1505 (B-1)

AN ABOVE THRESHOLD COMPRESSOR WITH ONE CONTROL

Leslie B. Tyler
dbx, Inc.
Newton, Massachusetts

Presented at
the 63rd Convention
May 15 through 18, 1979
Los Angeles

AN ABOVE THRESHOLD COMPRESSOR WITH ONE CONTROL

Leslie B. Tyler
dbx, Inc.
Newton, Massachusetts

ABSTRACT

An unconventional above threshold compressor is described. The compressor controls the static gain of the system by comparing its threshold to the desired output level and adding gain as necessary. The compression curve employs a soft threshold, allowing continuous variation in compression ratio with input signal. A single threshold control fulfills the functions previously ascribed to threshold, gain and compression ratio controls.

INTRODUCTION

There are many situations in recording speech or music which benefit from dynamic range modification. The dynamic range of live music may be greater than 80 dB, while recorded music must be stored and replayed through various media limited to less dynamic range than that present in the original. With modern recording and signal processing techniques, it is possible to preserve the dynamics of live music as far as the listening room, [1,2] but often other constraints make it necessary to reduce the dynamic range of a recording. Not the least of these constraints is often the background noise level in the listening environment, compared to the maximum permissible playback level.

Many different types of dynamic range reducers have been invented. The prototype dynamic range reducer is the recording or mixing engineer exercising control over level via a fader.[3] The engineer as a compressor has several points in his favor. For example, he is often able to anticipate crescendos and decrescendos before they occur, preventing overshoot. His time constants, compression ratio and threshold of compression are program dependent and intelligently controlled. Unfortunately, when many signals are to be independently and simultaneously compressed, it can be inconvenient to hire multiple engineers to function as "compressors", and asking one person to manually compress several tracks at once is inviting disaster.

Accordingly, electronic replacements for the engineer/compressor have been designed. In general, they fall somewhat short of offering the complexity of control which the prototype human compressor can exercise, but make up in convenience for what they lack in sophistication. It is the purpose of this paper to describe the invention of a new type of compressor which combines conceptual and operational simplicity with a sophisticated control algorithm.

TYPES OF COMPRESSORS

Compressors may be roughly divided into two basic types, those which are linear, and those which have one or more thresholds. Linear compressors (such as the dbx Model 118) are those which uniformly change the dynamics of a signal, regardless of the signal level.[4] For a 2:1 linear compressor, an input level change of 1 dB produces an output level change of ½ dB, no matter what the actual input level. Higher compression ratios produce less output change (1/c, where c is the compression ratio) for a 1 dB input change (see figure 1). Such a compressor is relatively easy to operate, since varying one parameter, the compression ratio, allows the operator complete control over the sound of the device.

Above threshold compressors (such as the dbx Model 160) pass a signal without modification (except, as we shall see, for a possible change in level) below the threshold, but apply various amounts of compression when the input exceeds the threshold.[5] For a 2:1 above threshold compressor, with the threshold at 0 dB, a 1 dB input level change produces 1 dB output change below 0 dB, but only ½ dB output change above 0 dB (see figure 2). Since the dynamics of the output signal will depend on the input level relative to the threshold as well as on the compression ratio, somewhat more care must be exercised in operating an above threshold compressor than in operating a linear one.

The above threshold compressor offers a somewhat closer approximation to the operation of the prototype human compressor than does its simpler, linear counterpart. The human compressor is likely to reduce gain when overloads are threatened, but will not increase gain during soft passages at the expense of an increase in noise. The above threshold compressor will mimic this action (no gain below 0 dB in Figure 2), while the linear compressor will blindly increase gain in quiet passages (10 dB gain at -20 dB input in Figure 1).

Compressors with multiple thresholds have also been invented. These devices include compressor/limiters with separate compression and limiting thresholds (such as the Allison Gain Brain), compressor/expanders with a high compression threshold and low expansion thresholds (for example, the EMT 156), as well as other variations. However, the majority of compressors designed for use in recording are of the single above threshold type.[3]

SETTING UP THE ABOVE THRESHOLD COMPRESSOR

The two primary controls for an above threshold compressor are the threshold and compression ratio controls (an input level control with a fixed threshold may be substituted for the variable threshold control). Their settings will be determined by the input dynamic range, the desired output dynamic range, and the characteristics of the sound to be compressed. In multi-track recording, individual instruments are often compressed separately, and the set up requirements here might be expected to differ from those for a compressor applied to the final, mixed product.

For a single track, a compressor would very likely be used to limit dynamic range to manageable proportions in order to record the sound without distortion or excessive noise. In this usage, the point would be to compress as unobtrusively as possible while still providing overload protection. Since unobtrusive compression is best achieved by manipulating the sound as little as possible, while overload protection is best achieved by very high compression ratios (limiting or drastic manipulation of the sound), this adjustment of a compressor becomes a tradeoff between sufficient overload prevention and excessive audibility of the gain change.

In many complex mixing situations, certain tracks must remain dominant in the final product (an example would be the vocal, bass, and drum tracks for most pop recordings). However, a musician who is really getting into the performance may not provide a stable level into his microphone or direct pickup at all times. In such a case, compression may be used specifically for its audible effect of stabilizing level. The compressed track should ideally sound as if no compression had been applied, but rather as if the original performer had skillfully controlled his own dynamic range. In other words, while some effect of compression must be audible, the compressed track should retain the natural quality of the original.

When a mix (or perhaps an original two track recording) is to be transferred to disc, it may be necessary to limit dynamic range in order to properly cut the record. This compression should again be as unobstrusive as possible, in order to retain most of the (presumably skillfully executed) balances of the recording. If the original balances were poor, some compression (or expansion!) and equalization might be used to set things right. For either case, the compression must not make the music sound unnatural.

All of the preceding scenarios in which compression might be used will involve varying the compression ratio and the threshold point in order to change the amount and type of compression applied. Each situation involves the judgement of the engineer and producer as to the proper amount of compression to be applied, where it will be applied, and how much audible effect it should cause. All these judgements must be made by listening to the compressed product, and comparing it to the original.

Let us suppose a vocal track is a candidate for some compression. The singer begins to rehearse his part, as the engineer fiddles with the compressor assigned to the vocal track. Both engineer and producer listen to the effect of moving the threshold and changing the compression ratio. (For multiple threshold compressors, this process is complicated by the additional adjustments to be made.) Once a seemingly appropriate setting has been found, an A/B comparison is made of the difference between the original and the compressed version. To everyone's dismay, the original is much louder than the compressed signal.

The observed level difference is caused by the fact that the compressor does its job by reducing gain (see figure 2). In order to make meaningful A/B comparisons, some means of statically restoring this gain must be provided. In many compressors (dbx 160 for example) static gain may be added or subtracted with an output gain control. The way in which such a control raises gain is shown in Figure 3 for a compressor with 0 dB threshold and 2:1 compression ratio. The output gain control allows the recording to be made at a level independent of the compression settings, and also facilitates meaningful A/B comparisons.

Now that we have provided the compressor with a gain control, let us return to the example above. The engineer increases the static gain of the compressor to match the input and output levels, and proceeds with the A/B comparison. If it seems that a bit more fiddling is necessary - the sound is still not quite right - he will return to changing the compression ratio and threshold settings, while listening to the results. But, every change in ratio or threshold produces two effects. Not only does the amount of compression change, so does the output level. The engineer may find himself moving the gain control every time he adjusts the other two controls. The problem here is that the secondary effects of changing threshold or ratio settings confuse the engineer's judgement at a time when all his attention should be focused on the primary effects of his fiddling.

A SOLUTION TO THE INTERACTION PROBLEM

Consider a specific above threshold compressor, which has characteristics as shown in Figure 4. The unit's characteristic is fixed at an infinite compression ratio and 0 dB threshold setting. This compressor has an interesting property: its output level is 0 dB so long as the input level is above 0 dB. Since we are proposing that the machine is to be used as a compressor, it seems safe to assume that input levels will be sufficiently high to cause compression, so we have a unit whose output level is nearly independent of the input level.

We could generalize that, for an infinite above threshold compressor without gain controls, the output level will tend to stay at the threshold level, so long as the input is above the threshold. Because of this simple, linear relationship, gain might be added inversely as the threshold is changed, and the output level would tend to stay put. This case is illustrated in Figure 5.

The compressor of Figure 5 allows change of the threshold position without change in output level. Different output levels (when desired) could be accomodated with an additional gain control, which modifies all the curves upward or downward by the same amount. Such a control might be calibrated in terms of nominal output level.

While the compressor described in Figure 5 eliminates the interaction between threshold position and output level, it does so at the expense of a drastic limitation in available compression ratios. Unfortunately for this device, infinite compression does not sound very good.[6] It would hardly be considered for the wide variety of applications for which the more general above threshold compressor of Figure 2 is useful. However, it is possible to change the characteristics of the curve in Figure 4 to obtain musically useful results, and still be able to vary gain as shown in Figure 5, as we shall see.

AN ALTERNATE ABOVE THRESHOLD CHARACTERISTIC

Let us consider in detail the action of the compressor shown in Figure 2 as the input signal increases in level from -10 dB to +10 dB. For the purposes of this discussion, examine the 10:1 compression characteristic illustrated. For the first 10 dB of increase, the output follows the input: no alterations in the dynamics are made. As the input passes 0 dB, the output dynamic range is suddenly reduced by a factor of 10:1. The effect of the sudden cut in gain would be quite noticeable, especially if the level increase occured at a point of increasing excitement in the music. It is just such a noticeable change in the character of the sound which makes high compression ratios undesirable from a musical point of view. Yet, high compression ratios can prevent drastic (and even more undesirable) overloads of succeeding elements in the signal path.

One method of reducing the audibility of the sudden gain change would be to make the threshold less abrupt. Such a situation is a natural outgrowth of the more complex compressor designs, which encompass both a compressor and a limiter with separate thresholds. Theoretically, an infinite series of thresholds could be constructed, giving rise to a compression curve which is smooth, without the sharp break point (threshold) of other compressors. This multiplicity of thresholds would be confined to a region of the compression curve, and we might reasonably regard the center of this region as the new compression threshold. Figure 6 illustrates a possible shape for the new compressor characteristic. Because the new curve allows the signal to go "over" the threshold area more "easily" (smoothly), we have nicknamed this characteristic the "over easy" curve.

That the compression curve of Figure 6 might sound better than the traditional one of Figure 2, for the same assymptotic compression ratio, appeared relatively easy to demonstrate. We built a prototype compressor according to the block diagram in Figure 7, which resulted in a transition region width of approximately 15 dB, as shown in Figure 6. The 15 dB figure was purely a guess at what might be required, affording us a starting point and no more. The dbx RMS detector (U.S. Pat. #3681618) was chosen for the similarity of its response characteristics to those of the human ear, and the time constants of the successful dbx 160 series of compressor/limiters were retained for the new design. (Among other considerations, keeping the same time constants insured that we were comparing apples to apples in our A/B tests!) The dbx VCA design (U.S. Pat. #3714462) was chosen for its wide dynamic range, linear/log control voltage characteristic, and ready availability.[7] The "over easy" circuit was designed to have an input vs. output characteristic as shown in Figure 8. The proper locations of the various controls was relatively straightforward and should be

obvious from the block diagram.

Equipped with examples of both traditional and "over easy" compressors, we began a series of listening comparisons. We found that the modified characteristic did indeed cause music to sound less "compressed" than the traditional characteristic, for the same amount of gain reduction. (Calibrated gain reduction meters were included in both compressors.) Listeners described the sound as being less "held back", less "constrained", or more "open" in the new compressor. The source material for these tests consisted of dubs of individual tracks from original 16 track master tapes, as well as mixed material. The tracks included vocals, snare drums, kick drums, drum kits, bass guitars (direct and live mike), guitars (electric - direct and live mike - and acoustic), pianos, violins and maracas. Considerable effort was made to secure uncompressed, wide dynamic range material so that we would be listening to the compressor under test, not an unknown compressor used in recording. All the tapes were dbx encoded copies of encoded originals.

While reaction to the sound of the compressor was generally favorable, we felt that there was still room for improvement. Our efforts were aimed at making the compression as unnoticeable as possible. We were striving to obtain at least 10 dB of compression without listener awareness of the dynamic modifications taking place. By extending the transition region width to about 30 dB, this goal could be met with almost all types of source material. The final compression curve on which we settled is shown in Figure 9, for an assymptotic compression ratio of ∞:1.

DETAILS OF THE "OVER EASY" CURVE

Several facets of the "over easy" characteristic are worth examining. The curve must not always approach infinite compression, but can be realized in more or less drastic final ratios. A control for varying the final compression ratio has already been shown in the block diagram of Figure 7. In fact, a version of the "over easy" compressor including this feature was later incorporated in the dbx Model 165 compressor/limiter. If, however, the curve were to be fixed at the infinite compression assymptote shown in Figure 9, every compression ratio from 1:1 to ∞:1 would be available without the need for a traditional compression ratio control. The compression ratio would instead be a function of the signal level relative to the threshold position.

The curve of Figure 9 is particularly desireable because it provides gentle compression (<3:1) for inputs below the threshold, but limits gross overloads to no more than ≈5 dB higher output levels than those outputs obtained with inputs at the threshold level. In other words, if normal program material appears at about the threshold level, a larger input overload will only cause a 5 dB increase in the output level, no matter how much the input increases (within the linear operating range of the compressor, of course).

A single threshold control could replace both the traditional threshold and compression ratio controls in the new compressor design. Varying the threshold position obviously varies where the threshold occurs relative to the input level, but also varies the instantaneous compression ratio for any particular input level. Note that the definition of threshold is less obvious for Figure 9 than for Figure 2, but nonetheless we can agree to call some point on Figure 9 the threshold reference. We have arbitrarily assigned the "over easy" threshold to the point at which 6 dB gain reduction occurs.

APPLYING THE "OVER EASY" CURVE TO THE INTERACTION PROBLEM

Because the "over easy" curve is assymptotic to infinite compression, the output level is limited by the threshold position. The output level can be no more than the threshold level plus 5 dB, as shown in Figure 9. This relationship is similar to the one which led to the curves of Figure 5, in which the idea of varying static gain with threshold was first presented. Gain may be added to the output of the "over easy" compressor just as it was in the compressor of Figure 5, generating a family of curves as shown in Figure 10.

At this point in the compressor's development, we were converging towards a single control compressor, in which the threshold control sets not only the threshold, but also compression ratio and output gain. In such a device, the maximum output level would be limited to +5 dB, relative to the internal calibration 0 dB point. For real world devices, 0 dB is a variable from system to system. An auxiliary control could be added to calibrate the compressor's internal 0 dB reference to the local standard. Fortunately for our goal of simplification, this control need be adjusted only once, when the compressor is installed. This adjustment would then tend to force the output level to remain at the system reference level.

The range of threshold adjustment should be large enough to allow infinite compression of fairly low level signals, and almost no compression of nominal level signals. The first criterion involves pushing the threshold far below the 0 dB reference, necessitating large amounts of static gain, and the second requires pushing the threshold above the 0 dB reference, which in turn calls for static gain reduction. The result would be a control which calls for positive gain at one end of its range, and negative gain at the other.

The realization that we would need to add both positive and negative gain nearly stopped the whole project, because human engineering of the compressor seemed to necessitate that the device revert to unity gain (straight wire) operation with a high threshold (very little compression) condition. At first we saw no way around the problem, for one requirement was unity gain, and yet another was negative gain. The solution was surprisingly simple, once we saw it: truncate the gain make-up process when it becomes gain reduction. At some point in the threshold control's motion, the threshold would become greater then 5 dB below the nominal output level. If the gain make-up were stopped at this point, no gain reduction would be implemented.

Again, with soldering iron in hand, we rushed to build a prototype of this latest conception, only to find on listening to it that the gain make-up action was not very smooth. If the input ranged around 0 dB, and the system were calibrated properly, moving the threshold continuously downwards from a high level caused at first a gain decrease, then an increase, and then a leveling out at the 0 dB level, along with increasing amounts of compression as the threshold was lowered. The changes in average level were less than 6 dB, but were noticeable and annoying. The cause of the problem was that we attempted to match an "over easy" curve, (in dynamic gain reduction - compression) with a sharp threshold curve (in static gain addition - make-up). The solution to this last problem was to make the gain make-up characteristics conform to the "over easy" curve.

CONCLUSION

The final realization of the compressor is embodied in the dbx Model 163 compressor/limiter. Figure 11 shows a block diagram of the unit. A single slide potentiometer serves as the threshold/compression/output level control. Its labeling is in keeping with the compressor's simplicity of operation: the direction of lower threshold is labeled "MORE". A rear panel switch allows selection of nominal output levels of +4 dBm, -10 dBm, or variable via a screwdriver adjust pot. Dynamic gain reduction is displayed via a 12 section LED display on the front panel.

The family of curves shown in Figure 12 illustrates the action of the front panel control for a -10 dBm output level setting. Moving the threshold control towards "MORE" includes more of the input dynamic range within the compression region, and increases the static gain of the system. At high threshold settings (away from "MORE"), gain is unity for inputs far below the threshold, and only decreases as the threshold is approached. The output level may therefore be above the nominal level for the combination of higher than nominal inputs, and high threshold settings. Moving the threshold downward in the presence of higher than nominal inputs will smoothly reduce the output level towards the nominal. Conversely, moving the threshold downward in the presence of lower than nominal inputs will smoothly increase the output level towards the nominal.

In normal operation of the compressor, an input signal is trimmed via the mic preamplifier gain adjustments approximately to the system nominal level. Varying the compression control setting will then have little effect on the compressor's output level. However, as the control is pushed towards "MORE", more compression and less dynamic range will be the result. The action of the control is smooth and intuitive. As the threshold is reduced, the compression ratio applied to any given input level is increased proportionately. By allowing the compressor to automatically compute the proper combination of compression ratios, threshold location and gain for various desired amounts of compression, the engineer and producer are freed from thinking about these details, and are permitted to fully concentrate on the quality of the sound.

The author would like to express his appreciation to the many people at dbx, Inc. who supported the project through both good and bad days. Special thanks are due to several people outside the company who helped at various stages of the project, including George Augspurger, whose remarks to Larry Blakely germinated the idea of the continuous threshold region, and Bill Riesman, who supplied the uncompressed evaluation material through the Northern Studio.

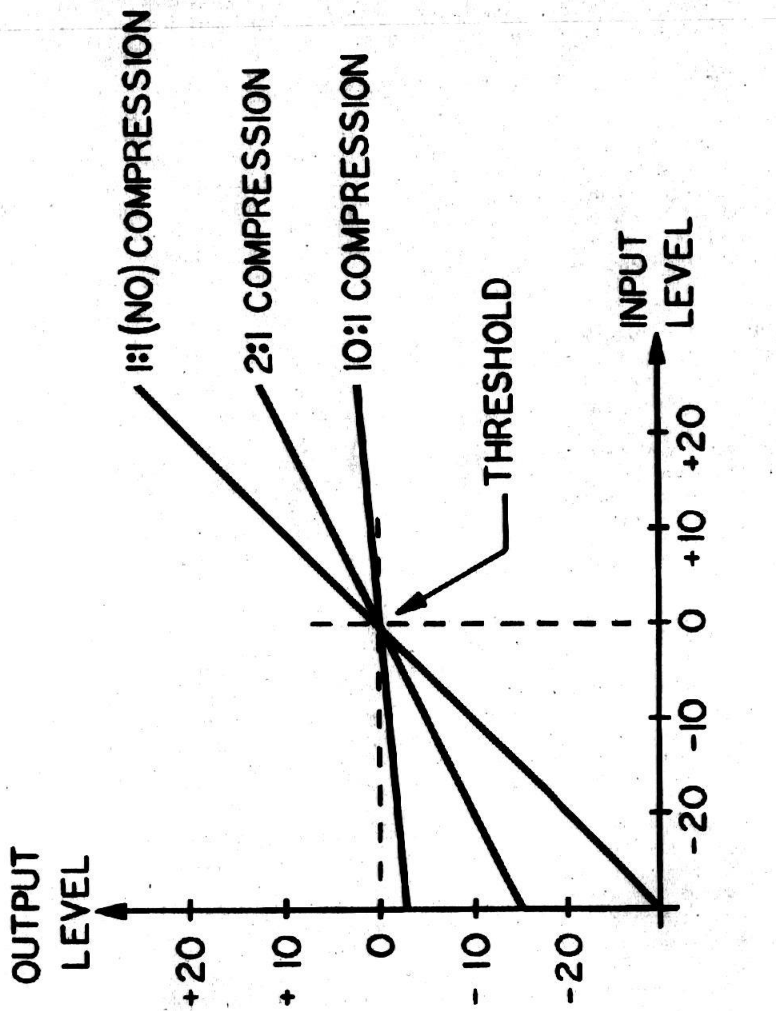

LINEAR COMPRESSOR CHARACTERISTIC CURVES

FIG-1

REFERENCES

1.) Blackmer, David E., "A Wide Dynamic Range Noise Reduction System", DB, The Sound Engineering Magazine, Vol. 6, No. 2 (1972)

2.) Blakely, Larry, "Using Noise Reduction to Reduce Disc Surface Noise", Recording Engineer/Producer, Vol. 4, No. 6, (1973)

3.) Eargle, John, Sound Recording, Van Nostrand Reinhold Co., New York, 1976, pp. 242-6

4.) dbx, Inc. "dbx Model 118 Instruction Manual", dbx, Inc. Newton, MA (1979)

5.) dbx, Inc. "dbx Model 160/161 Instruction Manual", dbx, Inc., Newton, MA (1979)

6.) Blesser, Barry A., "Audio Dynamic Range Compression For Minimum Percieved Distortion", IEEE Transactions on Audio and Electroacoustics, Vol. AU-17, No. 1 (1969) p. 32

7.) dbx, Inc., "Voltage Controlled Amplifiers (202)", dbx, Inc. Newton, MA (1979)

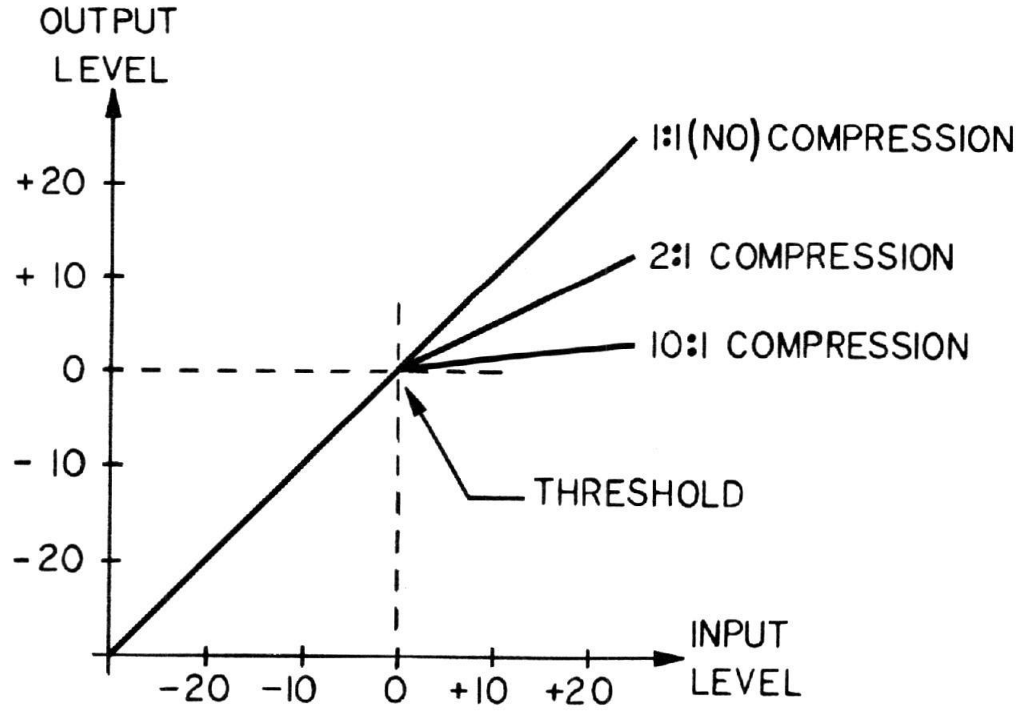

ABOVE THRESHOLD CHARACTERISTIC CURVES

FIG.2

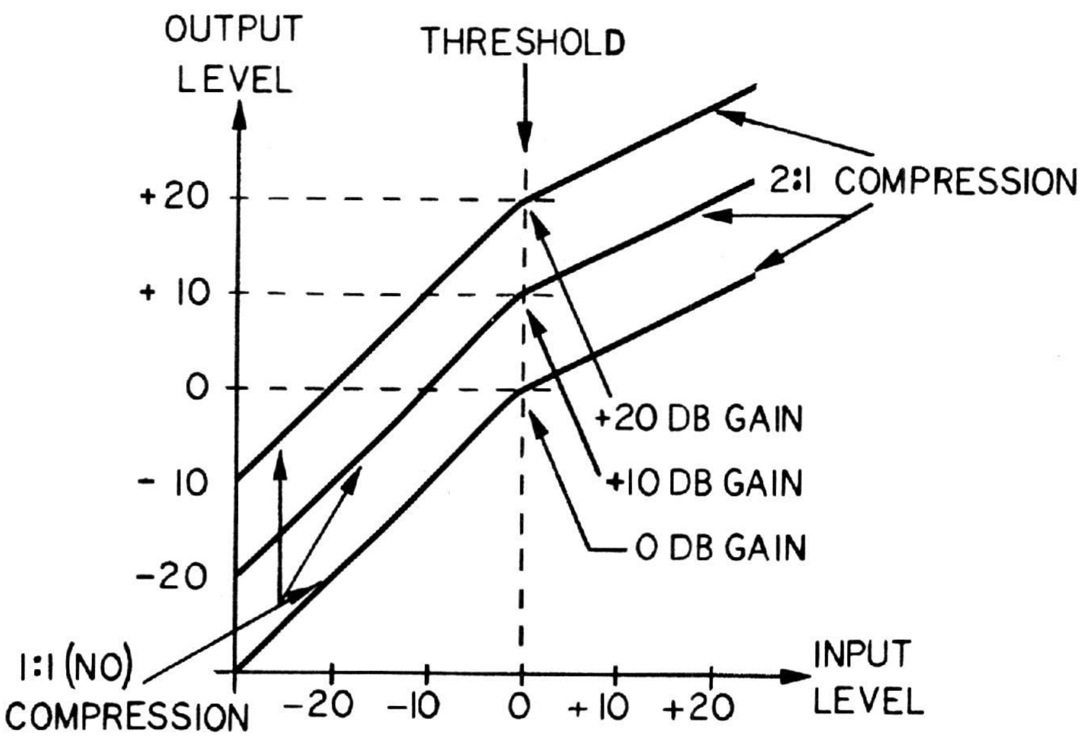

ABOVE THRESHOLD CURVES SHOWING GAIN CONTROL

FIG.3

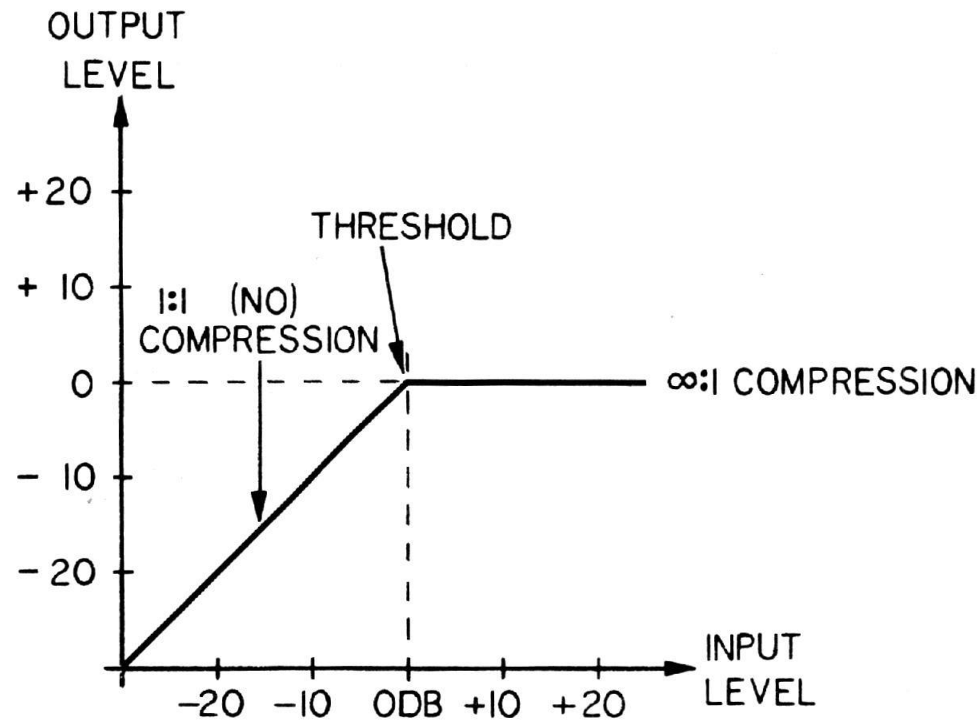

INFINITE ABOVE THRESHOLD COMPRESSOR CHARACTERISTIC

FIG. 4

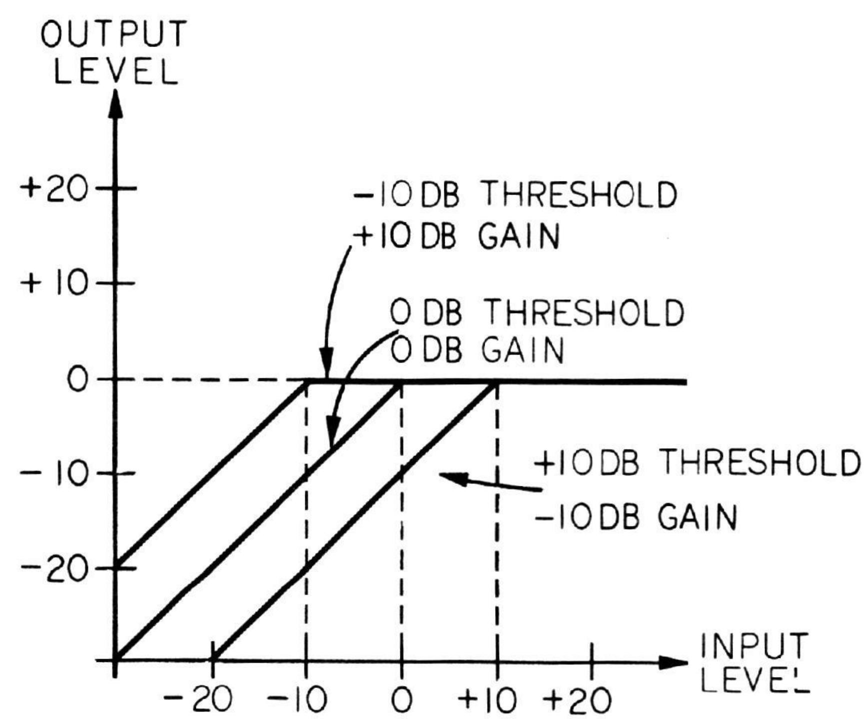

INFINITE ABOVE THRESHOLD COMPRESSOR
WITH GAIN ≃ -THRESHOLD

FIG. 5

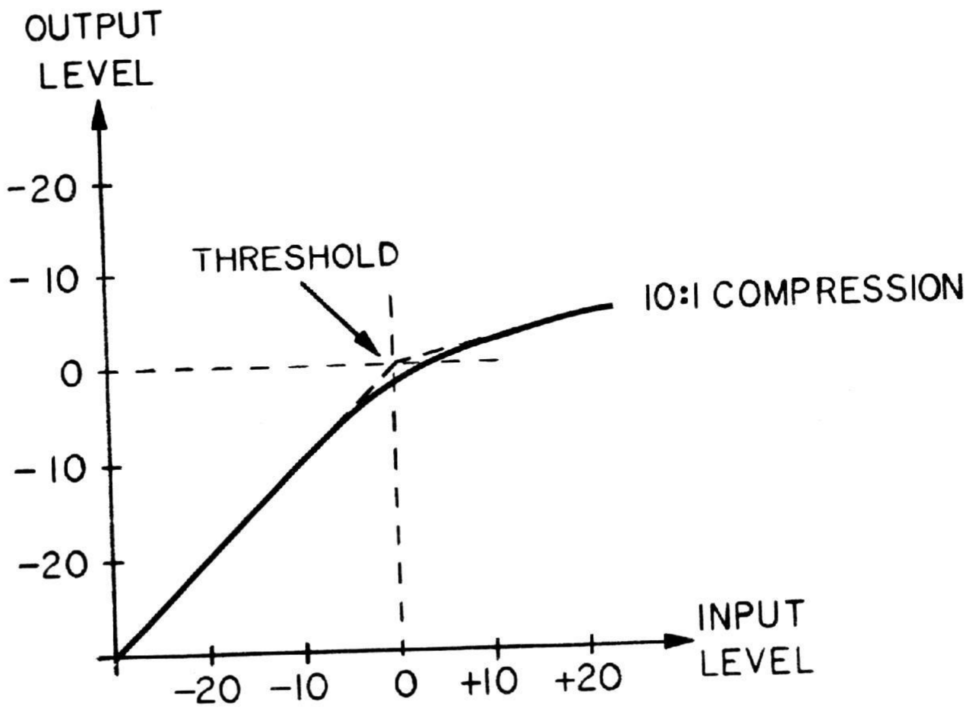

OVER EASY COMPRESSOR
SET TO 10:1 COMPRESSION, 0 DB THRESHOLD

FIG.6

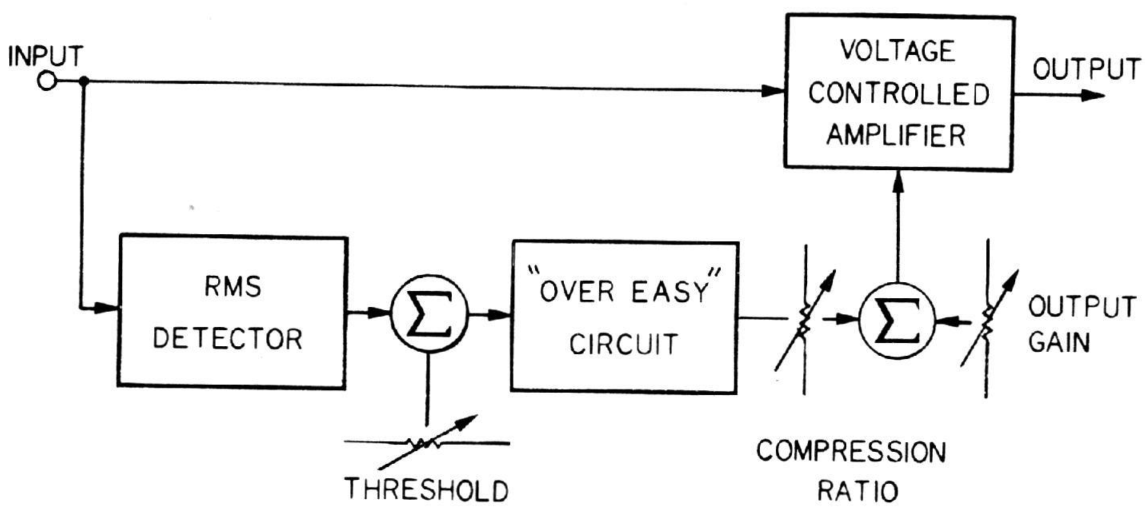

BLOCK DIAGRAM OF PROTOTYPE OVER EASY COMPRESSOR

FIG.7

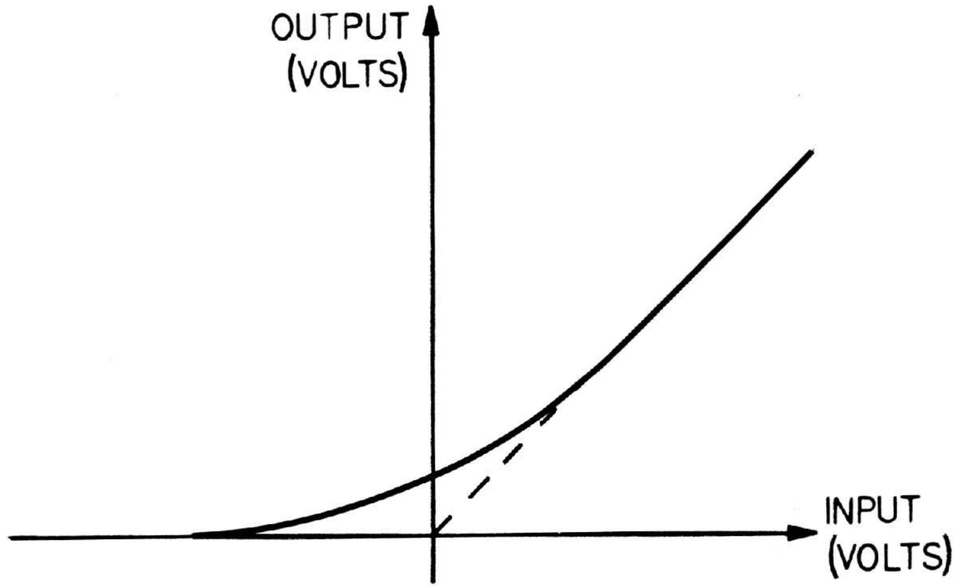

TRANSFER CHARACTERISTIC OF OVER EASY CIRCUIT

FIG.8

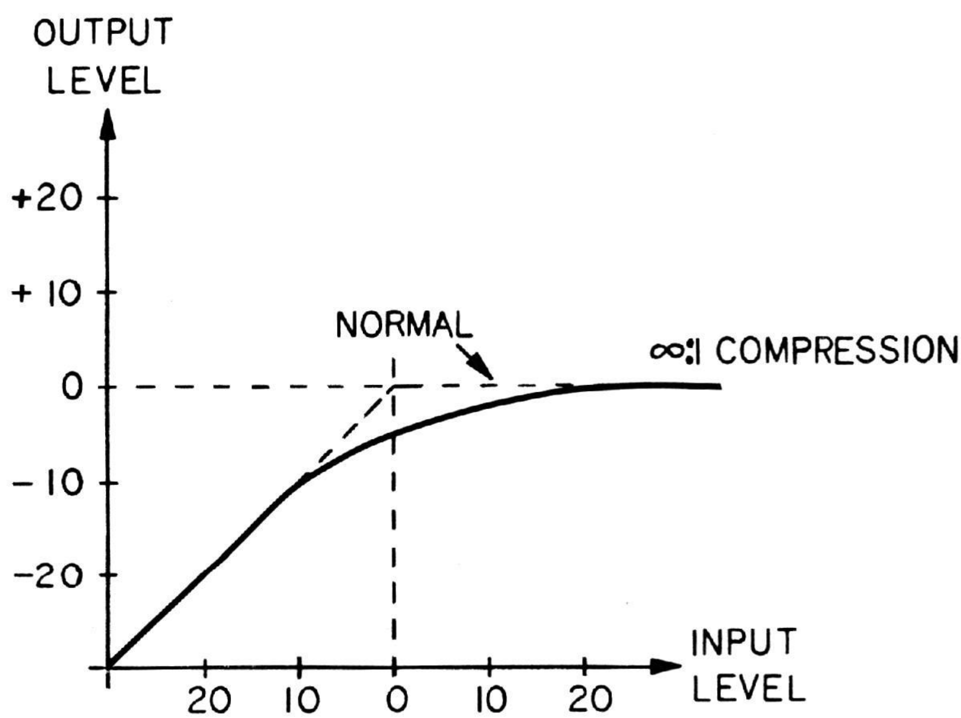

FINAL VERSION OVER EASY COMPRESSOR

FIG.9

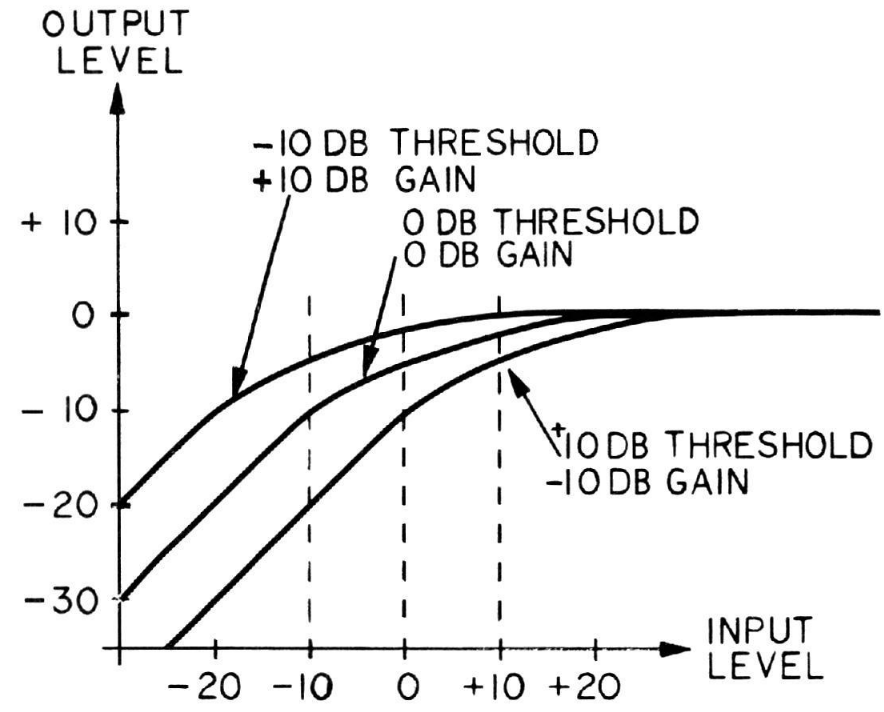

OVER EASY COMPRESSOR, GAIN=-THRESHOLD

FIG. 10

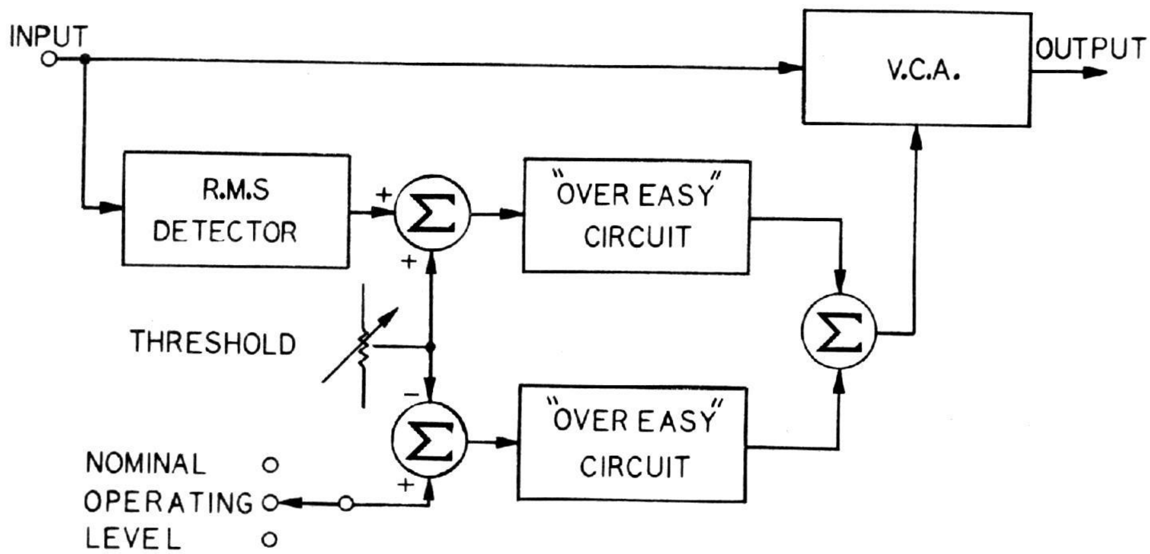

BLOCK DIAGRAM OF FINAL OVER EASY COMPRESSOR

FIG. 11

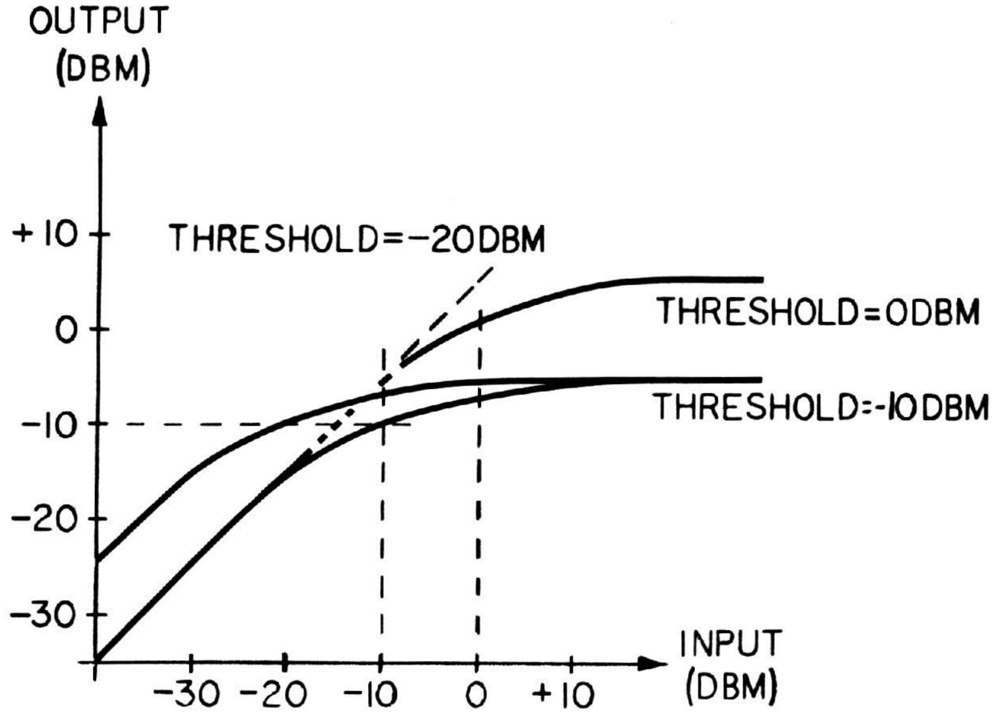

EFFECT OF CHANGING THRESHOLD ON OVER EASY COMPRESSOR

FIG. 12