DESIGNING NDFL AMPS

The use of nested differentiating feedback loops (NDFLs) is a new technique for reducing audible-frequency distortion in an amplifier to a vanishingly low level. As the name implies, NDFLs rely on negative feedback, but they use it in a new way. Edward M. Cherry, Associate Professor of the Department of Electrical Engineering, Monash University, explains the theory involved.

In order to understand just how far the new NDFL technique can improve an amplifier, we first need to know the fundamental limits to the reduction of distortion that can be achieved with conventional techniques. To begin with, we survey familiar negative-feedback theory.

Figure 1 is a block diagram of an amplifier with negative feedback. In this diagram, the forward path corresponds to the amplifier before feedback is applied, and its gain is traditionally designated by the Greek letter μ. The feedback network returns a fraction β of the output to the input circuit, where it is in some way subtracted from the true input to provide the actual input to the forward path.

In many practical amplifiers, the subtraction is accomplished by applying the input and feedback signals to the two inputs of a balanced differential first stage of the forward path. Figure 2 is an outline practical circuit. In this circuit the feedback factor β is the attenuation of the network comprising R_{F1} and R_{F2}

$$\beta = \frac{R_{F1}}{R_{F1} + R_{F2}} \qquad (1)$$

A typical value for an audio power amplifier might be 1/20. The forward-path gain μ in Fig. 2 corresponds to gain from input to output when the feedback network is removed. A typical value for a simple audio power amplifier might be 1000.

For Fig. 1, the overall closed-loop gain A is given precisely by

$$A = \frac{\text{Output}}{\text{Input}} = \frac{\mu}{1 + \mu\beta} \qquad (2)$$

The quantity $\mu\beta$ is called the loop gain. Physically, loop gain is the gain that would be observed if the feedback 'loop' in Fig. 1 was cut at some point, a signal was injected into one side of the cut, and the resulting signal at the other side of the cut was measured.

If the values of μ and β are such that loop gain is small compared with unity, the closed-loop gain is very nearly equal to the forward-path gain (that is, the gain without feedback)

$$A \rightarrow \mu \quad (\mu\beta \ll 1) \qquad (3)$$

However, if loop gain is large compared with unity, the closed-loop gain approaches the reciprocal of the feedback factor and becomes almost independent of the forward-path gain

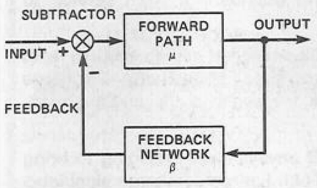

Fig. 1 Block diagram of a feedback amplifier.

$$A \rightarrow 1/\beta \quad (\mu\beta \gg 1) \qquad (4)$$

The quantity 1/B is often called the demanded gain, as it is the value the overall closed-loop gain would take in ideal circumstances.

As a numerical example, if we substitute the above values $\mu = 1000$ and $\beta = 1/20$ into Equation 2, the gain of our 'typical' audio power amplifier works out as A = 19.6. The approximate Equation 4 predicts A → 20, within 2% of the correct answer.

The quantity $1 + \mu\beta$ occurs often in feedback theory. It is called the return difference F.

$$F = 1 + \mu\beta \qquad (5)$$

Physically, return difference has the significance

$$F = \frac{\text{forward-path gain}}{\text{closed-loop gain}} \qquad (6)$$

For values of loop gain greater than about 10, loop gain and return difference are almost equal — in our 'typical' example the value are 50 and 51 respectively.

Simplified treatments of feedback theory show that, if the distortion generated in the forward path (that is, the amplifier without feedback) at a particular output signal amplitude is D_μ, then the resulting closed-loop distortion D_A at the same output signal amplitude is

$$D_A = D_\mu/F \qquad (7)$$

Distortion is improved when feedback is applied to an amplifier by a factor equal to the return difference. In our 'typical' amplifier, F = 51; if the distortion without feedback happened to be 10%, then feedback should reduce the distortion to 0.196%.

More rigorous treatments of feedback theory show that Equation 7 is no more than a poor approximation to the truth. In the first place, real amplifiers are far more complicated than Fig. 1 suggests, because several different feedback paths (not all intentional!) can be identified. For example, the collector-base capacitances of transistors inevitably provide some unintended feedback at high frequencies. There is a very real problem in interpreting just

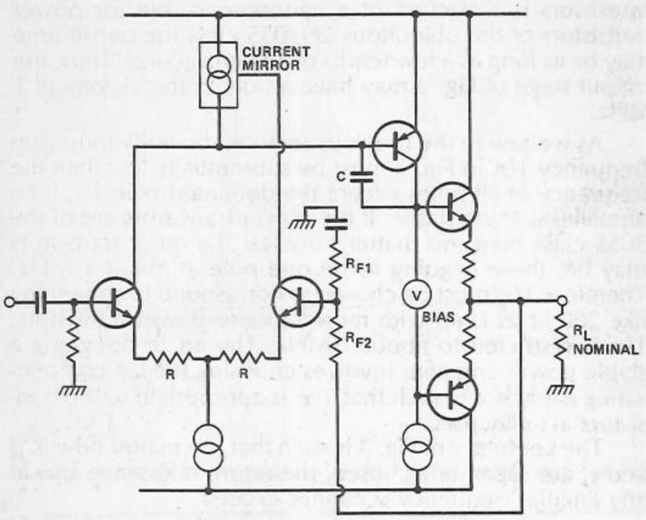

Fig. 2 Outline circuit of an audio power amplifier.

what loop gain and return difference mean when there is more than one feedback loop. Once the correct interpretation is established, return difference invariably turns out to be a function of frequency, and the reduction of distortion corresponding to Equation 7 depends on the value of return difference at the frequency of the distortion, not the frequency of the input. Feedback therefore, does not reduce all distortion components equally.

Finally, it is found that the closed-loop distortion of an amplifier can contain new components that were not present in the distortion that existed in the forward path before feedback was applied. These new distortion components initially increase as loop gain is increased, but they fall away again towards zero as loop gain is made large.

Despite all these complications, the fact remains that adequate negative feedback, properly applied, does reduce distortion. Why, then, do amplifier designers not simply apply some arbitrarily large amount of feedback and reduce amplifier distortion to the vanishing point?

TIM, IIM, PIM, . . .

In the last 10 years or so, readers of audio magazines have been made aware of a conjecture that goes something like this:

"Harmonic distortion and the usual intermodulation distortion decrease with increasing feedback. Transient intermodulation distortion (TIM) increases with increasing feedback, and is approximately directly proportional to the feedback. Therefore, there is an optimum value for the feedback at which the subjective distortion sensation is least. This optimum feedback is unlikely to exceed about 20 dB."

More recently, there has been conjecture that heavy overall feedback should be applied with caution if interface intermodulation distortion (IIM) is to be avoided. An amplifier should provide a low open-loop output impedance so that the need for feedback-generated loudspeaker damping is minimised.

There has also been conjecture that negative feedback, which reduces the usual intermodulation distortion, may increase phase intermodulation distortion (PIM) by converting amplitude nonlinearities into phase nonlinearities.

Unequivocally, none of these conjectures has any basis in the new NDFL amplifiers. As an aside, there is a substantial body of opinion that none of these conjectures has any basis, full stop; interested readers should refer to References 1-9.

Instability And Oscillation

A fundamental limit to the amount of feedback that can be applied to an amplifier is set by the onset of instability and oscillation.

If the magnitudes of the forward-path gain and demanded gain of the idealised Fig. 1 are plotted versus angular frequency ω (in radian/second) on logarithmic scales, the resulting graph looks something like Fig. 3. The 3 dB bandwidth of the amplifier without feedback is $1/\tau_\mu$, and the gain-bandwidth product (at which gain drops to unity) is $1/\tau_1$.

Because the graph is on logarithmic scales, the separation between the curves of forward-path gain and demanded gain is the loop gain (remember that, to divide two numbers, you subtract their logarithms; if you divide μ by $1/\beta$, you get $\mu\beta$). The magnitude of loop gain falls to unity at the frequency $1/\tau_x$ where the curves intersect and their separation is zero (remember that the logarithm of unity is zero).

By a similar argument, return difference is the separation between the curves of forward-path gain and closed-loop gain, as indicated in Fig. 3.

We could make a similar graph to Fig. 3, showing the phases of μ and $1/\beta$. Again, the phase of loop gain would turn out to be the separation between the two curves. However, there is a remarkable piece of mathematics due to Bode, who used a transformation evolved by Hilbert (1862-1943), which shows that there is a relation between the magnitude and phase of the response of any linear system. Subject to some qualifications, our proposed graph of the phases is completely predictable from Fig. 3 and contains no new information. Interested readers may refer to Chapter 14 of Bode's book (Reference 10) but are warned that it is anything but easy going!

As an example, many readers will know that, if the forward-path in Figs. 1 and 2 has a high-frequency cut-off rate variously described as single pole, 20 dB/decade, or 6 dB/octave, then its phase shift is 45° at the 3 dB cut-off frequency $1/\tau_\mu$, and is asymptotic to 90° at very high frequencies.

In 1932, Nyquist applied a theorem which dates back to Cauchy (1789-1857) to derive the condition for a feedback amplifier to be stable and free from oscillation. If a polar plot is made of the magnitude and phase of return difference as frequency is varied, a vaguely 'snail-shaped' curve results. Such a polar plot is called a Nyquist diagram. Subject again to some qualifications, the stability criterion for a feedback amplifier is that its polar plot of return difference should not enclose the origin. Figure 4

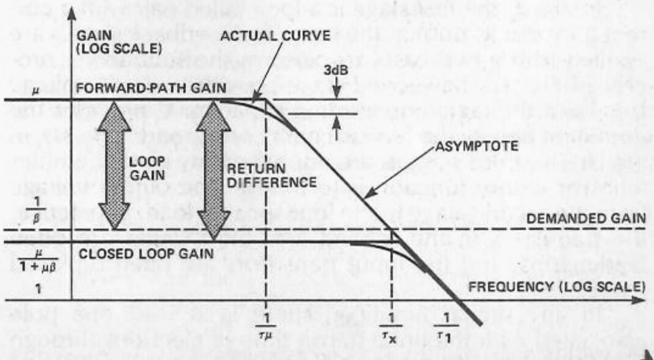

Fig. 3 Logarithmic plots of gain versus frequency for Fig. 1.

shows one example each of a stable situation and an unstable situation.

Because the phase of return difference can be predicted from Fig. 3 via Bode's result a Nyquist diagram can also be constructed from Fig. 3 and the onset of instability can be predicted. In 1945 Bode showed that Nyquist's criterion could in fact be expressed in terms of the gradients of the curves in Fig. 3, thereby eliminating the work of finding the phase explicitly and plotting the Nyquist diagram. Bode's exact rule is complicated, but a useful paraphrase is

"If in graphs such as Fig. 3 the separation between the forward-path gain and demanded gain decreases toward zero at a rate not exceeding 30 dB/decade, the amplifier is unlikely to oscillate."

This paraphrase makes no allowance for the tolerances on components. It assumes, in effect, that everything about the forward path is well known and constant. In the audio context, the paraphase takes no cognizance of the fact that the capacitance of the leads that connect an amplifier and loudspeaker is anything but well known. A more conservative rule, applicable to the audio context, is therefore

"In graphs such as Fig. 3, the separation between the forward-path gain and demanded gain should not decrease towards zero at a rate exceeding 20 dB/decade."

transistors is a fraction of a nanosecond, but for power transistors of the ubiquitous 2N3055 class the transit time may be as long as a few tenths of a microsecond. Thus, the output stage of Fig. 2 may have a pole in the vicinity of 1 MHz.

As we saw in the previous section, the unity-loop-gain frequency $1/\tau_x$ in Fig. 3 must be substantially less than the frequency of all poles except the dominant pole $1/\tau_\mu$ if an amplifier is to be stable. If the power transistors are of the 3055 class then, no matter how fast the other transistors may be, there is going to be one pole at about 1 MHz. Therefore $1/\tau_x$ must be chosen to correspond to something like 200 kHz. Even with more modern power transistors, $1/\tau_x$ is restricted to about 1 MHz. The art of designing a stable power amplifier involves choosing the lag compensating capacitor C such that $1/\chi_x$ is appropriate to the transistors actually used.

The geometry of Fig. 3 is such that, no matter how μ, β and τ_μ are separately chosen, the return difference $F(\omega)$ at any angular frequency ω cannot exceed

$$F(\omega) \leqslant 1/\omega\tau_x \qquad (8)$$

Thus, if $1/\tau_x$ is designed to correspond to 200 kHz, return difference at 20 kHz cannot exceed 10 (= 20 dB), and cannot exceed 200 (= 46 dB) at 1 kHz. An amplifier that boasts 80 dB of feedback (F = 10,000 at low frequencies) must have $1/\tau_\mu$ corresponding to about 20 Hz; return difference must begin falling above 20 Hz, and the former

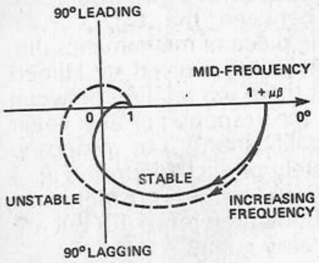

Fig. 4 Nyquist's stability criterion. The curves are polar plots of return difference for changing frequency.

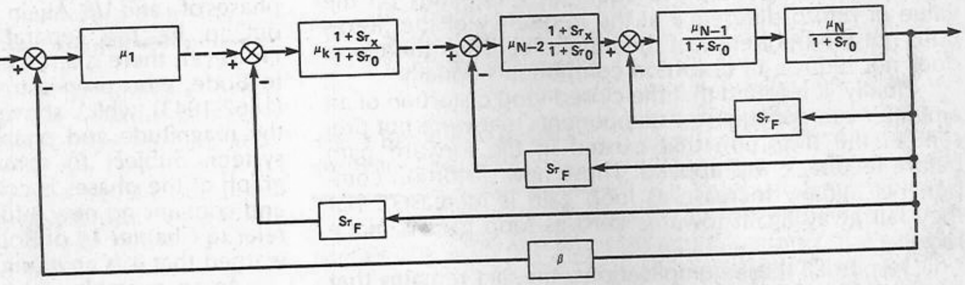

Fig. 5 Block diagram of an NDFL amplifier.

The practical consequence is that the forward path of an audio amplifier with conventional resistive feedback should have a single dominant pole which sets the fall-off of gain at frequencies above $1/\tau_\mu$. The second and subsequent poles should lie at frequencies substantially above $1/\tau_x$ (the frequency where the separation reaches zero), because each pole contributes a 20 dB/decade downwards slope to the graph of forward-gain path.

Maximum Available Feedback

In Fig. 2, the first stage is a long tailed pair with a current mirror at its output; the input and feedback signals are applied to the two bases to perform the subtraction process of Fig. 1. The second stage provides a large voltage gain, and the lag compensating capacitor C provides the dominant pole of the forward path corresponding to $1/\tau_\mu$ in Fig. 3. The third stage is a complementary class-B emitter follower whose function is to transfer the output voltage from the second stage to the loudspeaker load. In practice, the transistors in the second and third stages are often Darlingtons, and the input transistors are often replaced by FETs.

In any similar amplifier, there is at least one pole associated with the finite transit time of electrons through each transistor. The transit time for typical small-signal

values at 1 kHz and 20 kHz (46 dB and 20 dB) still apply.

Returning now to Equation 7, the effectiveness of feedback in reducing distortion is set by the frequency of the distortion, not the frequency of the input. The audible frequency range is generally reckoned to extend to about 20 kHz and, with the foregoing constraints, return difference at this frequency cannot exceed 10. Remembering that 20 kHz is the third harmonic of 6.667 kHz, we see that feedback cannot reduce offensive odd-harmonic distortion of mid-treble input signals by more than a factor of 10. Remembering too that 20 kHz is the seventh harmonic of 2.857 kHz, we see that feedback cannot reduce crossover distortion of mid-range input signals by more than a factor of 10.

Until recently there has been no way around this problem except to increase the unity-loop-gain frequency $1/\tau_x$, and this demands that the frequencies of the transistor poles must be increased if stability is to be preserved. Fragile, expensive power transistors, with narrow bases to achieve short transit times, become mandatory.

The NDFL Approach

There is, however, another solution to the stability problem. If the forward-path gain has two dominant poles, so that its gain falls at 40 dB/decade, the rate of closure

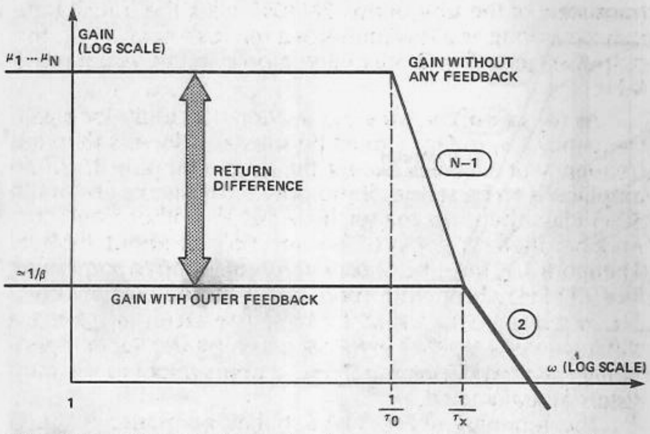

Fig. 6 Logarithmic plots of gain versus frequency for Fig. 5.

between the graphs of forward-path gain and demanded gain would still be 20 dB/decade provided the demanded gain itself were to fall at 20 dB/decade. In essentials, this requires that the usual frequency-independent resistive feedback factor β should be replaced by something having a frequency dependence of the form $\omega\tau_F$ (remember that the demanded gain is the reciprocal of the feedback factor). Mathematicians tell us that a linearly rising frequency response corresponds to differentiation with respect to time and, in hardware terms, a capacitive feedback network will perform just this action.

Figure 5 shows the outline of an amplifier incorporating nested differentiating feedback loops. Notice first that the forward path has been separated into a number of stages, whose mid-frequency gains are μ_1 to μ_N respectively. The variable s is what mathematicians call complex frequency; for sinusoidal signals its magnitude is equal to the angular frequency ω of the sinusoid. Factors of the form $(1 + s\tau_x)$ represent a frequency response that rises proportional to frequency above the frequency $1/\tau_x$ — that is, they represent a zero. Similarly, factors of the form $1/(1 + s\tau_.)$ represent a frequency response that falls inversely proportional to frequency above the frequency $1/\tau_0$ — that is they represent a pole. Thus, the stages in Fig. 5 have special frequency responses: all stages except the first have a pole at $1/\tau_.$, and all except the first and last two have a zero at $1/\tau_x$.

Notice also that there are differentiating feedback networks, each denoted by $s\tau_F$, linking the output back to various points in the forward path. The resulting feedback loops are arranged one inside another, like a nest of Chinese boxes — hence the name nested differentiating feedback loops.

The amplifier is completed by an overall resistive feedback network β.

If we removed all the feedback from Fig. 5, the forward-path gain would be shown in Fig. 6: constant up to the frequency $1/\tau_0$, then falling at an (N−1)-pole rate (20(N−1) dB/decade) up to $1/\tau_x$, and finally levelling off somewhat to a two-pole rate (40 dB/decade).

If we now applied just the overall resistive feedback β, the return difference would be as shown in Fig. 6. Distortion would be reduced by a constant large amount, approximately $\mu_1\ \mu_2\ \ldots\ \mu_N\ \beta$, at all frequencies up to $1/\tau_0$. Choosing $1/\tau_0$ to correspond to 20 kHz would virtually

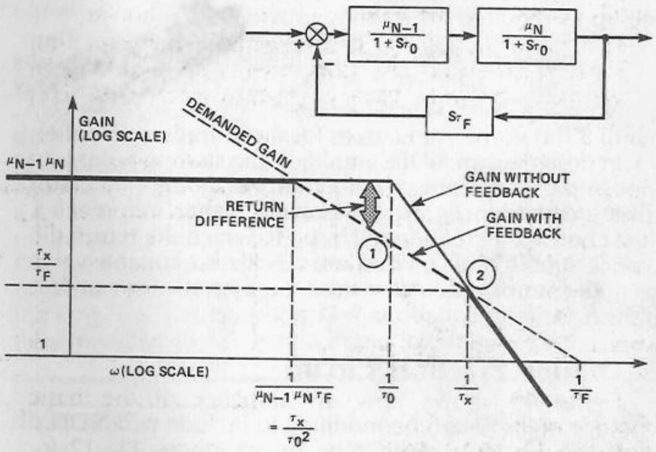

Fig. 7 The inner loop of Fig. 5.

eliminate audible-frequency distortion. *But the amplifier would be unusable because of oscillation.*

The rate of closure of the forward-path gain and demanded gain curves breaks the rule of 20 dB/decade. Let us see how inclusion of the nested differentiating feedback loops solves the problem.

Figure 7 shows just the last two stages and the inner differentiating feedback factor. This 'clump' is a feedback amplifier in its own right, and Fig. 7 shows its forward-path gain (that is, the gain of the last two stages without any feedback), the demanded gain, and the resulting closed-loop gain. Although the forward-path gain falls at a two-pole rate (40 dB/decade), the demanded gain falls at a one-pole rate (20 dB/decade), and their rate of closure is 20 dB/decade. By itself, this 'clump' is stable.

Figure 8 shows what happens when we add the antepenultimate stage and another differentiating feedback factor. Again this 'clump' can be considered as a feedback amplifier in its own right. Provided we choose.

$$\mu_{N-2} = \tau_0/\tau_x$$

the various gains line up as shown. The forward-path gain is the combined gain of stage (N−2) and stages (N−1) and N with their local feedback, and this is the middle solid curve in Fig. 8. The demanded gain is the dashed curve passing through $1/\tau_F$. Once again the forward-path gain and demanded gain close at 20 dB/decade, so the stability criterion is satisfied for this larger 'clump'.

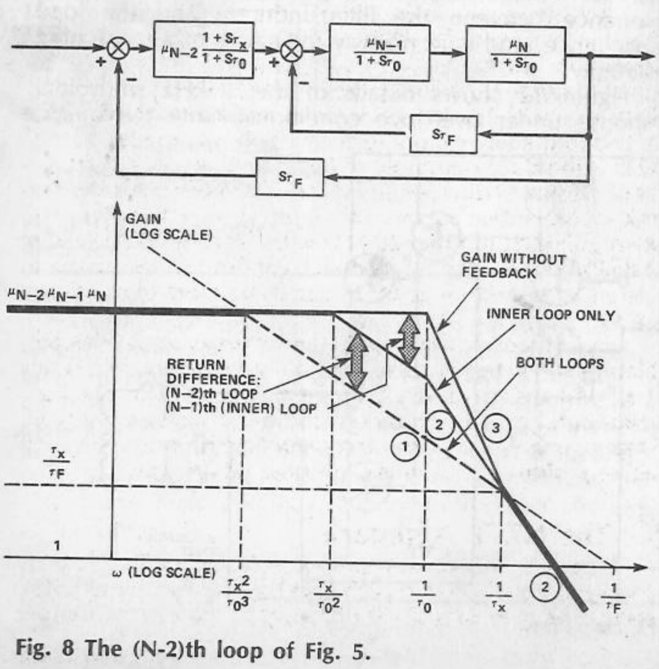

Fig. 8 The (N-2)th loop of Fig. 5.

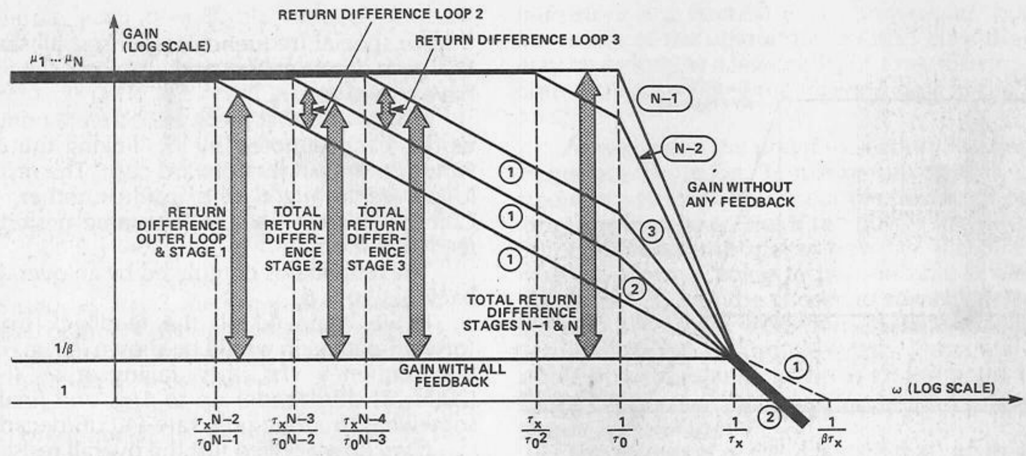

Fig. 9 Complete plots of gain versus frequency for Fig. 5.

And so it goes on. We can add more stages and differentiating feedback factors, and each time the curves line up as required for stability provided we choose

$$\mu_1\, \mu_{N-1}\, \mu_N B = (\tau_0/\tau_X)^2, \tag{9}$$

$$\tau_F = \mu_1\, \beta\tau_X, \tag{10}$$

$$\mu_K = \tau_0/\tau_X \text{ for } 2 \leqslant k \leqslant N-2. \tag{11}$$

Figure 9 shows the gain curves for the complete amplifier.

In designing an NDFL amplifier, the starting point is to choose the frequency $1/\tau_X$ so that the various transistor poles are sure to lie at substantially higher frequencies. Next choose the frequency $1/\tau_0$ up to which the return difference should remain constant; 20 kHz is a suitable value for audio amplifiers. After this, the circuit more or less designs itself via Equations 9-11 above.

Outline Practical Circuit

Figure 10 shows how an amplifier of the basic topology of Fig. 2 can be modified to include two NDFLs. Interested readers should refer to references 11, 12 for more details.

Notice first that the lag compensating capacitor, C, in the penultimate stage of Fig. 2 has been removed in Fig. 10. In its place are two capacitors (C) linking the output back to various points in the forward path. These capacitors are the feedback networks of the nested differentiating feedback loops.

The output stage has been changed to include a modified form of Thiele's load-stabilising network. Some form of LRC filter is required to locate one of the poles correctly, and with the circuit shown we get double value from the components.

The input stage itself is unchanged, but an inexpensive small capacitor in the overall feedback network β can be used to correct the group delay and improve the reproduction of transient waveforms.

Another essential addition is an amplifying stage between the two nested differentiating feeback factors. This rather peculiar circuit (which dates back to Rush in 1964) seems largely to have been forgotten. It uses one NPN transistor and one PNP to provide a well-defined gain (13).

As already suggested, once the demanded gain $1/\beta$ and the critical frequency $1/\tau_X$ are chosen, the circuit almost designs itself. The equations are:

$$\frac{R_{F1}}{R_{F1} + R_{F2}} = \beta, \tag{12}$$

$$RC = \beta\tau_X, \tag{13}$$

$$R_Y C_Y = \tau_X, \tag{14}$$

$$\tau_L = (\sqrt{3} - 1)\tau_X. \tag{15}$$

All stage gains and poles and zeros automatically look after themselves.

Figure 11(a) shows the 5 kHz square-wave response of Fig. 10 as built from 5%-tolerance resistors, 20%-tolerance capacitors, and unselected production transistors. Evidently the circuit is 'designable'; Equations 12-15 really do predict component values for good transient response.

A nice feature of the modified Thiele circuit in Fig. 10 is that, when the load is made capacitive (a well-known source of high-frequency oscillation in amplifiers), the voltage waveform at the FEEDBACK POINT is the waveform the amplifier would have delivered into its nominal resistance load. Figures 11(b) and (c) illustrate this; the violent ringing in Fig. 11(b) is simply an LC resonance between the filter inductor and the load capacitance, and is in no way indicative of approaching instability.

Figure 12 shows details of the 1 kHz sinusoidal response under overdrive conditions. Note the quick,

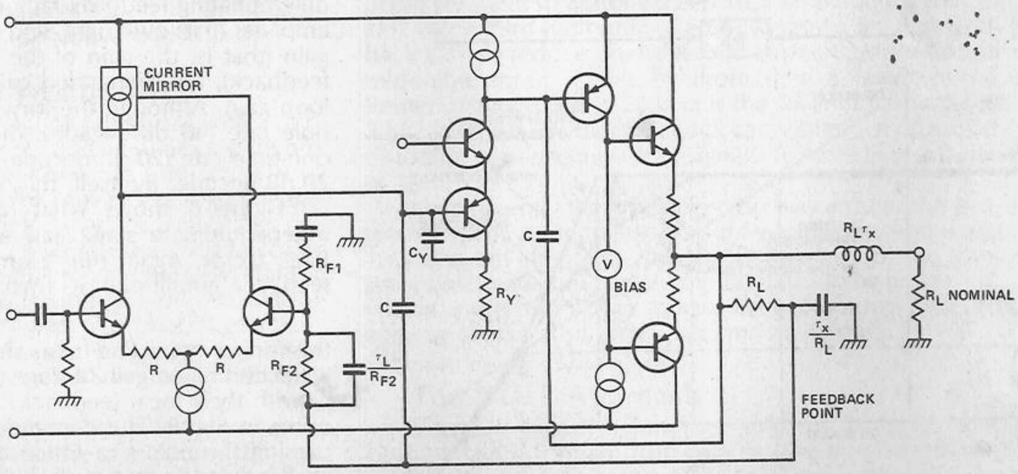

Fig. 10 Outline circuit for an NDFL amplifier.

Fig. 11 5 kHz square wave response of Fig. 10.

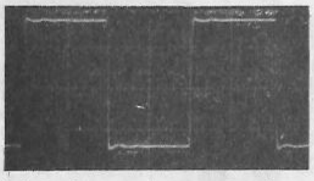

(a) 8 ohm resistance load.

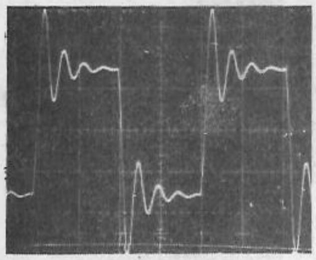

(b) 8 ohm and 2uF parallel load.

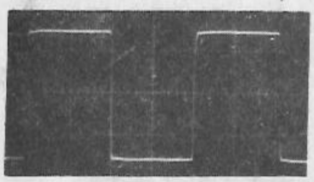

(c) waveform at feedback point for (b).

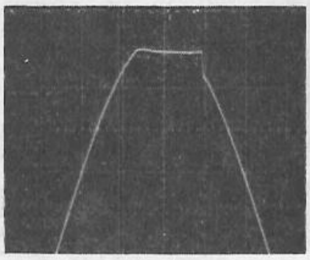

Fig. 12 Detail of output waveform from Fig. 10 under overdrive.

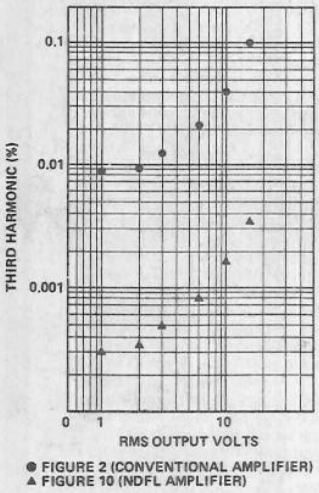

Fig. 13 1 kHz harmonic distortion.

Fig. 14 2 kHz crossover distortion when bias is set wrongly.

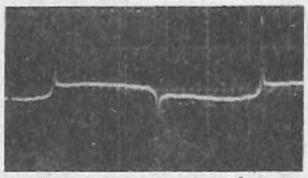

(a) Fig. 2 (conventional amplifier).

(b) Fig. 10 (NDFL amplifier).

clean recovery.

An amplifier has been built in which the circuit can be switched from Fig. 2 to Fig. 10, to illustrate the improvement in performance of adding two NDFLs. Figure 13 compares the measured third-harmonic distortions of 1 kHz. Notice how the distortion of Fig. 10 drops away to below three parts per million at small signal amplitudes. Such behaviour is more typical of class-A amplifiers than class-B amplifiers, and may account for the clean sound of NDFL amplifiers.

Crossover distortion associated with incorrect bias of the output stage is one of the most audibly annoying forms of distortion. Audio amplifiers based on Fig. 2 sometimes have a type of crossover distortion that does not show up in normal measurements. Correct biasing of the output stage relies on close tracking of the thermally-compensated biasing device and the power transistors. At best the biasing device can be thermally bonded to the power transistor case. More usually it is bonded to the heatsink, but there is no way it can simultaneously sense the actual junction temperatures of all the power transistors. Under rapidly-fluctuating dynamic signal conditions, the junction temperatures may be wildly different from each other and from the case or heatsink temperatures, and therefore the biasing may be wrong.

Figure 14 compares the static cross-over distortion of Figs. 2 and 10 when the bias is deliberately set 0V5 too low. Dynamic mistracking of the biasing circuit should not introduce audible crossover distortion in an NDFL amplifier.

One final point. The NDFL technique maximises the return difference (and hence minimises distortion components) at frequencies up to $1/\tau_0$. Above this frequency the return difference falls away rapidly, and distortion rises. Choosing $1/\tau_0$ to correspond to 20 kHz minimises audible-frequency distortion, but does not minimise ultrasonic distortion.

For example, a common specification for audio power amplifiers is their THD at 20 kHz. The harmonics of 20 kHz lie at 40 kHz, 60 kHz, 80 kHz, and so on. All are ultrasonic (and hence inaudible) and the NDFL technique does not minimise them. A measurement of THD at 20 kHz may therefore give a quite misleading indication of an NDFL amplifier's audible performance. Valid objective tests include the SMPTE and CCIF tests for two-tone intermodulation distortion, the proposed IEC test for TIM (14), Cordell's proposed three-tone test for TIM (15) and the proposed test for input-output intermodulation distortion IOD (6). The distinguishing feature of all these tests is that they measure the distortion at audible frequencies.

References

1. W. G. Jung, M. L. Stephens and C. C. Todd, *An overview of SID and TIM,* Audio, vol 63; part 1, pp 59-72, June 1979; part 2, pp 38-47, July 1979; part 3, pp 42-59, August 1979.
2. R. R. Cordell, *Open-loop output impedance and interface intermodulation distortion in audio power amplifiers,* 64th Audio Eng Soc Convention, preprint no. 1537, Dec 1979.
3. R. R. Cordell, *Another view of TIM,* Audio, vol 64; part 1, pp 38-49, Feb 1980; part 2, pp 42-59, March 1980.
4. E. M. Cherry, *Transient intermodulation distortion: Part 1 — hard nonlinearity,* IEEE Trans, vol ASSP-29, pp 137-146, April 1981.
5. R. R. Cordell, *Phase intermodulation distortion — instrumentation and measurement results,* 70th Audio Eng Soc Convention, preprint 1842, Nov 1981.
6. E. M. Cherry and G. K. Cambrell, *Output resistance and intermodulation distortion of feedback amplifiers,* J. Audio Eng Soc, vol 30, pp 178-191, April 1982.
7. E. M. Cherry, *Feedback, sensitivity, and stability of audio power amplifiers,* J. Audio Eng Soc, vol 30, pp 282-294, May 1982.
8. E. M. Cherry and K. P. Dabke, *Transient intermodulation distortion: Part 2 — soft nonlinearity,* IEEE Trans, to be published*.
9. E. M. Cherry, *Amplitude and phase of intermodulation distortion,* J. Audio Eng Soc, to be published*.
10. H. W. Bode, *Network analysis and feedback amplifier design,* van Nostrand (Princeton NJ) 1945.
11. E. M. Cherry, *Nested differentiating feedback loops in simple audio power amplifiers,* J. Audio Eng Soc, vol 30, pp 295-305, May 1982.
12. E. M. Cherry, *A new result in negative-feedback theory, and its application to audio power amplifiers,* Int J. Circuit Th, vol 6, pp 265-288, July 1978.
13. C. J. Rush, *New techniques for designing fast-rise transistor pulse amplifiers,* Rev Sci Instr, vol 35, pp 149-156, Feb 1964.
14. IEC Publication 269.3, *Part III — Amplifiers: Clause 22.6 "High-frequency intermodulation distortion".* (Proposal dated June 1981).
15. R. R. Cordell, *A fully in-band multitone test for transient intermodulation distortion,* J. Audio Eng Soc, vol 25, pp 578-586, Sept 1981.

*Manuscript copies available from the author.

ETI