Each stage above has a low-level gain of somewhat over 8 dB, whereby a total dynamic effect of about 16 dB is obtained at low frequencies, 24 dB at high frequencies. A further dynamic action of about 1 dB takes place above the reference level.

The spectral skewing network has the same purpose and function as in the C-type system, except that a spectral skewing action is provided at low frequencies as well. The spectral skewing networks desensitize the SR process to the influence of signal components at the extreme ends of the audio frequency band. This effect is particularly helpful if the recording or transmission system has an uncertain frequency response in these regions. The networks are also important in attenuating subsonic and supersonic interferences of all kinds. The spectral skewing action is compensated in the decoder, resulting in an overall flat frequency response.

Both high-frequency and low-frequency antisaturation networks are provided in the main signal path, again operating in substantially the same way as in the C-type system. There is an effective compounding of the antisaturation effects produced by the antisaturation networks and the spectral skewing networks. In this way the SR process achieves a significant increase in high- and low-frequency headroom.

2 GENERAL PRINCIPLES

2.1 Least Treatment Principle

A design philosophy used in the development of the new system is that the best treatment of the signal is the least treatment. The operating goal of the encoder is to provide fixed, predetermined gains for all frequency components of the signal, with corresponding attenuations in the decoder. If a large signal component appears at a particular frequency or frequencies, then the gains should be reduced at those frequencies only, in accordance with predetermined compression laws for restoration of the signal during decoding. In other words, the compressor should try to keep all signal components fully boosted at all times. When the boosting must be cut back at a particular frequency, the effect should not be extended to low-level signal components at other frequencies.

The audible effect of this type of compression is that the signal appears to be enhanced and brighter but without any apparent dynamic compression effects. (The ear detects dynamic action primarily by the effect of a gain change due to a signal component at one frequency on a signal component at some other frequency, somewhat removed.) If the ear cannot detect dynamic effects in the compressed signal, then 1) it is unlikely that noise modulation effects will be evident in the decoded signal, and 2) it is unlikely that signal modulation effects will be evident in the decoded signal if there should be a gain or frequency response error in the recording or transmission channel.

In the SR process two new methods are used that greatly reduce the circuitry required to achieve the design goal of a full spectrally responsive system. In particular, both fixed and sliding bands are used in a unique combination, called action substitution, that draws on the best features of both types of circuits. A further technique, called modulation control, greatly improves the performance of both the fixed and the sliding bands in resisting any modulation of signal components unless necessary.

The use of the new methods reduces the basic encoder to two frequency bands only (high frequency and low frequency), each with a fixed-band circuit and a sliding-band circuit (this combination being referred to as a stage). When the three-level action-staggering layout is taken into account, five fixed bands and five sliding bands are employed in the spectral recording process.

2.2 Action Substitution

A new type of compression and expansion action that is highly responsive to spectral changes can be achieved by superposing or overlaying the individual characteristics of different types of dynamic action circuits. One circuit may provide a quiescent characteristic or defining umbrella. A further characteristic is hidden until signal components appear that cause the hidden characteristic to be revealed and become active.

For discussion purposes let the gains in a compressor system be arranged such that subthreshold signals pass without attenuation. That is, the maximum possible action is that of providing a certain gain, unity, for instance. Somehow to achieve this gain over as broad a range of frequencies as possible, in the presence of higher level (dominant) signals, is the task of the system.

Thus, in a superposed action compressor circuit, represented in Fig. 2, a signal is fed into a first compressor circuit. The output from this circuit represents the completed part of the total potential action. The input signal minus the completed part is, therefore, the uncompleted part. This is so derived and fed into the next compressor circuit, which has some different characteristic. The output of the second circuit is then added to that of the first, augmenting the action of the first. In an extreme condition, in which the output of the first circuit may be negligible at a particular fre-

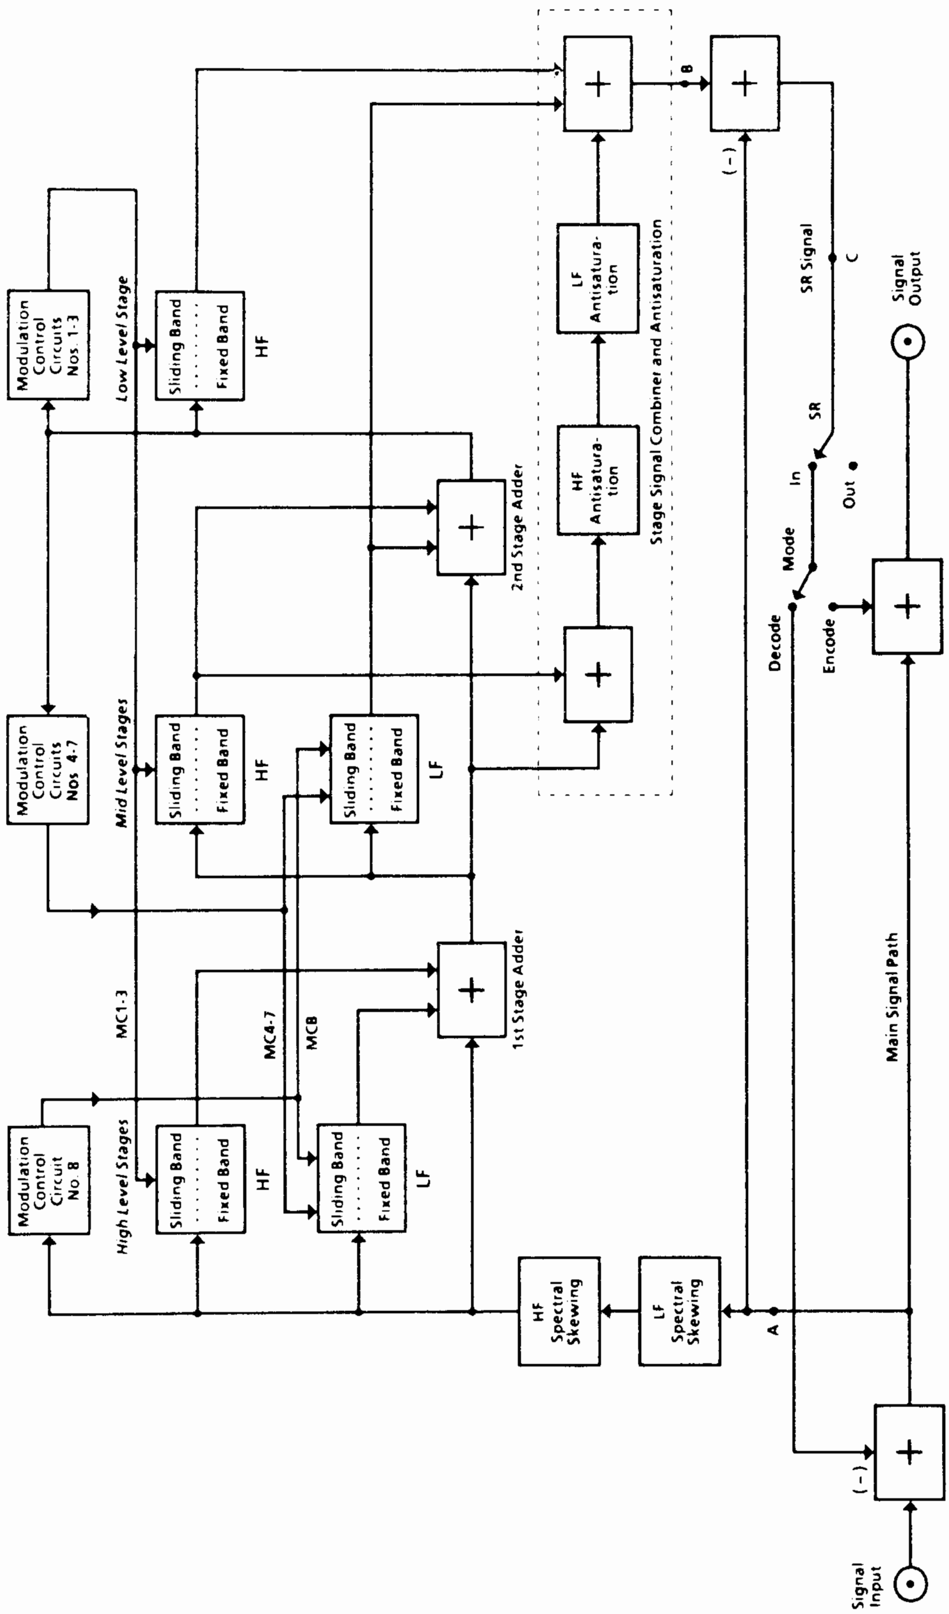

Fig. 1. Basic block diagram—spectral recording process. Diagram shows main features of process: multistage high-frequency and low-frequency dynamic action staggering, action substitution (fixed bands and sliding bands), modulation control, high-frequency and low-frequency spectral skewing, and high-frequency and low-frequency antisaturation.

quency, the action of the second circuit is effectively substituted for that of the first.

The operation of the action substitution compressor can be characterized directly from the above description. With an input signal V_i and an output signal V_o, a first compressor transfer function $F_1(s)$ and a second compressor transfer function $F_2(s)$, each being responsive to the respective applied signal, we have

$$V_o = V_i \, [F_1(s) + F_2(s) - F_1(s)\, F_2(s)] \; . \qquad (1)$$

This equation shows that the overall transfer function is the sum of the individual transfer functions minus their product. In other words, to the extent that the transfer functions may overlap, a factor is subtracted from the sum of the transfer functions.

The above type of action can be achieved with various circuit topologies, the one used in the initial implementation of the SR system being shown in Fig. 3. In this arrangement, the compressor circuits are arranged in a stack. Both circuits are fed in parallel and the output is taken from the top circuit, which is configured as a three-terminal network with terminals a, b, and c. The output of the bottom circuit is fed to the reference terminal c of the top circuit. It can be shown that the signal components at the output terminal b are those specified by Eq. (1).

The usefulness of the superposition technique can be appreciated by consideration of Fig. 4(a) and (b). The advantages of fixed-band compressor circuits [Fig. 4(a)] arise from the fact that all signal frequencies within the band are treated equally, in contrast with sliding-band action [Fig. 4(b)]. Thus the appearance of a dominant signal component actuating the compressor results in a loss of noise reduction effect that manifests itself in a uniform manner throughout the band, 2 dB in the example shown. The loss is not concentrated in any particular frequency region as it is in sliding-band circuits; note the 5-dB loss shown in the example of Fig. 4(b). The main significance of this is that if the recorder or transmission channel has an error in gain and/or frequency response, there is no undue exaggeration of the error at other, nondominant signal frequencies. In sliding-band circuits the amplification effect may be significant (the midband modulation effect, discussed in [2]).

In contrast, the advantages of sliding-band compression and expansion circuits derive from the fact that all signal frequencies are *not* treated equally. In particular, compression, expansion, and noise reduction action are well maintained above the frequency of the dominant signal component in high-frequency circuits, and below the frequency of the dominant signal component in low-frequency circuits. This action maintenance effect, except on a one-to-one basis, is absent in fixed-band circuits

Clearly it would be desirable to have the benefit of fixed-band operation on the stop-band side of the dominant signal frequency and sliding-band operation on the pass-band side. The action substitution technique provides this useful combination. In Fig. 4(c) the response of an action substitution compressor to the signal conditions of the two previous figures is shown. As is seen, the output is primarily from the fixed band for frequencies up to the dominant signal component and from the sliding band above that frequency. Conversely, for a low-frequency stage the output is from the fixed band for frequencies down to a low-frequency dominant component and from the sliding band below that frequency. This cooperative effect is particularly useful in the level region from the circuit threshold up to some 20 dB thereabove.

In the SR process, action substitution operation is used in both the high- and the low-frequency circuits. Thus both fixed-band and sliding-band dynamic actions are used in each of the five stages, a total of ten compressor circuits. While there is an effective interaction of the fixed and sliding bands in any particular stage, all of the stages operate independently. Depending on the levels and spectral conditions in each stage, fixed-band operation is used whenever it provides best performance; sliding-band operation is substituted whenever it has an advantage. The substitution is effective on a continuous and frequency-by-frequency basis.

Even though the frequency division of the stages is nominally 800 Hz, the use of what are effectively single-pole band-defining filters results in a significant overlap region between the high- and low-frequency stages. The high-frequency stages extend their effects down to about 200 Hz, the low-frequency stages extending their effects up to about 3 kHz. This overlap, together with the use of action substitution, contributes to the achievement of a very good spectral tracking effect

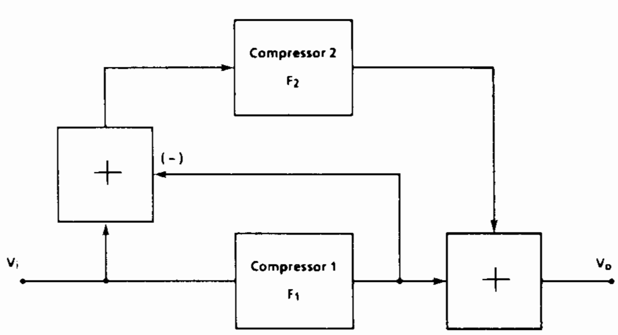

Fig. 2. Basic layout of action substitution compressor.

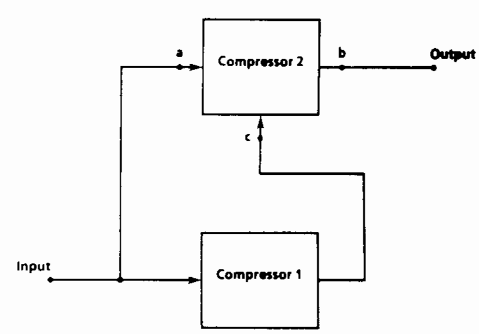

Fig. 3. Another action substitution compressor configuration.

under all frequency and level conditions. The practical significance is that an excellent noise reduction effect is obtained in the presence of signals, and that the system has a remarkable tolerance to gain and frequency response errors in the signal channel.

A further aspect of action substitution relates to the transient recovery characteristics of the system. A fixed-band compressor circuit has a recovery time that is essentially independent of frequency, at least in the pass band. A sliding-band circuit has a fast recovery time for nondominant signals at the pass-band end of the spectrum, and a slow recovery time for nondominant signals at the stop-band end of the spectrum. The choice of integrator recovery times is therefore a matter of compromise between this recovery time situation and the amount of steady-state and modulation distortion obtained. The compromise is made much easier by the use of the action substitution technique. In particular, the fixed band provides a definite and rapid recovery time for the overall system, so that the sliding band can employ longer time constants than would otherwise be desirable. This results both in low modulation distortion and a fast recovery time.

Thus the action substitution technique provides the advantages of fixed and sliding-band circuits while avoiding their disadvantages. In other words, there is a significantly improved adherence to the ideal of least

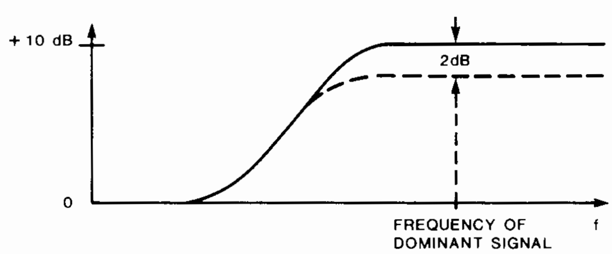

(a) Fixed-band compressor characteristic.

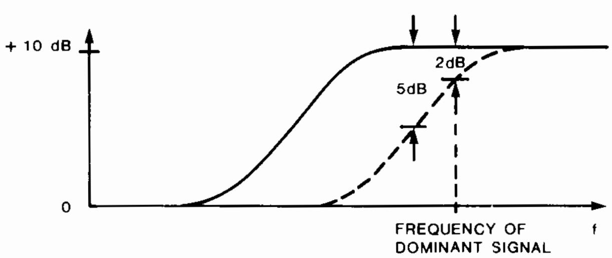

(b) Sliding-band compressor characteristic.

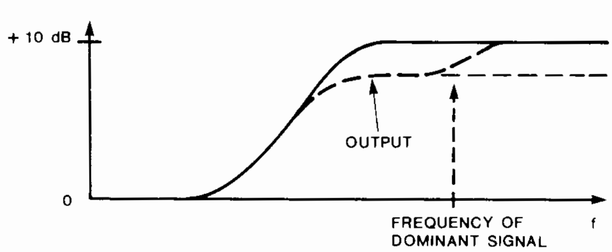

(c) Action substitution compressor characteristic.

Fig. 4. Types of compressor characteristics.

signal treatment. In the level region somewhat above the circuit threshold the signal more closely approaches fully boosted conditions in the encoding mode, with a consequently improved noise reduction effect in the decoding mode. For signals at higher levels, the technique of modulation control, described below, is employed.

2.3 Modulation Control

In the A-type, B-type, and C-type systems the signal from the side chain is highly limited under high-level signal conditions. This high degree of limiting, beginning at a low-level threshold, is responsible for the low distortion, low overshoot, and low modulation distortion which characterize these systems.

A closer examination shows that it is unnecessary to utilize such a low threshold and such a strong limiting characteristic under certain signal conditions. In particular, whenever the side chain signal departs from an in-phase condition with respect to the main path signal, then the threshold can be raised. Moreover, after an appropriate degree of limiting has taken place at a given frequency (in order to create the desired overall compression law), then it is unnecessary to continue the limiting when the signal level rises even further. Rather, the level of the side chain signal can be allowed to rise as a function of a further increase in signal level, whereby it stabilizes at some significant fraction of the main path signal level.

In the fixed-band portions of the spectral recording process the above arrangement results in conventional performance in the pass-band (in-phase) frequency region. However, in the stop-band region the modulation control scheme causes the limiting threshold to rise and the degree of limiting to be reduced. The possibility of doing this can be appreciated by consideration of the phasor diagrams of the two conditions shown in Fig. 5. In the pass-band (in-phase) condition the side chain signal and the main path signal add directly. Therefore a relatively low threshold must be maintained at all pass-band frequencies [Fig. 5(a)]. However, in the stop band the effective amplitude contribution of the side chain signal may be minimal due to the phase difference between it and the main path signal. Because of this it is possible to raise the threshold significantly and to reduce the limiting strength once the desired amount of attenuation has been obtained at a given frequency [Fig. 5(b)]. The result is that large signals in the stop band do not cause signal modulation in the pass band and consequently create an impairment of the noise reduction effect achieved during decoding.

Similar considerations apply in the SR sliding-band circuits. By way of introduction, in the B-type and C-type sliding-band circuits a variable filter follows a fixed filter, which has proved to be an efficient and reproducible arrangement. At frequencies outside the pass band a pure two-pole filter would result in overall amplitude subtraction from the main path signal because of the large phase angles created. Therefore the type of filter which has been employed is only quasi-two pole (a single-pole fixed filter plus a variable-shelf characteristic).

The same type of arrangement is used in the spectral recording process, with a one-octave difference (in the stop-band direction) between the variable filter turnover frequency (under quiescent conditions) and the fixed filter cutoff frequency. Above the threshold at a particular frequency the variable filter slides to the turnover frequency needed to create the overall (main path plus side chain signal) compression law. As the input level rises, and once an overall gain of about unity is obtained—when the variable filter cutoff frequency is about two to three octaves above the dominant signal frequency—there is no reason for further sliding of the variable filter. At this point the modulation control arrangement counteracts the sliding. As with the fixed-band circuits, this technique prevents unnecessary modulation of the signal.

The above effects in both the fixed and the sliding bands are created by circuits called modulation control circuits. Suitably filtered or frequency-weighted signals from the main signal path are rectified, and in some cases smoothed, and are fed in opposition to the control signals generated by the control circuits of the various stages. The result at higher signal levels, relatively (beginning at about 20 dB above the threshold of the relevant compressor circuit), is to tend to create a balance or equilibrium between the compressor circuit control signals and the modulation control signals. Under these conditions there is a significantly reduced gain reduction or sliding of the variable filters as a function of increasing input signal levels.

Fig. 6 illustrates the action of modulation control with a high-frequency fixed-band compressor circuit. The circuit has a low-level gain of about 8 dB and an 800-Hz high-pass characteristic. Fig. 6(a) shows the response of the circuit in the absence of modulation control. Ideally there should be no attenuation in response to a 100-Hz signal because the overall shape of the envelope is such that there is negligible signal boosting at 100 Hz. Nevertheless, with a conventional compressor circuit as shown here, when the 100-Hz signal increases in level, there is a reduction of low-level signal boosting over the whole frequency band. The unnecessary attenuation has two effects: 1) substantial noise reduction action is lost during expansion, and 2) when the amplitude of the 100-Hz signal varies, it can modulate low-level signal components at higher frequencies, resulting in possible incorrect restoration of the signal by the expander if the recording channel has an irregular frequency response in the vicinity of 100 Hz.

Fig. 6(b) shows the operation of the same circuit with modulation control. A greatly reduced attenuation occurs when the 100-Hz signal is varied over the same

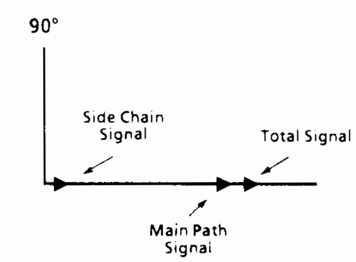

(a) In the pass-band, a low threshold and strong limiting characteristic are required.

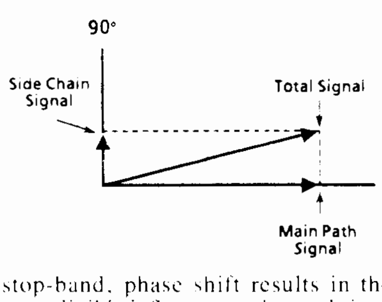

(b) In the stop-band, phase shift results in the side chain signal having negligible influence on the total signal amplitude; therefore a higher threshold and weaker limiting can be used.

Fig. 5. Phasor diagrams, dual path compressor.

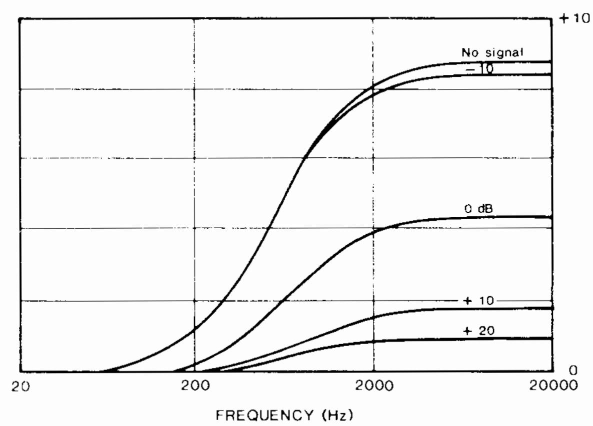

(a) Frequency response curves with 100-Hz signal at the levels indicated, no modulation control.

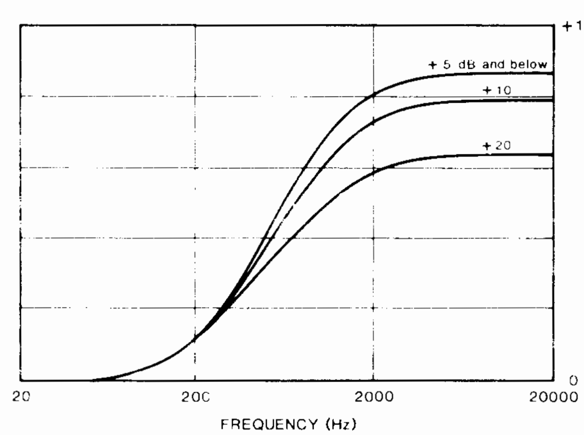

(b) Same as 6(a), with modulation control.

Fig. 6. Effect of modulation control on fixed-band compressor circuit.

range of levels as in Fig. 6(a). Thus a significant immunity to strong signals in the stop-band frequency region is achieved, the effect decreasing as the dominant signal frequency approaches the pass-band frequency region of the circuit.

In Fig. 7(a) the operation of a sliding-band circuit under comparable conditions is shown. As with the fixed-band circuit, ideally there should be no sliding in response to a strong 100-Hz signal. Nonetheless, as the 100-Hz signal increases in level, the band slides upward. As with the fixed-band circuit, the unnecessary sliding results in a loss of noise reduction action and the modulation of signals at higher frequencies when the sliding band varies under the control of the 100-Hz signal.

Fig. 7(b) shows the operation of the same circuit with modulation control. Minimal sliding occurs when the 100-Hz signal is varied over the same range of levels as in Fig. 7(a). Thus the sliding-band compressor is also made essentially immune to strong signals outside its pass band.

The effect of modulation control is further illustrated by Fig. 7(c) and (d), except that the frequency of the dominant signal is changed to 800 Hz, a frequency within the pass band of the circuit. Ideally, sliding is required to go only so far as not to boost the 800-Hz signal above the 0-dB reference level. Thus in the Fig. 7(c) response, without modulation control, the sliding produced by the 800-Hz signal at levels above −10 dB is excessive. Fig. 7(d) shows the response of the circuit with modulation control: sliding at and above the 0-dB level is greatly reduced. The effect is progressively reduced for low signal levels, but is observable to some extent at the −10-dB level.

The use of modulation control techniques also has advantages under transient conditions, both from their use in the steady-state circuits and also because of their use in the overshoot suppression circuits. Modulation control generally prevents any further fixed-band attenuation or sliding of the variable filter than is required to respond to a given signal situation. Therefore 1) signal modulation is reduced, 2) the SR process is rendered very tolerant of channel errors, 3) subsequent noise modulation during decoding is reduced, and 4) recovery from transient signal conditions is faster. The electronic reality in both the steady-state circuits and transient control circuits is that the integrator capacitors are prevented from charging to voltages as high as they normally would in the absence of modulation control. With lower fully charged voltages, recovery is faster.

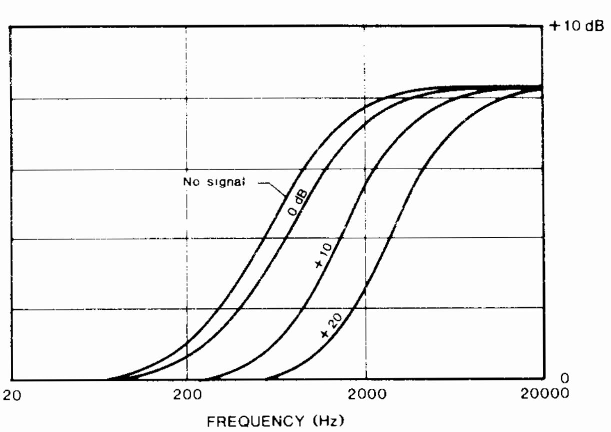

(a) Frequency response curves with 100-Hz signal at the levels indicated, no modulation control.

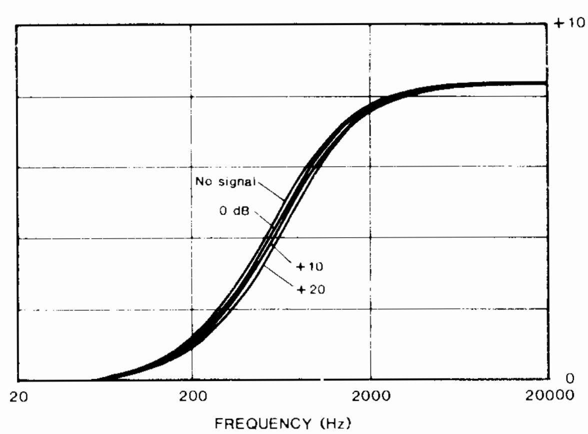

(b) Same as (a), with modulation control.

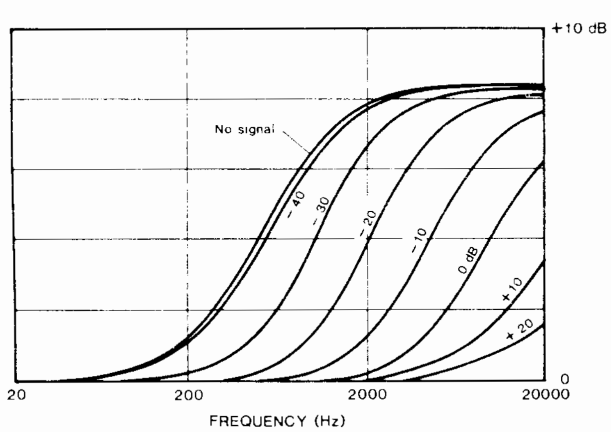

(c) Frequency response curves with 800-Hz signal at the levels indicated, no modulation control.

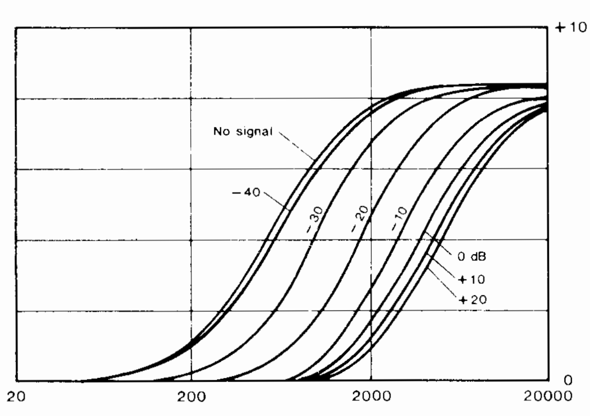

(d) Same as (c), with modulation control.

Fig. 7. Effect of modulation control on sliding-band compressor circuit, with a signal in the stop-band (a), (b), and a signal in the pass-band (c), (d).

The modulation control aspects of the SR process result in an encoding action which is remarkably free of noticeable signal-related modulation effects. Working together with action substitution, modulation control contributes to the goal of least treatment, in providing a highly boosted, audibly stable encoded signal.

2.4 Overshoot Suppression

A side effect of the modulation control scheme is that at high signal levels the amplitudes of the signals from the several stages are relatively high in comparison with the situation in the A-type, B-type, and C-type systems. Because of this it is not possible to employ simple overshoot suppression diodes as in these previous systems. A more flexible but necessarily more complex scheme, operating directly upon the control signals, is used.

In common with the A-type, B-type, and C-type systems, the SR process features overshoot suppression thresholds that are significantly higher than the steady-state thresholds; this results in low modulation distortion. The overshoot suppression thresholds are set about 10 dB above the relevant steady-state thresholds. The net result is that for most musical signals the overshoot suppressors rarely operate; the compressors are controlled by well-smoothed signals. When the suppressors do operate, the effect is so controlled that modulation distortion is minimal.

Under extreme transient conditions, such as from a subthreshold signal situation, the overshoot suppression threshold is set at its lowest point. The overshoot suppression effects are then phased in gradually as a function of increasing impulse level.

Under relatively steady state but nonetheless changing signal conditions the overshoot suppression effects are gradually phased out as a function of increasing signal levels, this action further ensuring low overall modulation distortion from the system. The phasing out effect is achieved by increasing the overshoot suppression thresholds. The thresholds are controlled by signals that are the same as or derived from the modulation control signals used to control the steady-state characteristics, whereby a tracking action between the transient and steady-state behavior is obtained. This arrangement results in both well-controlled overshoots and low modulation distortion.

Both primary and secondary overshoot suppression circuits are employed, the latter acting as fall-back or long-stop suppressors. In the high-frequency circuits the secondary overshoot suppressors improve the performance just inside the stop band—that is, in the 400–800-Hz region. In the low-frequency circuits these additional suppressors improve the performance under extreme complex signal conditions (such as high-level low and mid-frequency transient signals in combination with high-level high-frequency signals). In the low-frequency circuits a further overshoot suppressor (LF O/S) is used for very-low-frequency signals. This is a very gentle, slow-acting circuit which reduces low-frequency transient distortion.

2.5 Staggered Action Multilevel Format

The principles discussed above are incorporated into each stage of the multilevel staggered action encoder and decoder. See [2] for a detailed discussion of staggered action circuits. In the SR system two stages are employed at low frequencies, three at high frequencies. The thresholds used are approximately −30 dB, −48 dB, and −62 dB below reference level (about 20 dB below SR peak signal level). In the series-connected staggered action format there is a compounding of the actions of the individual stages; the transfer functions of the several stages are multiplied, whereby the dB characteristics add. In this way a large total dynamic action can be achieved with low modulation distortion, low overshoot, and good manufacturing reproducibility. An important additional result is that there is an overall steepness enhancement of the frequency discrimination abilities of the circuit, further inhibiting signal modulation and noise modulation effects.

2.6 Spectral Skewing

The spectral skewing networks employed in the SR process comprise both high-frequency and low-frequency sections, with the same rationale and mode of operation as discussed in [2]. The spectral distributions of the signals processed by the encoder are altered or skewed, well within the pass band, such that the encoder action is significantly less susceptible to the influence of signals beyond the abrupt roll-off frequencies of the spectral skewing networks.

The high-frequency network is a low-pass filter with an attenuation characteristic similar to that of a 12-kHz two-pole Butterworth filter, but with a limiting attenuation of about 35 dB (that is, a shelf). The low-frequency network is a 40-Hz high-pass filter, connected in series with the high-frequency network, also with a two-pole Butterworth-like characteristic, but with about a 25-dB limiting attenuation. These shelves do not interfere with the attenuation within the audio band, but provide phase characteristics that are essential in the decoding mode.

2.7 Antisaturation

The general principle of antisaturation was described in [2]. Briefly, by placing a fixed attenuation network, usually a shelf, in the main path of a dual path compressor, it is possible to create an effective antisaturation characteristic at the extremes of the audio band without undue adverse effects on the noise reduction achieved during decoding. In the SR process high- and low-frequency networks are operative above about 5 kHz and below about 100 Hz, respectively. In addition, the spectral skewing networks have a secondary but very useful antisaturation effect, especially at very low and very high frequencies.

3 BLOCK DIAGRAMS

3.1 Basic Block Diagram

As mentioned previously, Fig. 1 shows the basic layout of an SR processor. While the whole system comprises an encoder and a complementary decoder, the figure as drawn shows a switchable configuration, which generally is the most useful one. The main signal path transmits high-level signal components. To this is added in the encoding mode, and subtracted in the decoding mode, the output of the side chain circuitry, called the SR signal, point C. The stage circuits, as well as the spectral skewing networks and antisaturation networks above, are driven from point A. See [1] for a mathematical explanation of these arrangements.

A secondary main path which does not include any antisaturation is employed as the basis of the side chain, to which the outputs of the high-level and mid-level stages are added in the first-stage and second-stage adders, respectively. The low-level stage and modulation control circuits 1–7 are driven directly from the output of the second-stage adder. Modulation control circuit 8 is driven from the output of the spectral skewing network, as will be discussed.

The antisaturation effects are created in the dashed block labeled stage signal combiner and antisaturation. The arrangement shown provides a high-frequency deemphasis effect on the secondary main path signal, which includes the output of the high-level stages, and on the high-frequency mid-level stage signal. This deemphasis is effective not only on the steady-state aspects of these signals, but on all transient effects as well. The output of the low-level stage is then added directly. For low-frequency antisaturation the low-frequency deemphasis is effective on the secondary main path signal, including the high-level stage outputs. The low-frequency mid-level signal is then added directly. The mathematical basis for these arrangements is provided in [2].

With the final combination of signals in the last adder, an SR encoded signal appears at point B. The encoded signal can be considered to comprise an unmodified component from the input plus an SR signal which carries all of the SR characteristics. Thus the SR component can be derived by subtracting the unmodified input signal at point A from the SR encoder output at point B. This provides an SR signal at point C that can be handled and switched the same way as in the A-type system. This simplifies practical use of the system.

3.2 Modulation Control Circuits

Fig. 1 shows how the inputs of the various modulation control circuits are connected and how the resultant signals are distributed. Modulation control signals MC1–MC7 are derived from the output of the second-stage adder. In this way the modulation control signals begin to have a significant influence at relatively low levels, such as at −30 dB (because of the contributions of the HLS and MLS stages); the phase relationships between the modulation control signals and the signals in the control circuits of the several stages are also optimized. In the generation of MC8, which is used for low-frequency stage overshoot-suppression inhibition under high-frequency transient signal conditions, the influences of the HLS and MLS stages are undesirable. MC8 is therefore derived from the first feed point of the stages, just following the spectral skewing networks.

Fig. 1 also shows the distribution scheme of the modulation control signals. MC1–MC3 are used for the high-frequency stages; MC5–MC8 are used for the low-frequency stages.

In Fig. 8 the basic layout of the modulation control circuits is shown. MC1 controls the high-frequency sliding-band circuits. The signal from the takeoff point is fed through a 3-kHz single-pole high-pass filter, full-wave rectified (all rectifiers in the system are full wave), and fed in opposition to the control signals generated by the high-frequency stages. An all-pass phase shift network is used to optimize the phase of the MC signal in relation to the stage control signal at low frequencies; this reduces control signal ripple. MC1 is also smoothed by a two-stage 1-ms integrator and is employed, as MC2, to oppose the operation of the high-frequency sliding-band overshoot suppression circuits; the overshoot suppression thresholds thereby track the steady-state thresholds. MC2 must be smoothed because the phase relationships of MC1 and the signals in the stages vary (because of the sliding-band action) throughout the audio band, being a function of frequency and level.

MC3 controls the high-frequency fixed-band circuits. The signal from the takeoff point is weighted by cascaded 400-Hz and 800-Hz single-pole low-pass filters, rectified, and fed in opposition to both the steady-state and the transient control circuits of the high-frequency fixed-band circuits. There is no need to provide a smoothed MC signal for the overshoot suppressors of the high-frequency fixed-band stages because a fixed phase relationship exists between the stage signals and the control signals throughout the audio band.

MC4 controls the sliding-band circuits of the low-frequency stages. The signal from the takeoff point is fed through a 200-Hz single-pole low-pass filter, rectified, and fed in opposition to the sliding-band control signals generated in the stages. The phase of the modulation control signal is optimized by the use of an all-pass phase shifter, as with MC1; low-frequency control signal ripple is thereby reduced. MC4 is smoothed by a two-stage 2-ms integrator to form MC5; this signal is used to control the low-frequency sliding-band overshoot suppressors.

MC6 controls the low-frequency fixed-band circuits. The signal from the takeoff point is weighted by cascaded 800-Hz and 1.6-kHz single-pole high-pass filters, rectified, and used to oppose the steady-state fixed-band control signals. MC6 is also smoothed in a two-stage 2-ms integrator, forming MC7, which is used to control the low-frequency fixed-band overshoot suppressors. This smoothing is necessary in the low-frequency fixed-band stages because, unlike the situation in the high-frequency fixed-band stages, there is no fixed phase relationship between the stage signals and the overshoot suppression signals. MC7 is also used in a supplemental way to control the low-frequency sliding-band overshoot suppressors.

MC8 is used to control the overshoot suppression circuits of both the fixed- and the sliding-band low-frequency circuits. MC8 compensates for the fact that no frequency weighting is used in the generation of the low-frequency primary overshoot suppression signals. High-frequency transient signal components are detected and used to oppose the operation of the LF primary overshoot suppression circuits. The signal from the MC8 takeoff point is fed through a 5-kHz high-pass filter, rectified, double differentiated with 15-μs time constants, and peak held with about a 30-ms time constant. The resultant high-frequency transient modulation-control signal MC8 is then employed to oppose the low-frequency overshoot suppression action.

3.3 High-Frequency Stage

Fig. 9 shows both the steady-state and the transient control aspects of the high-frequency stages. The diagram shows only the basic parameter determining elements; the practical circuits of course contain other details such as buffering, amplification, and attenuation. The high-level, mid-level, and low-level stages have the same basic block diagrams and schematics. The main distinctions are that the ac and dc circuit gains are increased for the mid- and low-level stages.

Referring to the block diagram, each stage comprises a fixed-band section on the bottom and a sliding-band section on the top, each with its own control circuits. The fixed- and sliding-band circuits are fed in parallel, and the output signal is taken from the sliding-band circuit. The sliding-band variable filter is referenced to the output of the fixed band; that is, the fixed-band output is fed directly to the bottom end of the sliding-band variable resistance RV_s. This connection results in the action substitution operation discussed previously. At any given frequency the overall output will be the larger of, or some combination of, the fixed- and sliding-band contributions. If there is a signal situation in which the fixed-band output is negligible, then the sliding band predominates. Conversely, if there is little or no sliding-band contribution, the output from the fixed band will still feed through to the output through RV_s. In this way the action of one circuit augments that of the other, and, as the occasion requires, may be substituted for that of the other.

The incoming signal is fed through an 800-Hz single-pole band-defining filter. This is followed by a 400-Hz single-pole filter which attenuates the low-frequency signal levels fed to both the fixed- and the sliding-band circuits; this reduces waveform distortion and complex-signal transient distortion at high signal levels. The filter also forms part of the fixed-band control signal weighting network. The output signal is taken from the sliding-band stage and is fed through a 400-Hz network having a reciprocal characteristic to that of

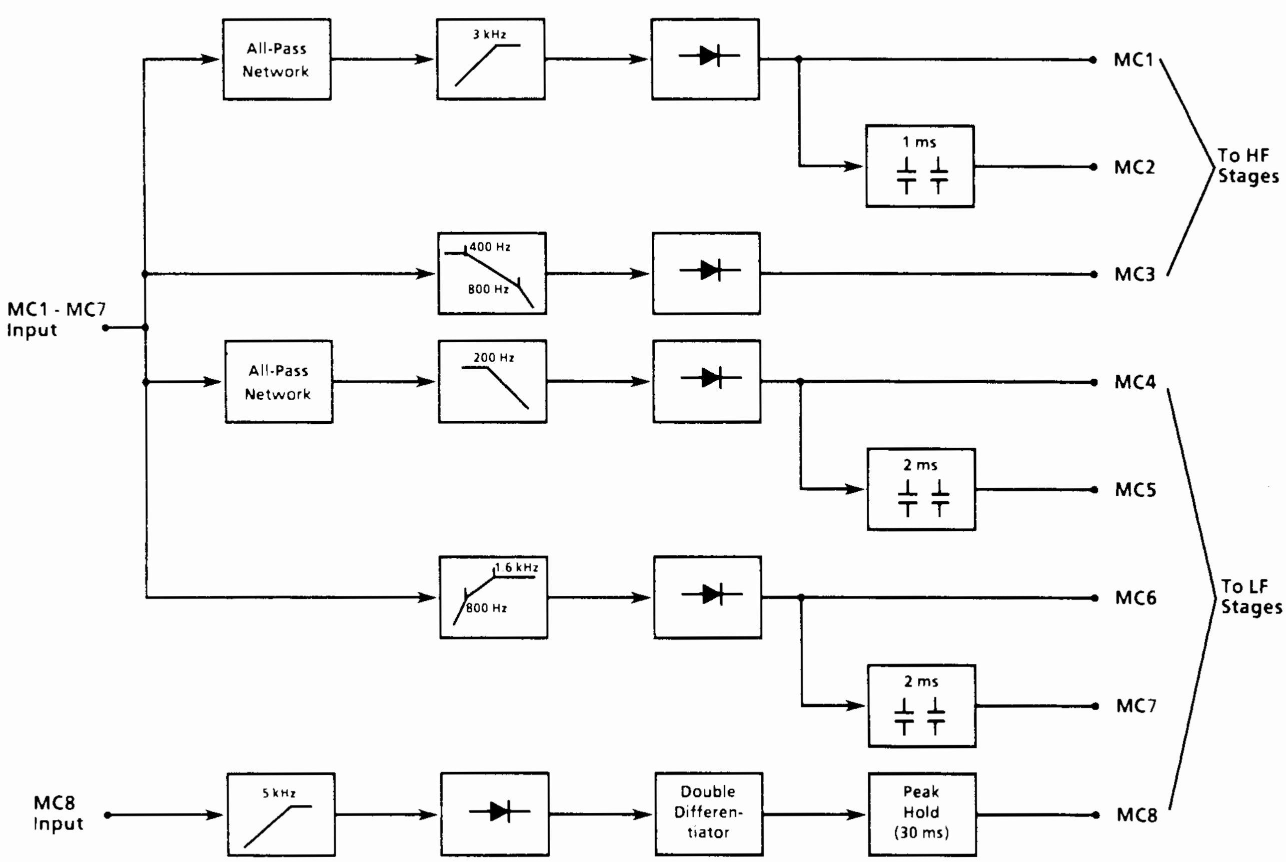

Fig. 8. Basic layout of modulation control circuits. These circuits reduce modulations of the gains and frequency response characteristics used in system, especially at signal levels significantly above the compressor thresholds.

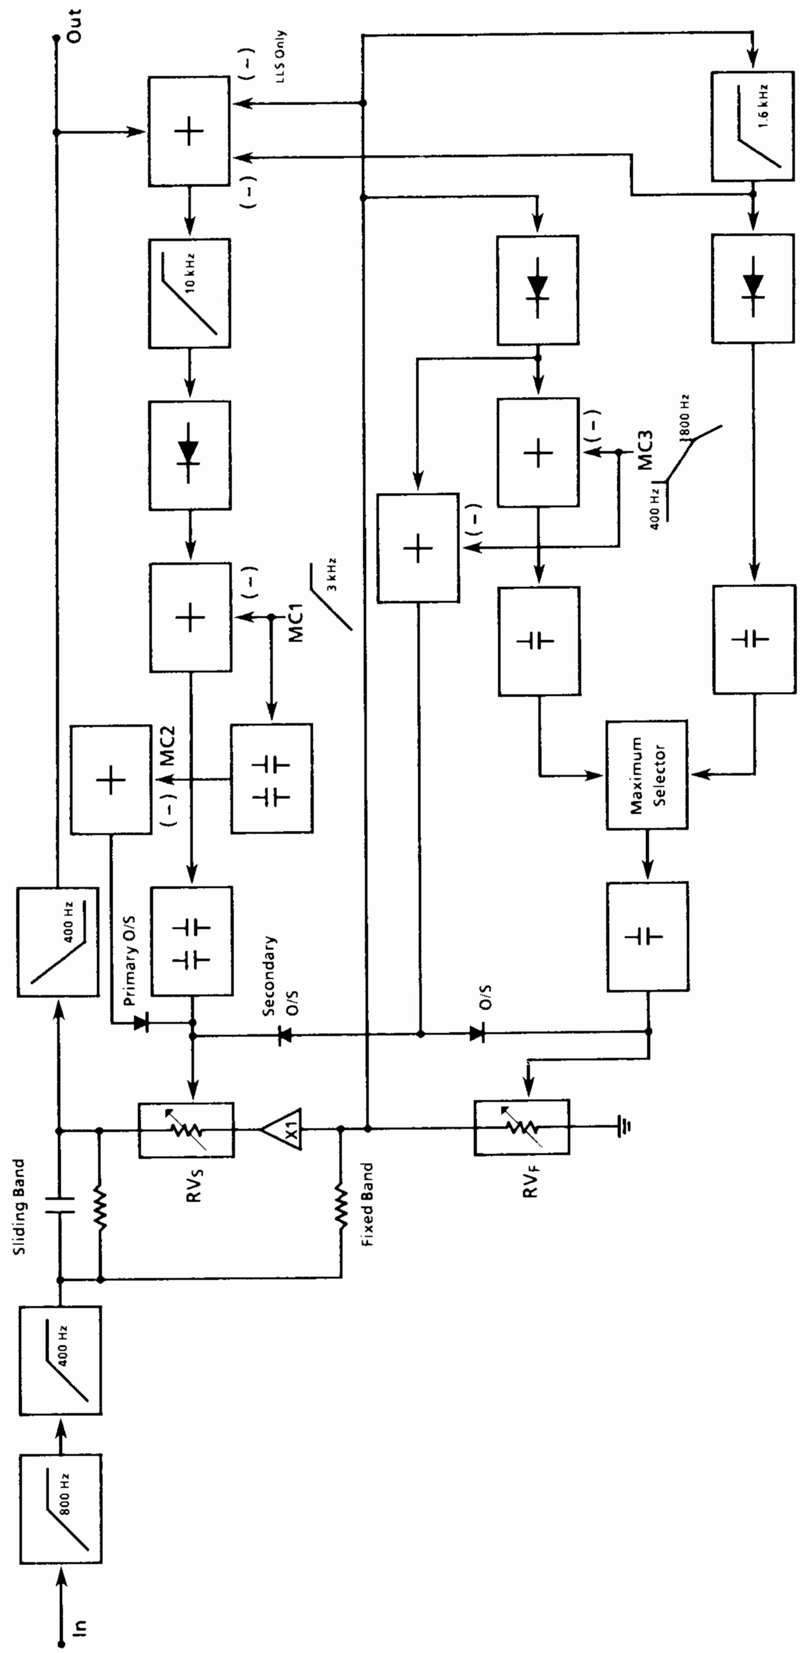

Fig. 9. High-frequency stage block diagram.

the 400-Hz high-pass filter at the input. Thus the overall quiescent (subthreshold) frequency response of the circuit is that of a single-pole 800-Hz high-pass network. The low-frequency stages have a complementary 800-Hz single-pole low-pass characteristic, which overall results in optimal combination of the signals from the high- and low-frequency stages.

The fixed-band output from RV_f (that is, variable gain circuit) is fed to two control circuits, the main control circuit (middle of Fig. 9) and the pass-band control circuit (bottom of Fig. 9). In the main control circuit the signal is rectified and opposed by the modulation control signal MC3. The resulting dc signal is smoothed by an integrator circuit with a 15-ms time constant (the overall steady-state control signal characteristic in this and all other stages is average responding). The control signal is then fed to one input of a maximum selector circuit, which passes to its output the larger of two signals applied to the input.

The fixed-band output is also fed to the pass-band control circuit (bottom of Fig. 9), which comprises a 1.6-kHz single-pole high-pass filter, a rectifier, and a smoothing circuit (15 ms). The pass-band control signal is applied to the other input of the maximum selector circuit. The output of the maximum selector circuit is further smoothed by a 160-ms time constant and is used to control the fixed-band variable resistance RV_f or other variable gain means.

The dual control circuit arrangement described above is employed to obtain optimal performance under both simple-signal (single dominant signal) and complex-signal (more than one dominant signal) situations. The modulation control signal MC3 is optimized in frequency weighting and amount for simple signal conditions, in which the modulation control action is most useful. Under complex signal conditions, however, the modulation control signal developed becomes larger, and the subsequent modulation control action is then greater than necessary; that is, the dc control signal output from the main control circuit is less than required. Under this condition the control signal from the pass-band circuit is phased in, via the maximum selector circuit, to control the overall action of the fixed-band compressor circuit.

The output of the fixed band is fed through a buffer with an overall gain of unity to provide the reference for the sliding band filter. This is the only signal output of the fixed-band compressor circuit.

The sliding-band control signal is derived from the stage output (top of Fig. 9). The signal is fed through a single-pole high-pass weighting network (about 10 kHz, different for each stage) and is rectified. The rectified signal is opposed by modulation control signal MC1. Since MC1 also has a single-pole high-pass characteristic, the ratio between the rectified control signal and MC1 monitors the signal attenuation (this ratio creates an end-stop effect on the sliding-band action). The result is smoothed first by a time constant of about 5 ms (different for each stage), and finally by a time constant of 80 ms. The smoothed control signal is then used to control the sliding-band variable resistance RV_s. A single control circuit suffices in the sliding-band circuit because the 10-kHz high-pass control weighting network tends to offset the effect of complex signals on the modulation control voltage developed (MC1).

A modification is made in the sliding-band control characteristic at low levels. Signals from the fixed-band circuit are combined in opposition with the sliding-band output signal (combining circuit at right of Fig. 9). The effect is in the direction of simulating the derivation of the sliding-band control voltage from the voltage across the sliding-band variable filter only (that is, from the voltage across RV_s). This tends to raise the sliding-band threshold at high frequencies, which reduces unnecessary sliding of the band. (The 10-kHz control weighting network provides the correct amount of control signal for the variable filter at medium and high levels, but it produces the undesirable side effect of lowering the threshold at high frequencies. The differential control signal derivation method counteracts this threshold-lowering effect.)

The overshoot suppression (O/S) arrangements are also shown in Fig. 9. In the high-frequency circuits a general feature is that unsmoothed rectified signals from the control circuit rectifiers are opposed by appropriate modulation control signals and are fed via diodes to the final integrator circuits. The low-frequency arrangements follow the same pattern, with some modifications.

Referring to the middle of the diagram, in the high-frequency fixed-band circuit the overshoot suppression signal is derived from the rectifier of the main control circuit. As with the steady-state control signal, the rectified signal is opposed by MC3, so that the overshoot suppression threshold is appropriate for conditions in the steady-state regime. The resultant overshoot suppression signal is coupled by a diode to the final integrator circuit.

In the sliding-band circuit (top of Fig. 9) two overshoot suppression signals are used, primary and secondary. The primary overshoot suppression signal is derived from the control circuit rectifier and opposed by MC2, a smoothed version of MC1 (MC1 controls the steady-state characteristics). The smoothing is necessary because, unlike the situation in the fixed-band circuit, there is no constant and favorable phase relationship between the signal in the control circuit and MC1 (because of the sliding band); the smoothing enables reliable bucking action to take place.

The secondary overshoot suppressor supplements the action of the primary overshoot suppressor under certain conditions. The primary overshoot suppression signal is derived from the same rectifier used in the steady-state control circuit, with the consequence not only of economy but of a favorable phase relationship between the overshoot suppression impulse and the signal to be controlled; this results in low transient distortion. However, the control circuit frequency weighting responsible for this situation also causes a reduction of control signal amplitude with falling frequency. A dc bias is used in the overshoot suppression circuit to create the required suppression threshold; when the signal amplitude in the overshoot suppression circuit decreases, the bias results in the overshoot suppression effect falling away faster than the signal amplitude. For frequencies below about 400 Hz a reduced overshoot suppression effect is appropriate because of the attenuation and phase shift of the stage input filter (see Fig. 5). However, in the 400–800-Hz region there is an overshoot suppression deficiency; this is compensated by feeding an appropriate amount of overshoot suppression signal from the fixed-band circuit into the sliding-band circuit. This supplemental signal is called the secondary overshoot suppression signal.

Regarding recovery times, the use of action substitution and modulation control both contribute to rapid action, as already mentioned. Nonetheless, reverse-biased recovery speed-up diodes are used in a fairly gentle way (series resistors) to provide a further increase in speed.

3.4 Low-Frequency Stage: Steady-State Aspects

Fig. 10 shows only the steady-state layout of the low-frequency stages. As with the high-frequency stages, only the basic parameter determining elements are shown. The high-level and mid-level low-frequency stages have the same block diagrams and circuits, but the ac and dc gains are increased for the mid-level stage; there are also some other minor differences.

Referring to the block diagram, certain similarities and differences may be noted with respect to the high-frequency diagrams. The dual-layer arrangement of the fixed band on the bottom and the sliding band on the top is similar. However, the sliding band acts downward, using a simulated inductance (gyrator circuit). As with the high-frequency stages, the fixed- and sliding-band circuits are fed in parallel, and the output signal is taken from the sliding-band circuit. The fixed-band output is coupled to the bottom of the sliding band to provide the action substitution operation described previously.

A notable difference from the high-frequency circuit is that the fixed 800-Hz band determining filter follows, rather than precedes, the variable filter. This arrangement has several advantages: 1) overshoot suppression signals can be generated without the delay inherent in a low-pass filter, resulting in lower transient distortion; 2) any transient distortion produced by the circuit is attenuated by the 800-Hz low-pass filter; and 3) noise generated by the gyrator circuit is attenuated by the filter. The price to be paid for these advantages is the resulting high signal levels that the variable resistances RV_F and RV_S must be capable of handling at high frequencies. (There is no active attenuation at all at very high frequencies, since sufficient passive attenuation and phase shift are provided by the 800-Hz low-pass filter.) Special control arrangements, called high mode, comprising complementary bootstrapping and control circuit gain boosting, enable the fixed- and sliding-

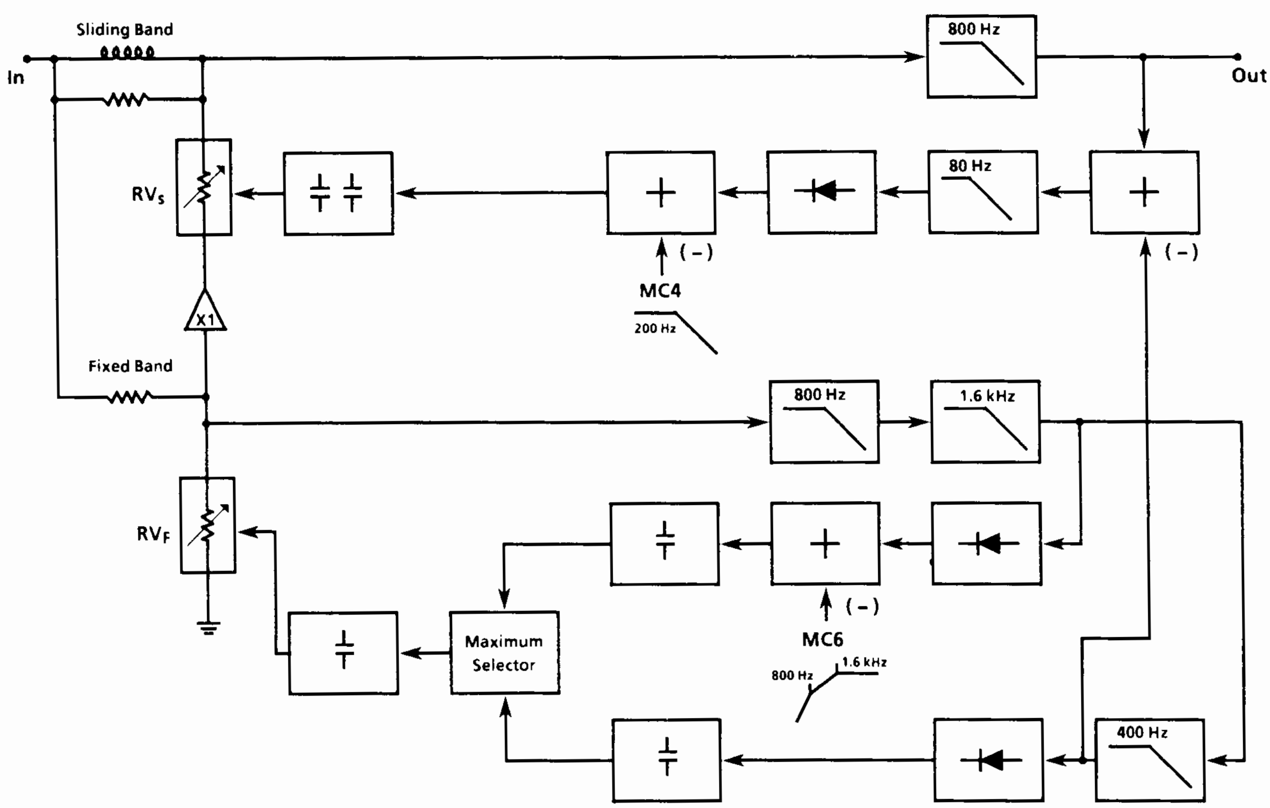

Fig. 10. Low-frequency stage block diagram—steady-state aspects.

band signal circuits to handle the required levels with low distortion and low noise.

Referring to the fixed-band section (lower half of Fig. 10), the incoming signal is applied directly to the variable gain circuit. Control circuit frequency weighting is provided by cascaded single-pole 800-Hz and 1.6-kHz low-pass filters. (The corresponding filters in the high-frequency stages are the 800-Hz and 400-Hz filters at the inputs of the circuits.) The main control circuit rectifies the filtered signal; the resulting dc signal is bucked by modulation control signal MC6, smoothed by a 15-ms integrator, and fed to one input of the maximum selector circuit. The maximum selector circuit has the same purpose and mode of operation as in the high-frequency circuits.

The 800-Hz and 1.6-kHz frequency-weighted output of the fixed-band circuit is also fed to the pass-band control circuit (bottom of Fig. 10). Here the control signal is further weighted by a 400-Hz single-pole low-pass filter, rectified, smoothed by a 15-ms integrator, and fed to the other input of the maximum selector. As in the high-frequency stages, the larger of the two signals is passed to the final integrator (300 ms) to become the fixed-band control signal applied to RV_F. In this way, both simple and complex signals are accommodated.

The sliding-band control signal, as in the high-frequency circuits, is derived from the stage output—that is, from a point following both the fixed 800-Hz band determining filter and the variable filter. The signal is frequency weighted by an 80-Hz single-pole low-pass filter, rectified, and bucked by modulation control signal MC4 (which also has a single-pole low-pass characteristic, with the same type of sliding-band end-stop effect as in the high-frequency circuits). The result is smoothed by a 7.5-ms integrator and finally smoothed by a 150-ms integrator to become the sliding-band control signal applied to RV_S. As in the high-frequency stages, a single control circuit suffices for the sliding band.

The same type of low-level control characteristic modification is made in the low-frequency circuits as in the high-frequency circuits. Namely, a signal from the fixed band is combined in opposition with the sliding-band output signal (see combining circuit at right of Fig. 10). This differential control modification raises the sliding-band threshold at low frequencies.

3.5 Low-Frequency Stage: Transient Control Aspects

The transient control (and steady-state) aspects of the low-frequency stages are shown in Fig. 11. In a manner generally similar to that of the high-frequency circuits, unsmoothed rectified signals from the outputs of the variable elements are opposed by appropriate modulation control signals and are fed via diodes to the final integrator circuits.

Both the fixed and the sliding bands each have primary and secondary overshoot suppressors, which operate at frequencies above about 100 Hz. In addition, each has a gentle and slow-acting low-frequency overshoot suppressor, operating at frequencies below about 200 Hz; there is a crossover effect between the two types of overshoot suppression in the 100–200-Hz region. The primary overshoot suppressors provide the earliest and strongest suppression effect in simple transient situations. With more complex signals the primary overshoot suppression thresholds rise, and eventually the secondary overshoot suppression circuitry takes control.

In contrast with the high-frequency situation, the low-frequency general strategy is to derive the primary overshoot suppression signals from signal points that do not include any control circuit frequency weighting. This is because the required control circuit weighting networks of the low-frequency stages are low-pass in character, resulting in delays. (The high-pass networks used for control circuit weighting in the high-frequency stages do not introduce delays.) However, because of the lack of a weighting factor in the primary overshoot suppression signal, there is no inherent tracking between the steady-state and overshoot suppression thresholds of the circuits involved, particularly in the stop bands. Therefore further modulation control techniques are employed to obtain the required tracking. The secondary overshoot suppression signals are derived from a point in the fixed-band circuitry that provides adequate tracking in both the fixed and the sliding bands.

Referring to Fig. 11, the fixed-band primary overshoot suppression signal is generated by passing the variable attenuator output through a 200-Hz single-pole high-pass filter (middle of figure). This filter reduces the influence of the primary overshoot suppressor at low frequencies, allowing the more gentle low-frequency overshoot suppressor to take over the transient control function. The signal is rectified and then opposed by modulation control signal MC7, a 2-ms smoothed version of MC6 (the fixed-band steady-state modulation control signal); the effect is in the direction of improving the steady-state and overshoot suppression threshold tracking on a steady-state basis. However, the thresholds must also track on a transient basis. This is the function of the high-frequency transient modulation control signal MC8, which is a high-frequency-weighted, peak-detected signal that opposes the primary overshoot suppression signal in the time interval before MC7 becomes effective. The overshoot suppression signal is then diode coupled to the final integrator circuit of the fixed-band circuit.

In the generation of the sliding-band primary overshoot suppression signal the output of the variable filter is fed through a 200-Hz single-pole filter (top of Fig. 11) in order to reduce the effect of the circuit at low frequencies (as in the fixed-band circuit). The signal is rectified and then opposed by modulation control signals MC5 and MC7 to provide an adequate degree of tracking between the steady-state threshold and the overshoot suppression threshold on a steady-state basis. As in the fixed-band circuit, MC8 provides the required degree of tracking on a transient basis. The resultant overshoot suppression signal is diode coupled to the

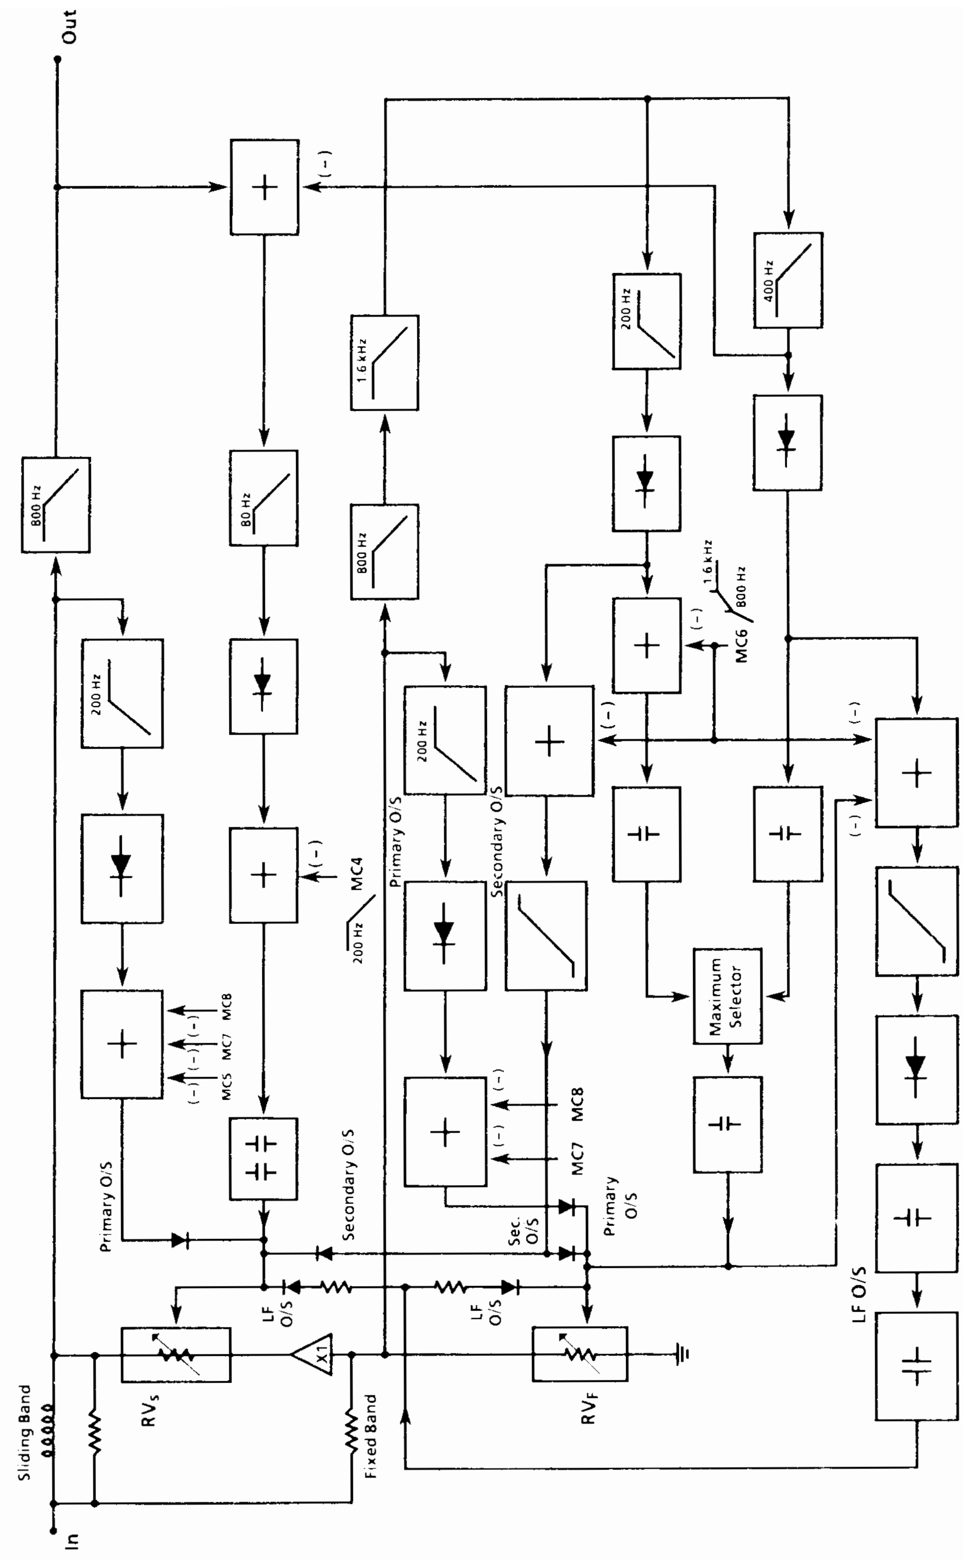

Fig. 11. Low-frequency stage block diagram—steady-state and transient control aspects.

sliding-band final integrator circuit.

In both the fixed- and the sliding-band circuits, the effects of the primary overshoot suppression circuits are maximized for the most significant transient signal situation—that is, a single impulse or toneburst starting from a subthreshold signal level. A side effect of the use of smoothed MC5 and MC7 signals is that the overshoot suppression levels for low- and medium-frequency transient signals are raised under certain complex signal conditions, especially those in which relatively steady-state high-frequency signals at high levels are also present. To compensate for this effect, secondary overshoot suppression signals are derived from the fixed-band main control circuit rectifier and are diode coupled to the fixed-band and sliding-band final integrator circuits. The secondary overshoot suppressors have higher thresholds than the primary suppressors and operate only rarely because of the unusual circumstances for which they are designed.

The secondary overshoot suppression signals are generated from the frequency-weighted point (800-Hz and 1.6-kHz pass) in the fixed-band steady-state control circuit (lower middle of Fig. 11). To prevent interference with the low-frequency overshoot suppression circuit at low frequencies, the signal is further filtered by a 200-Hz single-pole high-pass network, as in the primary overshoot suppression circuits; the filtered signal is then rectified. (Note that in Fig. 10 this filter is not shown, for clarity. On a steady-state basis the pass-band control circuit controls the circuit at very low frequencies, via the maximum selector circuit; this arrangement allows the main control circuit rectifier to serve a double function.) The dc signal is opposed by MC6 in order to phase out the secondary overshoot suppression effect at high frequencies. An optimal phase relationship is obtained between the rectified signal and MC6, apart from the effect of the 200-Hz filter (which is negligible). An ideal tracking effect is achieved between the steady-state and secondary overshoot suppression thresholds.

The effect of the 800-Hz and 1.6-kHz frequency-weighting networks is to introduce a time delay into the secondary overshoot suppression signal. The effective delay is significantly reduced by increasing the gain used in the secondary overshoot suppressor circuit and applying limiting. The resultant overshoot suppression signal is more in the nature of a nearly fixed amplitude impulse, applied in the rare circumstances when necessary, than it is a proportional response. The signal is coupled through a diode to the fixed-band final integrator circuit and is also used, suitably biased, for secondary overshoot suppression in the sliding-band circuit, also coupled through a diode.

The low-frequency overshoot suppression signal (bottom of Fig. 11) is developed by tapping the rectified, but unsmoothed, output of the pass-band control circuit of the fixed-band circuit. The signal is opposed by MC6 to desensitize the circuit to high-level, high-frequency components. The signal is further opposed by the resulting fixed-band smoothed control signal, in a negative feedback fashion. (When the fixed-band control signal has risen to a sufficient level, there is no further need for any low-frequency overshoot suppressor action.) The signal is then highly amplified and limited, peak rectified, and smoothed by an integrator with about a 20-ms decay time constant. The resulting high-amplitude pulses are fed through a differentiating network, with a time constant of the same order as the integration time constant, to provide low-frequency overshoot suppression impulses of defined strength for distribution to the fixed-band and sliding-band final integrators, via high-value resistors and series diodes. The result is a decaying "constant current" charging of the capacitors of the final integrator circuits. This is in contrast with the higher peak currents and correspondingly more abrupt control voltage changes produced by the relatively low-impedance primary and secondary overshoot suppressors. The use of the low-frequency overshoot suppression method results in low waveform distortion of relatively slowly changing low-frequency signal impulses applied to the system.

The control signal recovery characteristics are similar to those of the high-frequency circuits, although being about half as fast because of the longer time constants employed. As in the high-frequency stages, the recovery time is favorably affected by the use of action substitution and modulation control, but is further augmented by the use of speed-up diodes.

4 OPERATING CHARACTERISTICS

The practical results of the various methods and circuits which have been described are given in Figs. 12–16, which show some of the main measurable operating characteristics of the SR process.

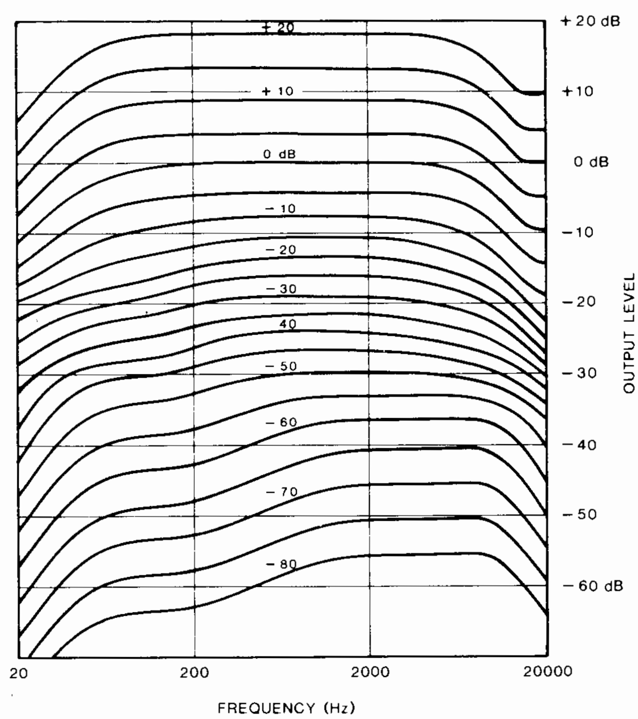

Fig. 12. Encoder characteristics—single tone.

4.1 Dynamic (Compression) Action for Steady-State Dominant Signals

Referring to the single-tone encoder curves shown in Fig. 12, several features may be noted.

4.1.1 Low Frequencies

Dynamic action occurs in the range from about −48 dB to −5 dB (with respect to reference level). That is, there is no action (but full boosting) in the lower 35–40 dB of the dynamic range. Similarly there is no action in the top 25 dB of the total dynamic range. A linear dynamic characteristic prevails in these two regions (a bilinear characteristic).

4.1.2 High Frequencies

Dynamic action occurs in the range from about −62 dB to −5 dB. That is, there is no action in the lower 20–25 dB (but full boosting), or the top 25 dB of the dynamic range (a linear dynamic characteristic in these regions).

In the intermediate level regions of dynamic action the effects of the multilevel stages are joined together to create a compression ratio of about 2:1.

Referring to the high-level low- and high-frequency portions of the curves, the effective antisaturation of the system can be seen, with the combined effects of the spectral skewing and antisaturation networks, the SR stages, and about 1 dB of wide-band level compensation built into the coefficients of the stage signal combiner (Fig. 1). The overall result is an antisaturation effect of about 2 dB at 5 kHz, 6 dB at 10 kHz, and 10 dB at 25 Hz and 15 kHz. At high frequencies this amount of antisaturation significantly reduces distortion, reduces signal compression effects, and, with tape recording, improves the long-term stability of the recording. The high-frequency improvements are especially significant with 35-μs CCIR recordings. The antisaturation effect at low frequencies usefully counteracts tape overload, particularly with 3180-μs NAB recordings.

4.2 Quiescent (Subthreshold) Signal Characteristic

The very low level, or subthreshold, characteristics of the SR process are shown in Fig. 13. The general shape of this characteristic was determined in a way that takes good advantage of the properties of human hearing. First, there is less of a noise generation and perception problem at moderately low frequencies (such as 200 Hz) than at moderately high frequencies (such as 3 kHz). Therefore two low-frequency stages are employed, but three high-frequency stages are used.

Second, at very low and very high frequencies even less noise reduction is needed (below 50 Hz and above 10 kHz). Strong spectral skewing actions can therefore be used in these regions, resulting in accurate decoding even when the recording medium has response irregularities. In addition the spectral skewing networks provide for good immunity to high- and low-frequency interference (supersonic audio components, tape recorder bias; subsonic noise components arising from wind, traffic, or other rumble sources).

Note that the overall shape of the low-level SR decode characteristic resembles the low-level Fletcher–Munson and Robinson–Dadson curves; the encode characteristic resembles the subsequently derived CCIR noise-weighting curve.

Thus the SR system is designed to reduce only those noises that can be heard. The prevention of action in inaudible signal regions promotes accuracy in the audible region.

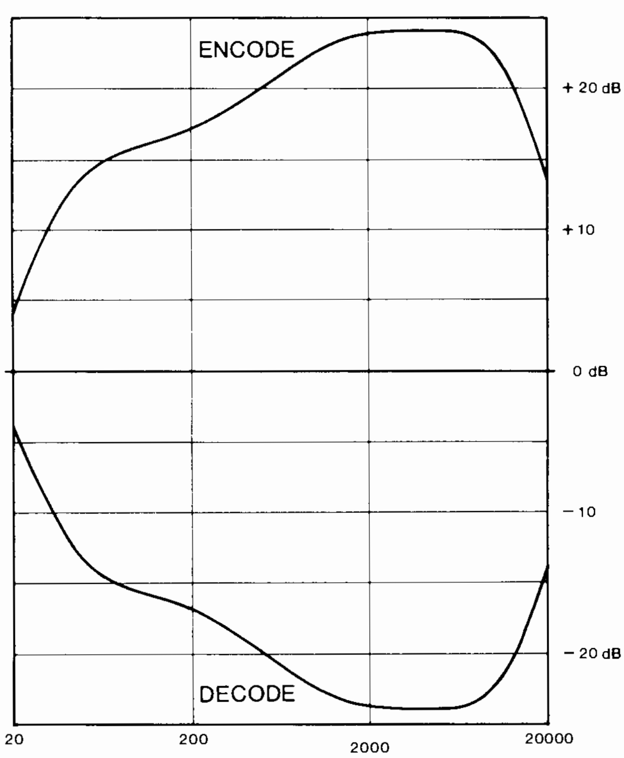

Fig. 13. Low-level (subthreshold) encoding and decoding characteristics.

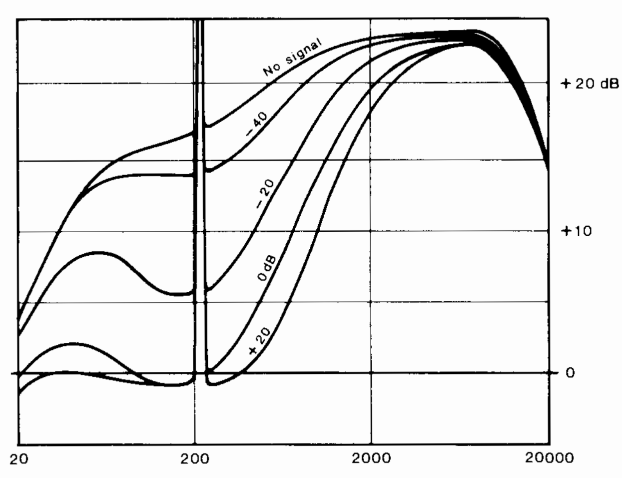

Fig. 14. Low-level encoding characteristic in the presence of a 200-Hz signal at the levels indicated.

4.3 Treatment of Nondominant Signal Components

The behavior of the system with low-level nondominant signal components in the presence of higher level dominant signal components can be simulated by the use of probe tones. Such a representation is significant because it is an indicator of the noise reduction effect achieved with signals. Refer to the curves shown in Figs. 14–16, which were obtained by adding a swept frequency probe tone at levels between −60 dB and −80 dB into the encoder input signal and detecting the tone at the output with a tracking wave analyzer.

Toward the two spectrum ends, nondominant signal components are boosted more than the dominant signal by high- and low-frequency sliding-band actions. If there are two dominant signals, a fixed-band compression effect prevails for the nondominant signal components between the frequencies of the dominant signal components.

Thus nondominant signal components are boosted by an amount at least equal to the amount of boost of the dominant signal. The boosting of the nondominant signals is maintained toward the spectrum ends, even though the level of the dominant signal is relatively high (in the range −30 dB to +20 dB). This boosting action spectrally tracks the dominant signal frequency or frequencies.

It is advantageous to have a steeply rising boosting effect away from the frequency of the dominant signal component. In this connection the SR circuit profits from the steepness enhancing effect of cascaded stages. The low frequencies have two stages of steepness compounding; the use of three stages at high frequencies further improves the effect. These characteristics are particularly evident in the high-level areas of the probe tone curves.

The curves show that the encoder circuit tends toward keeping all low-level signal components boosted at all times. Only those components above the threshold are subject to a reduction of boosting. With regard to the overall encode–decode system, the advantages of this type of characteristic are 1) a powerful noise reduction effect in the presence of signals, and 2) a relative tolerance of level and frequency response errors in the channel between encoder and decoder.

4.4 Audible Results

As with the investigation of other psychoacoustic devices such as loudspeakers, measurements of the SR process can provide only a partial characterization. The rest must be obtained by detailed listening, using a wide range of source material and a variety of recording and transmission media.

Generally, the audible encoding effect of the system is to create a dense, bright sounding signal (as sent to the recording channel), but with little or no apparent dynamic action. Harmonics, overtones, and small-scale components of the sound, including noise, are all enhanced. At high signal levels the antisaturation characteristics cause a high- and low-frequency audible dulling of the encoded signal; when applied to the recording channel, this treatment results in a significant reduction of recording distortion.

The overall audible encoding and decoding effect of the SR process is simply to create a clean and accurate sounding replica of the input signal. Tape bias noise and modulation noise are significantly reduced. Also, a reduction of intermodulation distortion is achieved by the low-frequency noise reduction capabilities of the process, as well as by the effects of the antisaturation characteristics used during encoding.

Furthermore, the decoding portion of the system reduces harmonic distortion generated by the recording channel. Steady-state third-harmonic distortion is typically reduced to less than one-half, fifth-harmonic distortion to less than one-quarter; higher order harmonics are even further reduced. Thus, especially if the recording medium has a hard clipping characteristic, the subjective cleanliness of the signal at high recorded levels is significantly improved.

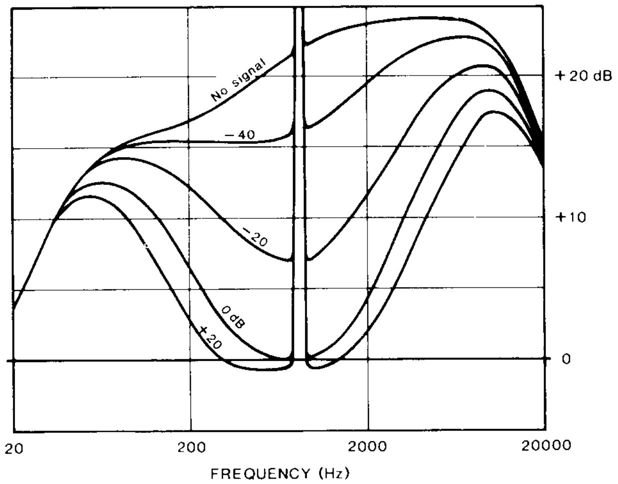

Fig. 15. Low-level encoding characteristic in the presence of an 800-Hz signal at the levels indicated.

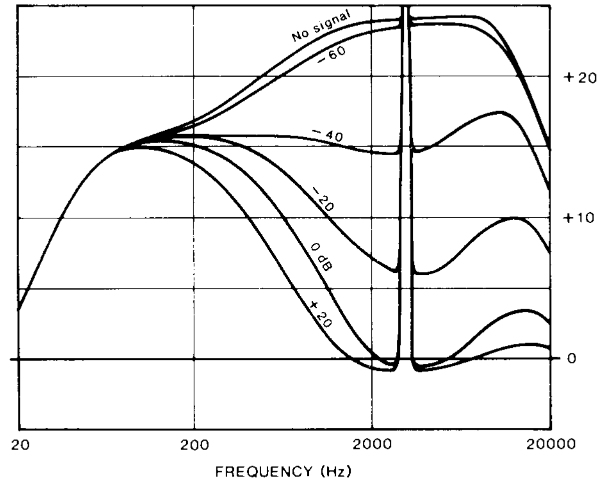

Fig. 16. Low-level encoding characteristic in the presence of a 3-kHz signal at the levels indicated.

5 CALIBRATION

The spectral recording calibration procedures are conceptually similar to those of the A-type, B-type, and C-type systems. That is, signal levels in the decoder circuit ideally should match those in the encoder circuit, even though the SR process has been designed to be more tolerant of gain and frequency response errors than these previous systems. For tape interchange standardization it is also preferable if, at least within a given organization, the reference level of the encoder and decoder corresponds to a known and fixed flux level. Whether or not a standardized flux is used for this, the matching of the decoder to the encoder is accomplished by a calibration signal generated in the encoder and recorded on the tape; this allows the tape replay gain to be set correctly, using the meter in the decoder unit.

Most problems in the studio use of noise reduction, and indeed analog recording in general, can be traced to incorrect level settings and/or frequency response errors in the recorder. This may be because checking these factors is a time-consuming and boring process. A faster and more interesting method of accomplishing these checks would be more likely to produce reliable and consistent results. For this reason, practical embodiments of the SR process include pink-noise generators which are used for both level and frequency response calibration, instead of single-tone sine-wave oscillators. For identification, the pink noise is interrupted with 20-ms "nicks" every 2 s. During recording this signal is fed to the tape at a level of 15 dB below reference level, a level low enough not to cause saturation problems with low-speed tape recording or highly equalized transmission channels.

During playback the tape signal is automatically alternated with internally generated reference pink noise (uninterrupted) in 4-s segments (8-s total cycle time) and passed to the monitor output. An audible comparison can thus be made between the reference pink noise and the calibration noise coming from the tape. Any discrepancies in level and/or spectral balance are immediately noticeable and can be corrected or at least taken note of. If desired, the signal can also be fed to a spectrum analyzer.

In using the new calibration method it is important to be able to tell when the 4-s tape segments are being passed to the monitor and when the signal heard is from the reference pink-noise generator. Differentiation of the tape segments from the reference segments is accomplished in two ways. First, the reference segments are 4 s of continuous pink noise, and the tape segments begin with a nick, have a nick in the middle, and end with a nick; this time sequence is easily identified with a little practice. Second, colored lights identify the two different signals.

The new calibration facility gives the recording and production personnel a useful control of the recording process. At any time a check of the recorder can be made; the result can be heard immediately and conclusions drawn about whether adjustments might be necessary.

With tape and signal interchanges it is possible to tell quickly whether there is any error or misunderstanding about levels, equalization, azimuth, and the like. If the original recording of calibration noise stays with the tape, the quality of the ultimate playback, even after copying, can be retained. Thus the comparison function serves to ensure that the recorder and spectral recording process provide on a routine basis the signal quality and reliability of which they are capable.

6 CONCLUSION

A new professional recording format, designated spectral recording, has been described. The objective of the new encoding and decoding process is to record and reproduce audio signals with a high degree of audible signal purity.

The system employs a dual-path multilevel, staggered action arrangement of two low-frequency compressor stages and three high-frequency compressor stages, each with a fixed band and a sliding band. The outputs of the bands are combined in a unique way, called action substitution, which results in an unusually responsive treatment of the signal with respect to both frequency and level; a technique referred to as modulation control augments the spectral tracking abilities of the system. Spectral skewing contributes to a tolerance of channel errors, and the employment of both high- and low-frequency antisaturation techniques results in a significantly improved channel overload characteristic.

7 ACKNOWLEDGMENT

The author is indebted to Douglas Greenfield for the printed-circuit layouts, mechanical design, and general assistance with the first practical embodiments of the new system. Thanks are also due to many other individuals, particularly to Henry Bottino for the design of testers, to Brad Teague for preproduction testing, and to Michael Smith, Christopher James, and Martin Fried for production engineering.

8 REFERENCES

[1] R. M. Dolby, "An Audio Noise Reduction System," *J. Audio Eng. Soc.*, vol. 15, pp. 383–388 (1967 Oct.).

[2] R. M. Dolby, "A 20 dB audio Noise Reduction System for Consumer Applications," *J. Audio Eng. Soc.*, vol. 31, pp. 98–113 (1983 Mar.).

[3] R. M. Dolby, "A Noise Reduction System for Consumer Applications," presented at the 39th Convention of the Audio Engineering Society, *J. Audio Eng. Soc.* (*Abstracts*), vol. 18, p. 704 (1970 Dec.).

[4] R. Berkovitz and K. J. Gundry, "Dolby B-Type Noise Reduction System," *Audio*, pp. 15–16, 33–36 (1973 Sept., Oct.).

THE AUTHOR

Ray Dolby was born in Portland, Oregon, in 1933, and received a B.S. degree in electrical engineering from Stanford University in 1957. From 1949–52, he worked on various audio and instrumentation projects at Ampex Corporation, and from 1952–57 he was mainly responsible for the development of the electronic aspects of the Ampex video tape recording system. After he was awarded a Marshall Scholarship, followed by a National Science Foundation graduate fellowship, he left Ampex in 1957 for further study at Cambridge University in England where he received a Ph.D. degree in physics in 1961 and was elected a Fellow of Pembroke College (Honorary Fellow, 1983). During his last year at Cambridge, he was also a consultant to the United Kingdom Atomic Energy Authority.

In 1963, he took up a two-year appointment as a United Nations adviser in India and returned to England in 1965 to establish Dolby Laboratories in London. Since 1976 he has lived in San Francisco, where his company has established further offices, laboratories, and manufacturing facilities.

Dr. Dolby holds a number of patents and has written papers on video tape recording, long wavelength X-ray analysis, and noise reduction. He is a fellow and past-president of the AES and a recipient of its Silver Medal Award. He is also a fellow of the British Kinematograph, Sound and Television Society, and of the SMPTE, and a recipient of its Samuel L. Warner Memorial Award, the Alexander M. Poniatoff Gold Medal, and the Progress Medal. In 1979 he and his colleagues received the Scientific and Engineering Award of the Academy of Motion Picture Arts and Sciences. In 1986 he was awarded the British title of Officer of the Most Excellent Order of the British Empire (OBE).