Envelope Demodulators For AM

Updated 8 December 2012

A discussion of numerous topologies, pointing out
their problems , offering solutions to those problems
and showing their performance in various ways.
A series of commercial circuits is also ciscussed.

by James L. Tonne
W4ENE
Tonne Software
www.TonneSoftware.com

This paper is targeted toward technically-inclined personnel who have a serious interest in the design of circuitry to demodulate the traditional amplitude-modulated signal.

The beginning part discusses the ordinary diode-type envelope demodulator as used in typical AM radio receivers, in a tutorial fashion.   This part of the paper uses only ideal diodes and related parts. Illustrated first are two major faults which are commonly observed by critical users and which are seldom discussed in the literature, especially in introductory tutorials. Those faults and the reasons for them are somewhat subtle and tend to complicate the usual discussions of how the envelope demodulator functions.

Then we go on to explore various topologies using real components.

Finally we look at some actual commercial implementations of these demodulators.

The depth of modulation used throughout this paper has been fixed at 85% with a modulating frequency of 1000 Hz. The carrier frequency has been set initially to 25 kHz; this low carrier frequency is used here (as in most introductory tutorials) to clearly illustrate the operation of the ordinary diode type of envelope demodulator. After this introductory part of the paper we move the carrier upward to 500 kHz

Synchronous and digital demodulators are outside the scope of this paper.

LTspice users will recognize the schematics and waveforms in this paper as having been drawn and the waveform displayed using that program. LTspice also provided information about distortion of the demodulated signal. That distortion data is presented using a graphing program written for this paper by *ye scribe*. We have no connection with Linear Technology Corporation or LTspice other than being a grateful user.

Table of Contents

The Simple Textbook Ideal Diode Demodulator

The basic AM demodulator of the diode type is shown in Figure 1. The component values are typical of many solid-state radios. By the simple operation of dividing the capacitor value by 100 and multiplying the resistor value by 100 they will then be typical of many vacuum-tube demodulators as well. The item now under discussion, the controlling factor, is the time-constant (the R and C product) of the load seen by the diode.

In Figure 1 we have an amplitude-modulated RF signal applied to a textbook diode demodulator.

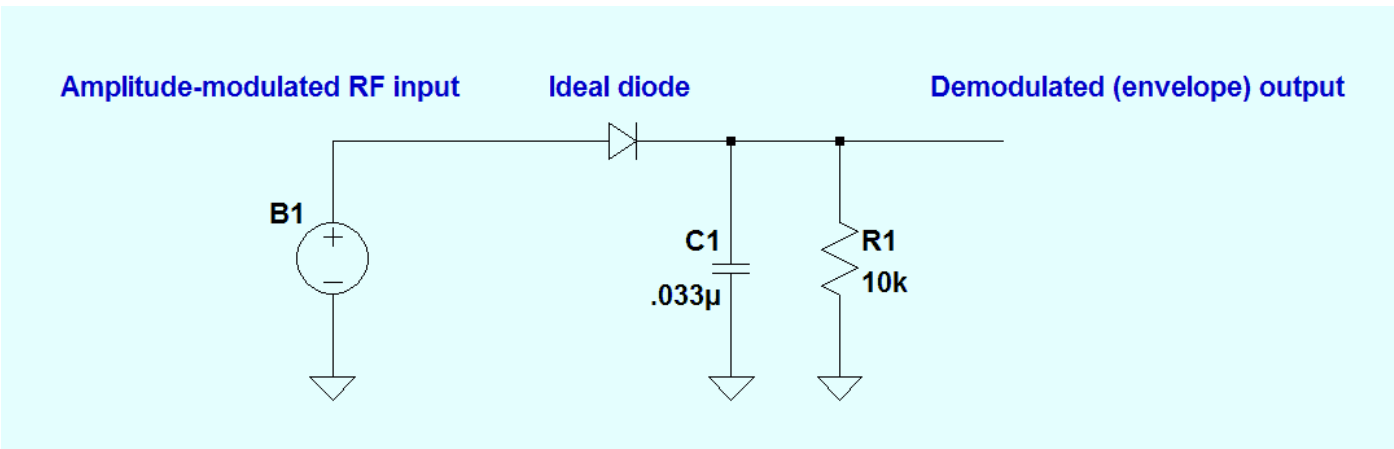

Figure 1 - the textbook diode demodulator

The carrier to be modulated - at 25 kHz for this illustration - is illustrated in Figure 2.

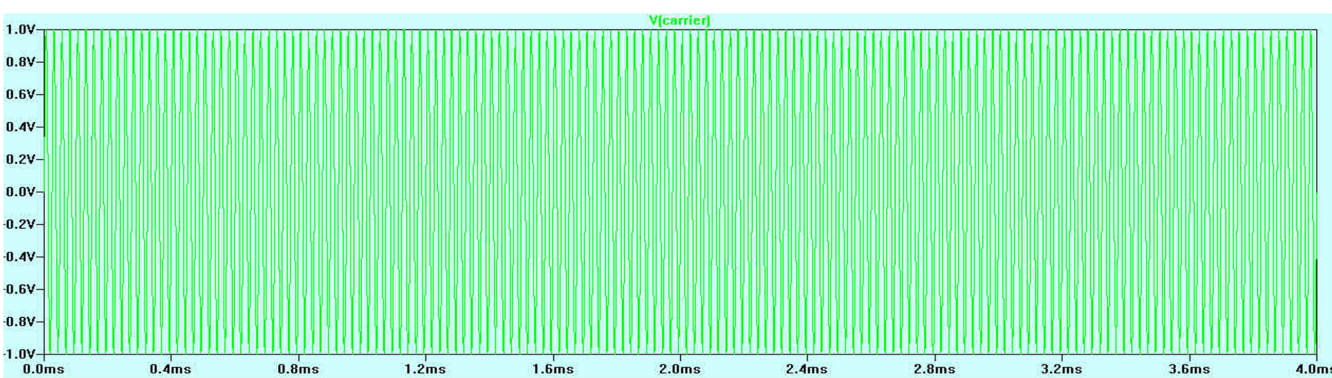

Figure 2 - 25 kHz carrier

The modulating signal is shown in Figure 3. This signal is composed of a steady DC component of 1 volt onto which has been added a 1000 Hz sinusoid with an amplitude of 0.85 volts peak.

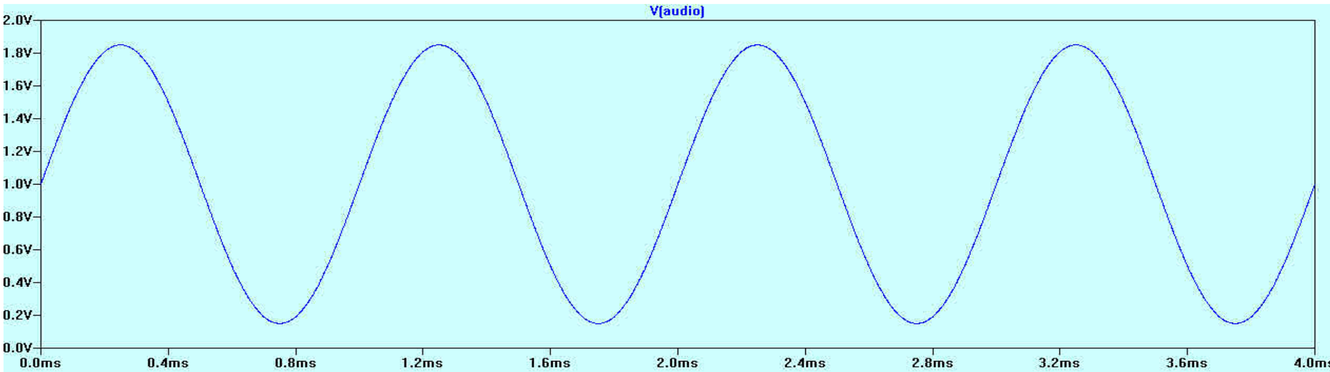

Figure 3 - modulating waveform

Figure 4 shows the modulated carrier - the signal as transmitted or as it would be delivered to the demodulator by the receiver's I.F. strip. This signal was generated by multiplying the modulating wave by the carrier wave. The "idling" carrier level (no audio, just the steady DC component to set carrier level) would be 1 volt peak. With modulation (audio added to the steady DC component) we have an output amplitude which varies over the course of the modulating wave.

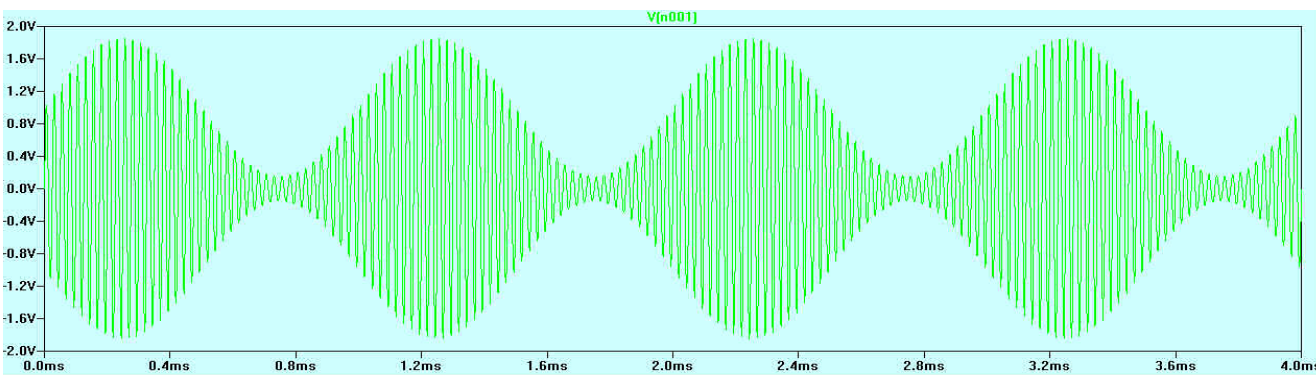

Figure 4 - the modulatd carrier

Figure 5 shows both what we would ***like*** to see at the output of the ideal diode demodulator - the sinusoid - along with what we ***actually*** see - the trace with residual RF. Each cycle of RF charges the capacitor via the diode. The capacitor can only discharge through the resistor across it. The capacitor charges up to the peak value of the modulated wave and in this simple circuit it discharges in an exponential fashion toward ground at a very clearly-defined rate (which is a function of the RC time-constant, R1 and C1 in Figure 1).

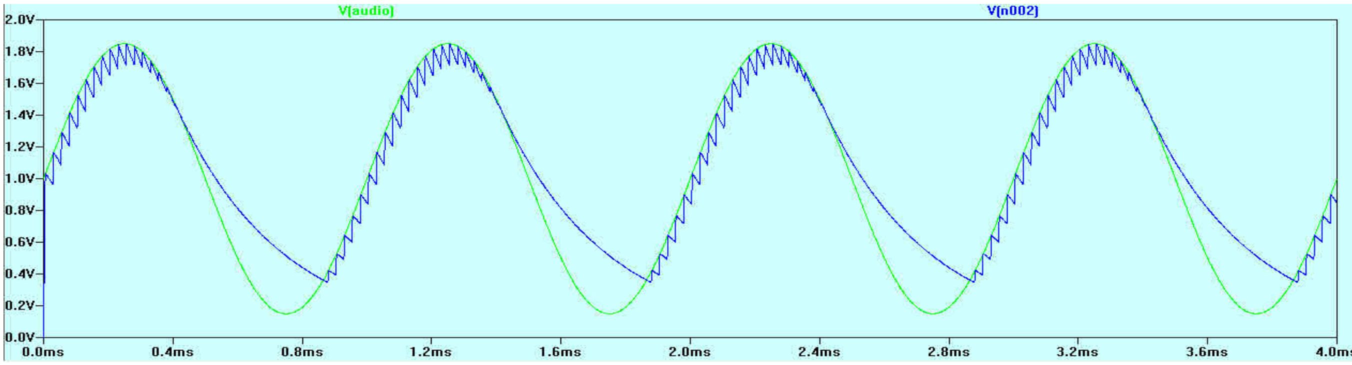

Figure 5 - Ideal and actual outputs from the demodulator in Figure 1

Increasing the time constant of the load (parallel R1-C1) will reduce the level of the residual or "feedthrough" RF. It will also increase the negative-peak "diagonal clipping" (as Terman calls it). Here we will call it slew-rate distortion.

The residual RF is more obvious on the positive peaks of the waveform and the slew-rate problem is more evident on the negative-going portion. The slew-rate distortion becomes worse as the modulating rate increases, as the modulation level increases and as the time-constant increases. This latter situation will be looked at shortly.

Typically, the selectivity in the receiver's I.F. strip reduces the amplitude of the sidebands at the higher audio frequencies, reducing the modulation depth as seen by the demodulator. As a result, slew-rate distortion at the higher modulating frequencies is generally less of a problem in practice.

We have used a carrier frequency of only 25 kHz for these displays so that the operation of the diode demodulator can be seen quite easily. When the carrier frequency is changed upward to 500 kHz, to more closely approximate a practical scenario, it would appear as shown in Figure 6.

The Simple Demodulator at 500 kHz

Now we are going to raise the carrier frequency up to 500 kHz; this is shown in Figure 6.

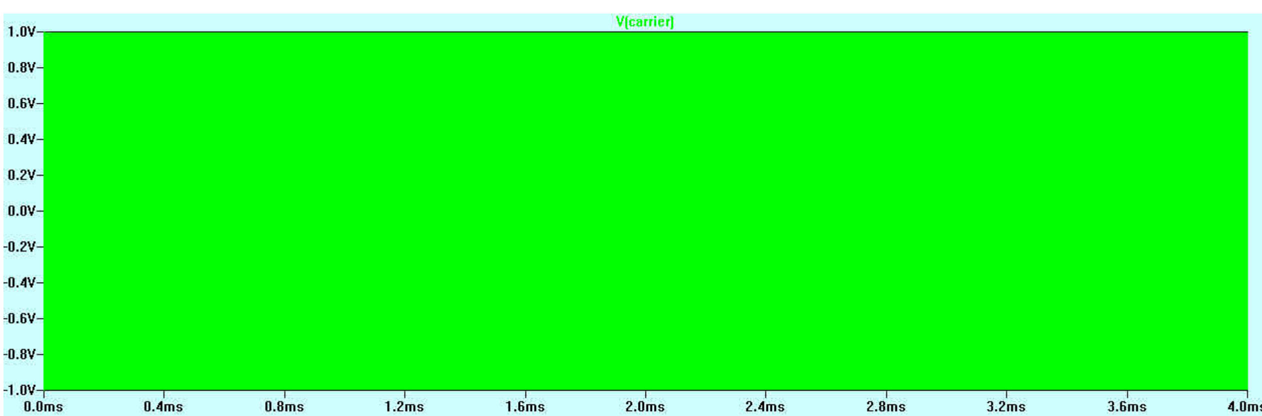

Figure 6 - the 500 kHz carrier

The modulating signal, just as in the introduction. It is still the steady DC component of 1 volt onto which has been added a 1000 Hz sinusoid with a peak amplitude of 0.85 volt for a modulation percentage of 85%. This modulating waveform is shown in Figure 7.

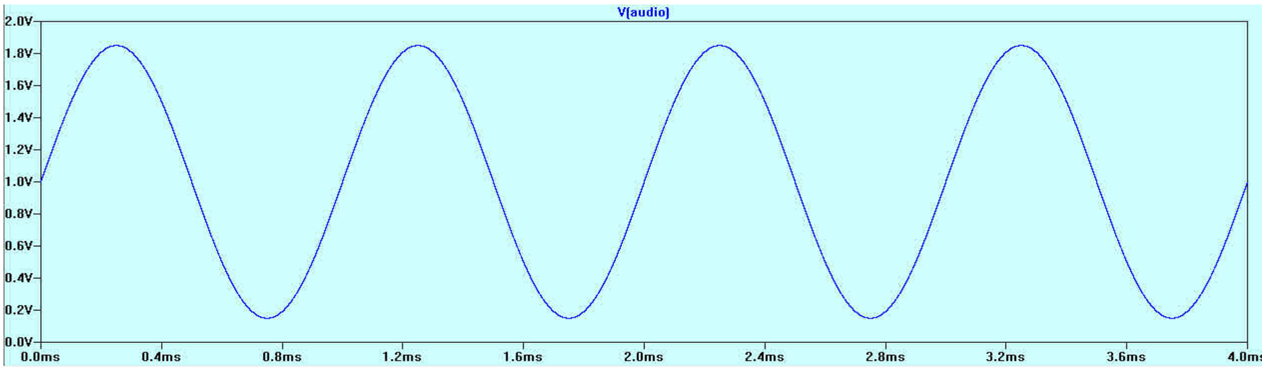

Figure 7 - the modulating signal

The resulting modulated carrier is shown in Figure 8.

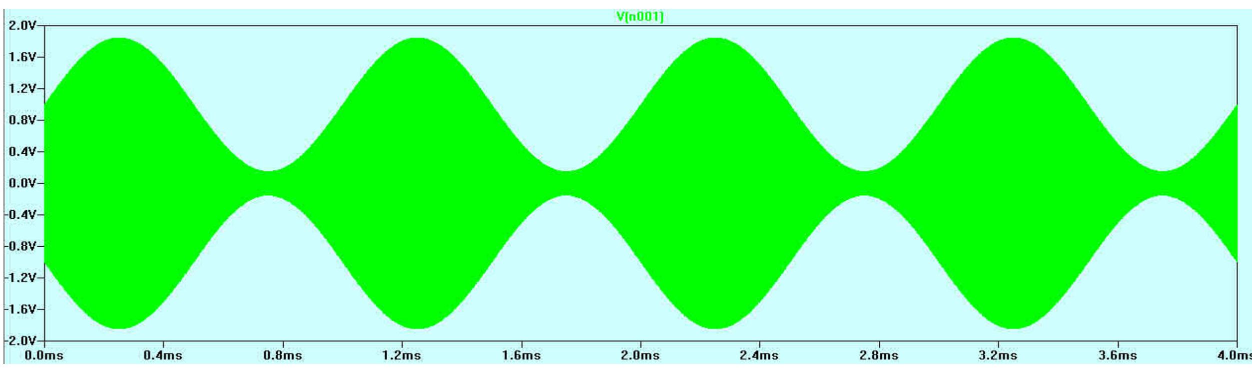

Figure 8 - the modulated 500 kHz signal

Using a 500 kHz carrier, what we would **like** to see at the output of the ideal diode demodulator is shown along with what we will **actually** see, both in Figure 9.

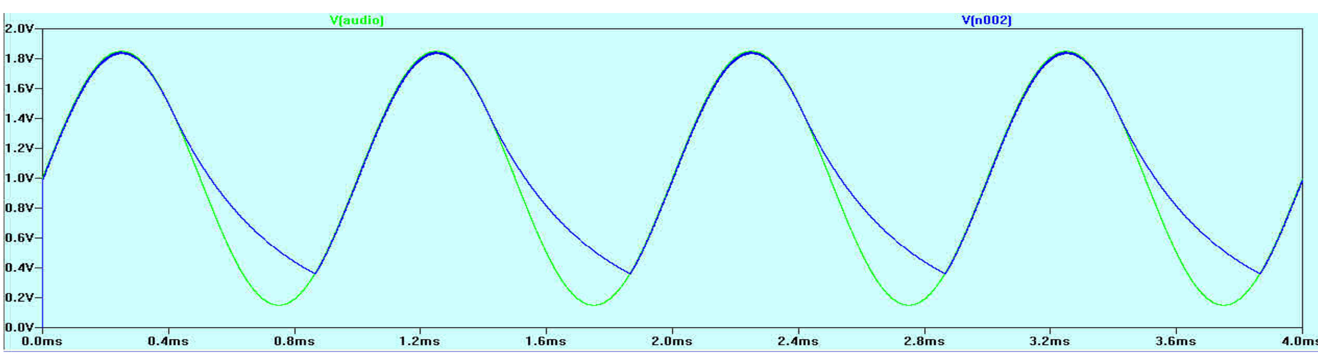

Figure 9 - the slew-rate problem clearly seen

The amplitude of the residual RF is now lower because of the increased carrier frequency.

Now let us look at this slew-rate problem again while changing the time constant. Figure 10 illustrates what happens at the upper audio frequencies (5000 Hz here) when the slew-rate problem enters into the picture.

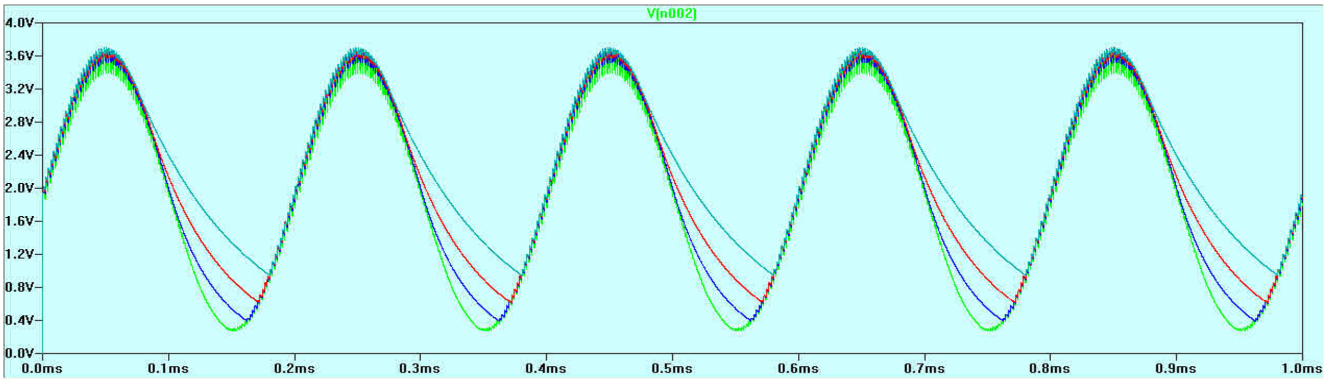

Figure 10 - 500 kHz carrier modulated by a 5000 Hz sinusoid.

Figure 10 illustrates the effect of four different time-constants (32, 56, 85 and 127 microseconds) The modulating frequency of 5000 Hz is used to emphasize the slew-rate phenomenon. As the time-constant is increased the amount of residual RF is lowered but the slew-rate distortion becomes greater. Logically the best course of action is to use a minimum of time-constant and them remove the RF, should it be problematical, later in the system.

The first major point to be noted in this paper is that with this rather common AM demodulator we can have a slew-rate problem on the negative-going excursions. The output of the demodulator decays at an exponential rate toward ground in this circuit and so in theory will never actually get to zero volts.

This phenomena will become less of an issue as the modulation percentage is reduced and especially at low modulating frequencies.

The Simple Demodulator with an AC Load

The design of these demodulators as they are seen in typical radio receivers involves not only the parallel R-C load on the diode demodulator but also a blocking capacitor prior to a resistive volume control. In some cases the volume control forms the resistive part of the diode load but its output is capacitively coupled to the following audio amplifier. In either case the diode sees a different (lower) load for modulating frequencies than it does for the unmodulated signal.

A typical schematic illustrating this is shown in Figure 11. The thing we should be concerned about is that in parallel with the diode load we have an added load which is capacitively coupled to the resistive part of the diode load.

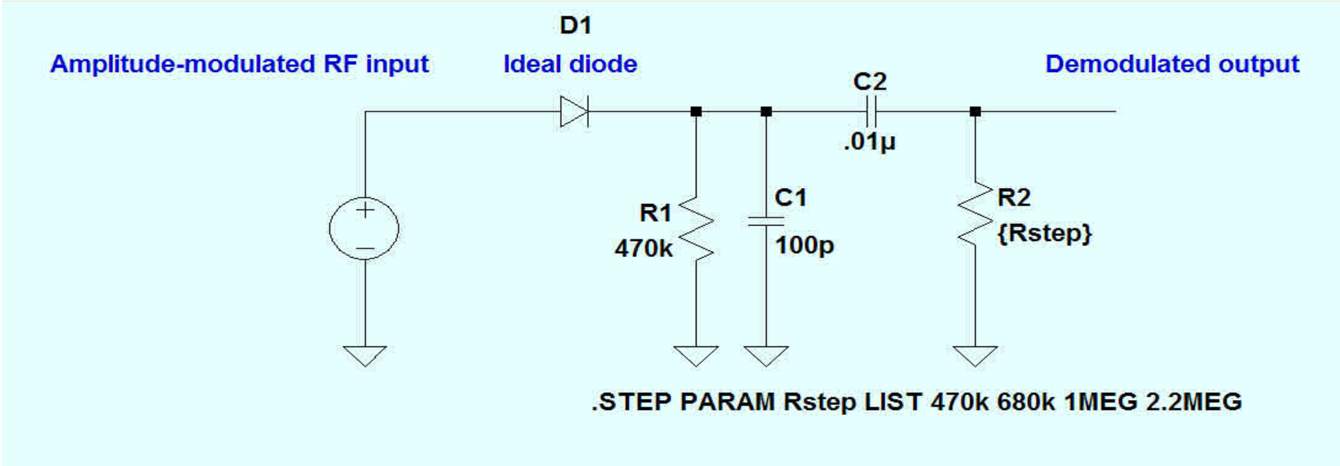

Figure 11 - AC load in parallel with the original diode load

The result of adding such an AC load in parallel with the DC load is shown in Figure 12. This waveform is the signal as it appears at the cathode of the diode using several values for "R2."

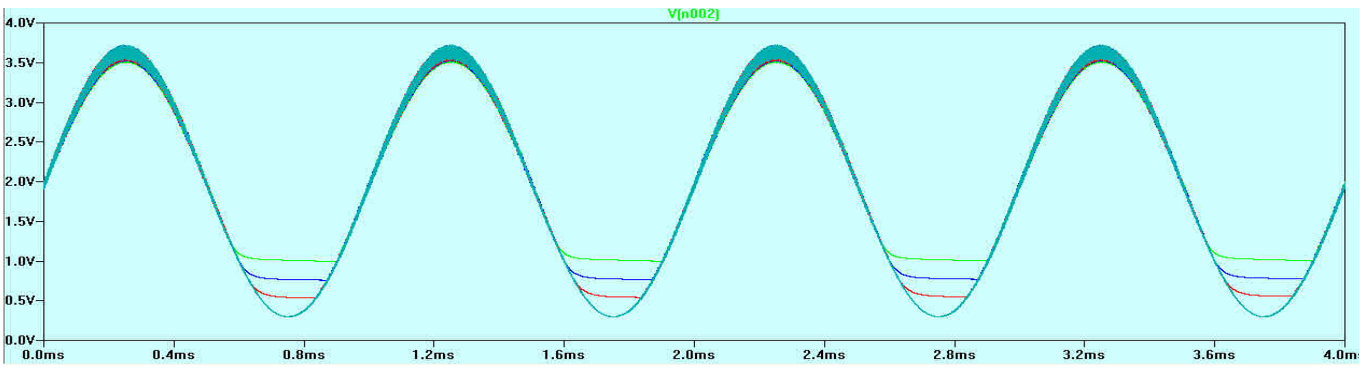

Figure 12 - Output from the diode with several values of AC load

Generally that same signal is applied to a simple R-C lowpass filter the output of which is used to provide AGC (or "AVC") to the I.F. and perhaps R.F. amplifier in the receiver. That load is yet another load that is placed in parallel with the diode's basic RC network, making the negative-peak clipping situation even worse.

Notice that the waveform shown in Figure 12 starts at a particular value (2.0) and goes up to some more positive value (3.7) on "positive" peaks.   On the negative modulation peaks the signal does not go down to 0.3 volts: we have what might be called "carrier shift" because the area under the curve above the resting (no-modulation) value does not equal the area under the curve below the resting value. It is this phenomena that causes the signal-strength meter on inexpensive receivers to move upward with modulation. Of course the clipping as seen here results in audible distortion as well.

The second major point to be noted here is that distortion of the demodulated signal is caused by an AC load on the diode demodulator. This distortion has nothing to do with diode nonlinearity. That nonlinearity will of course make the situation even worse. Note that so far in this discussion we have used only ideal diodes. Now let us examine circuitry with actual, not ideal, components.

Simple Demodulator With Semiconductor Diode

Figure 13 shows the ordinary diode demodulator but now with a semiconductor diode.

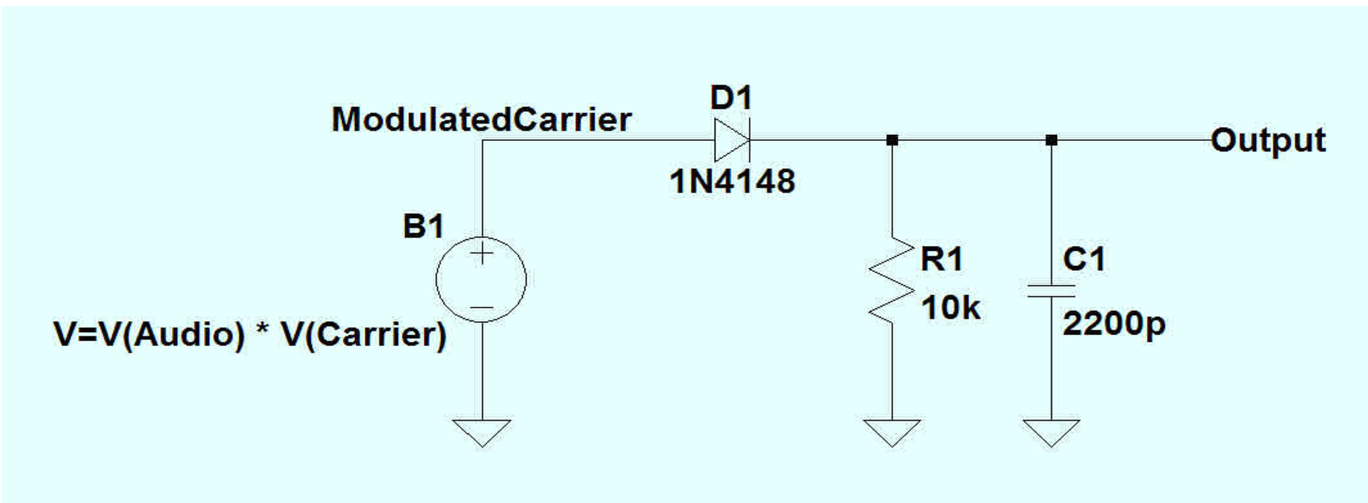

Figure 13 - Simple diode demodulator with semiconductor diode

The distortion as function of input signal level for that simple demodulator is shown in Figure 14.

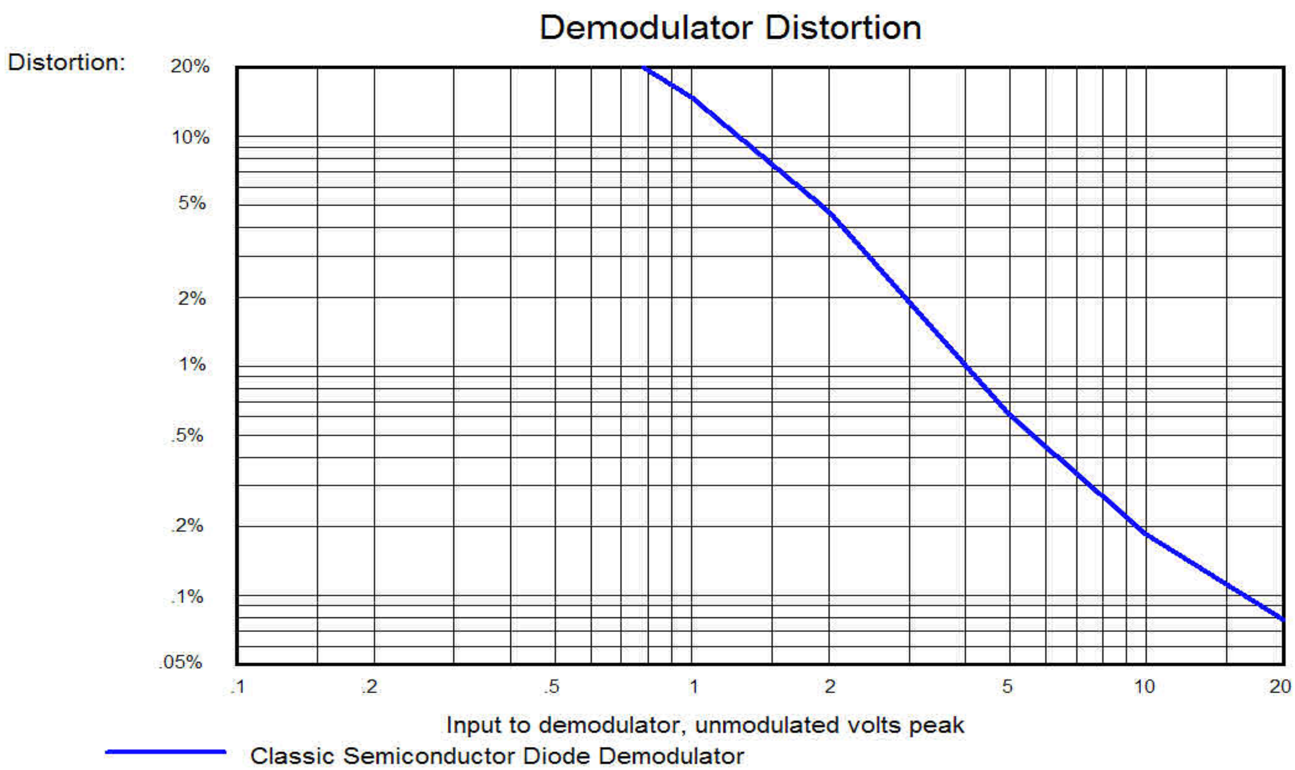

Figure 14 - Distortion vs input level for the simple semiconductor diode demodulator

To make these distortion plots we adjusted the carrier level in steps from .1 to 20 volts peak and did an FFT to recover the total harmonic distortion. In general the distortion in these demodulators will be largely even-order, mostly second-harmonic.

The results as indicated on this plot are about what might be expected from the simple semiconductor diode demodulator. With low input signal levels the distortion is severe. At high signal levels the distortion drops to a much lower value.

Forward-biased Semiconductor Diode Demodulator

Commonly, however, even in inexpensive little "transistor radios," the diode demodulator is slightly forward-biased. A circuit illustrating this is shown in Figure 15.

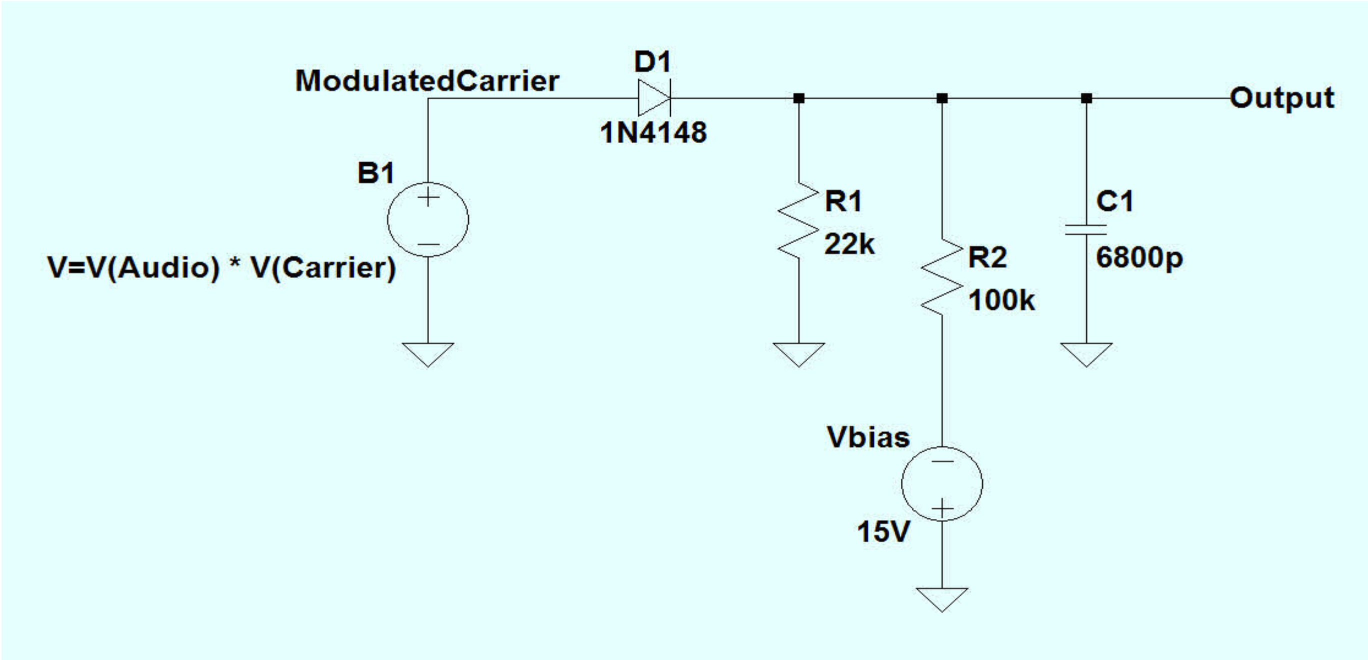

Figure 15 - A forward-biased semiconductor diode demodulator

The distortion as a function of signal level of this demodulator is shown in Figure 16.

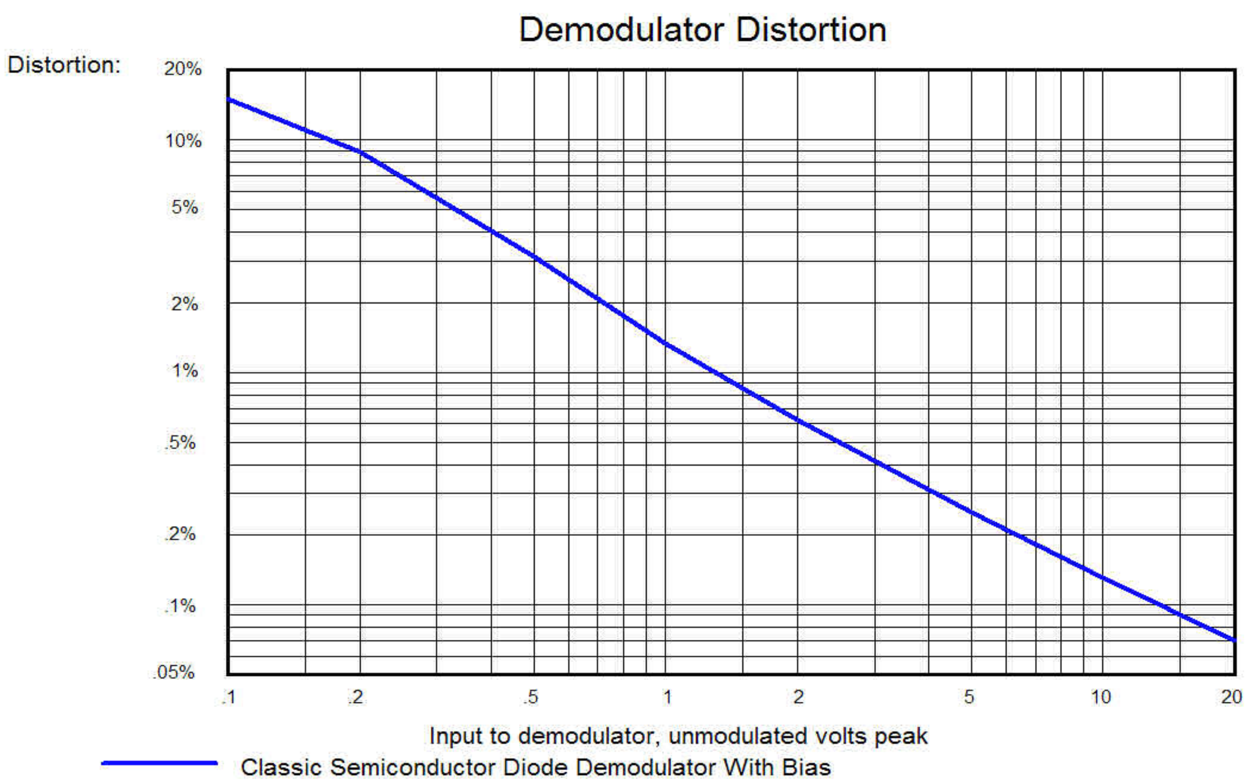

16 - Distortion vs input level for the biased diode demodulator

The slight forward bias shown in Figure 15 significantly reduces the distortion in the recovered audio, especially at lower input signal levels. Compare the distortion here with the original circuit whose distortion was shown in Figure 14.

AC Load On The Biased-diode Demodulator

Unfortunately that circuit with its fairly good distortion and wide dynamic range (it accommodates signals from a fraction of a volt on up to several volts well) is shown with no AC load on it. When we add a useable load we have the schematic as shown in Figure 17.

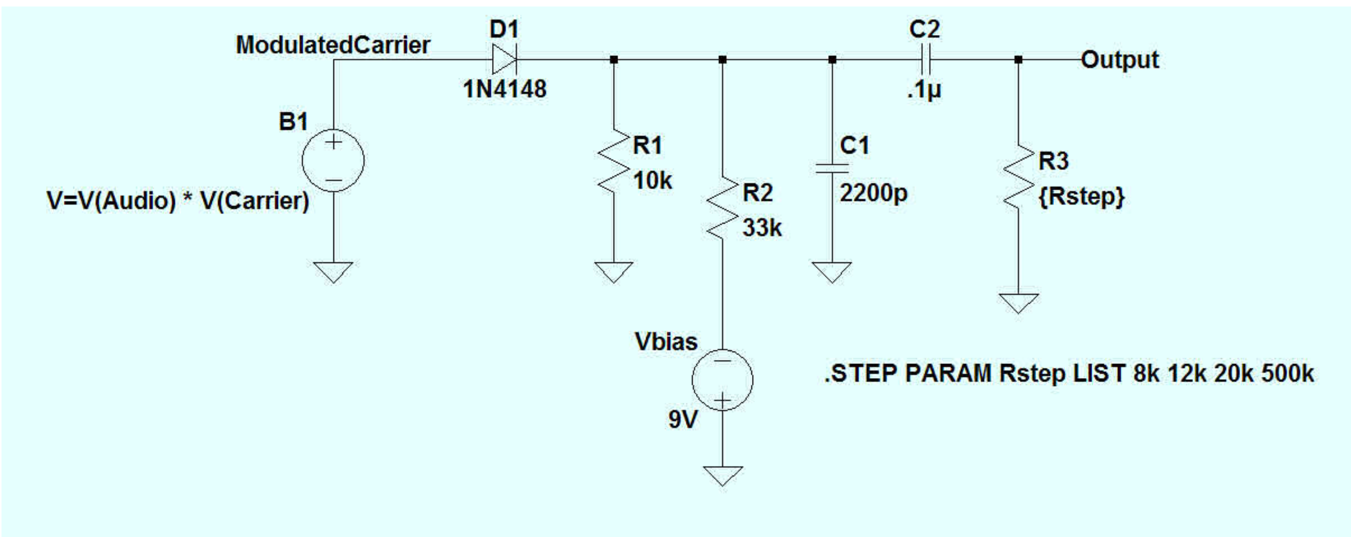

Figure 17 - Biased semiconductor diode with added AC load

In Figure 17 the load is being stepped from 8k ohms (slightly less than the DC load on the diode) on up to 12k, 20k and finally to 500k (essentially an open circuit). The harmonic distortion that results is shown in Figure 18.

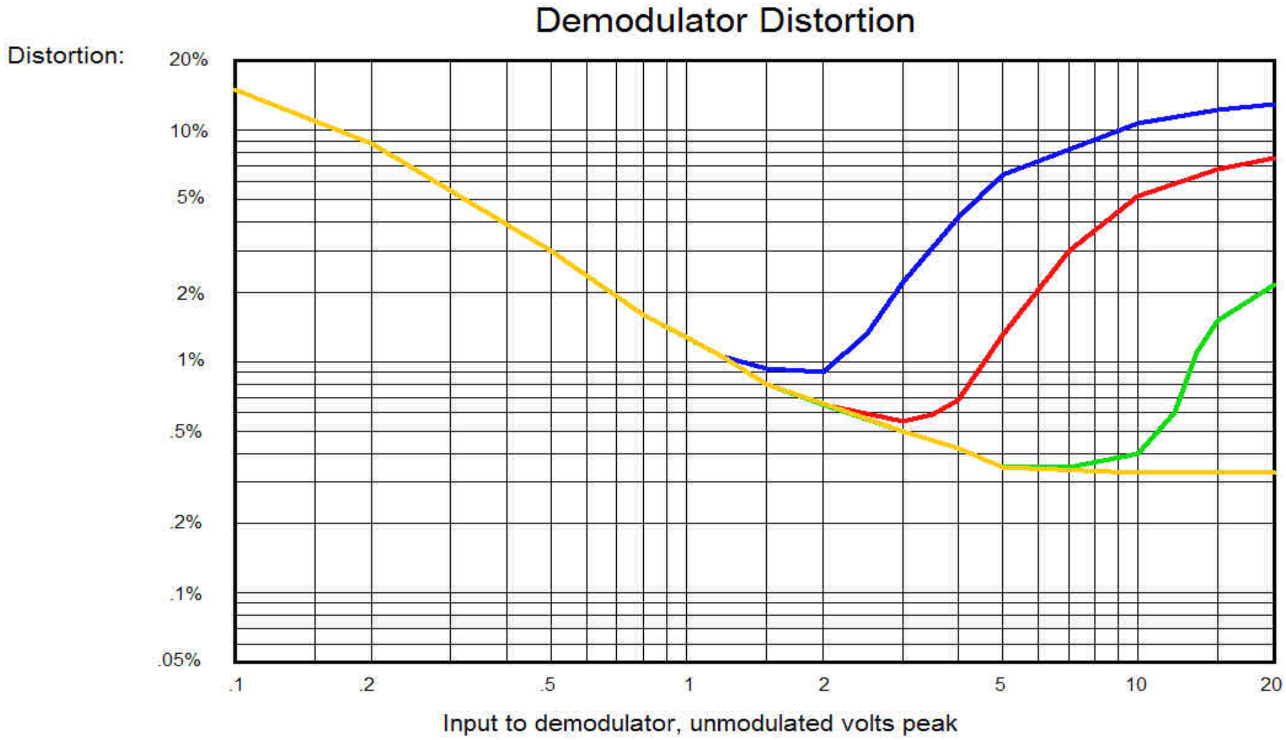

Figure 18 - The biased-diode demodulator's distortion with various values of AC load

At low signal levels the harmonic distortion is similar to that of the original demodulator with no AC load at all. As the input signal level rises the distortion increases. There are four traces on this plot, with the upper one showing the effect of the greater load (the lowest value of load resistance). The lowest trace shows the effect of a minimal load. As the AC loading gets greater (the impedance gets lower), the distortion in the recovered audio increases at high signal levels.

The Bipolar Infinite-impedance Demodulator

The infinite-impedance family has performance similar to the original diode demodulators but has a relatively high input impedance (perhaps not quite *infinite*) and perhaps better linearity. They are all basically "followers" such as an emitter follower or cathode follower. A caution to be aware of in these circuits is that followers with a capacitive load tend to be unstable and oscillate at some very high frequency. They all operate by virtue of charging a capacitor (C2 in Figure 29) from a low-impedance source (Q1 in Figure 29), which capacitor can only discharge into a relatively high value of resistance. A bipolar transistor version of the infinite-impedance demodulator is shown in Figure 19.

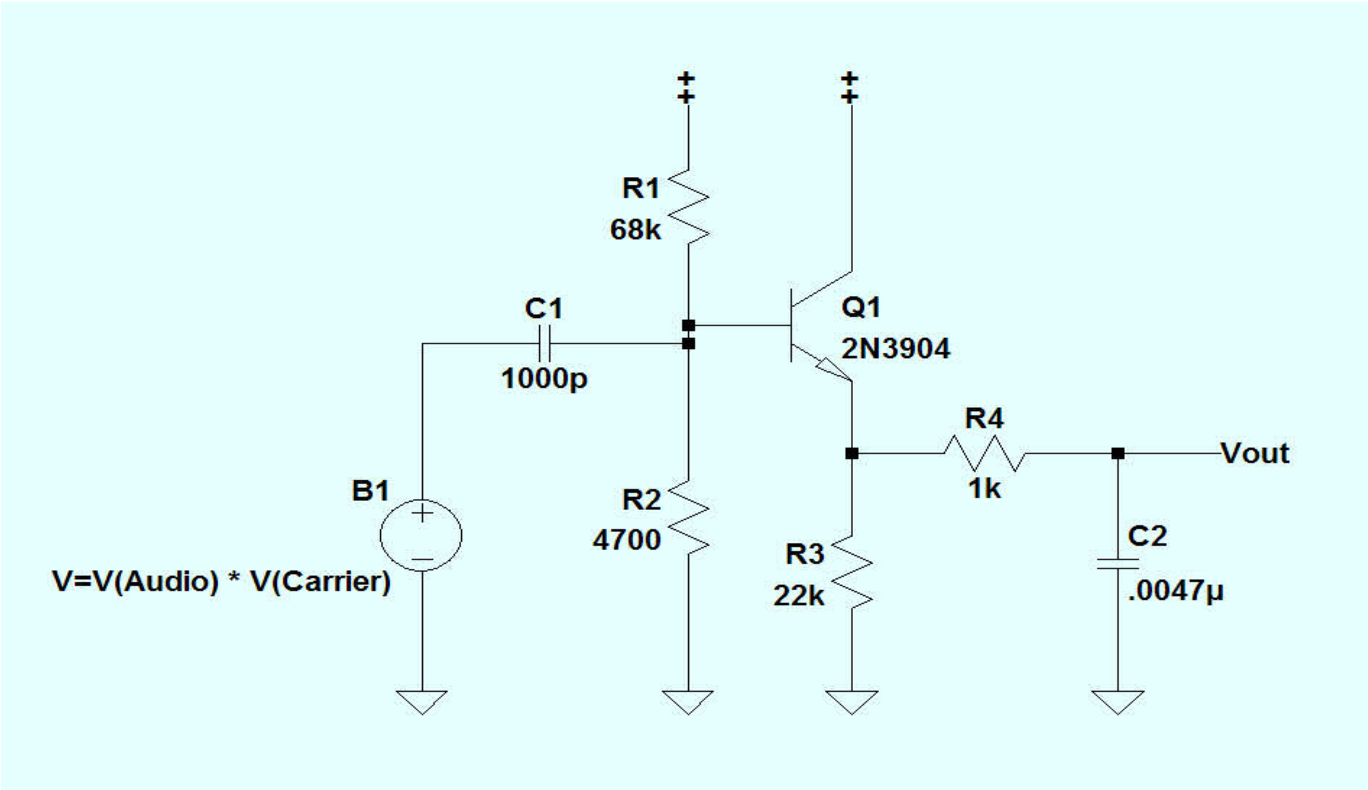

Figure 19 - A bipolar transistor infinite-impedance demodulator

The transistor is forward-biased; without that bias the distortion would be intolerable. The 1000 ohm resistor between the emitter and the capacitive load is to preclude oscillation. It does increase the distortion in the demodulated waveform a small amount.

The output from this demodulator is shown in Figure 20.

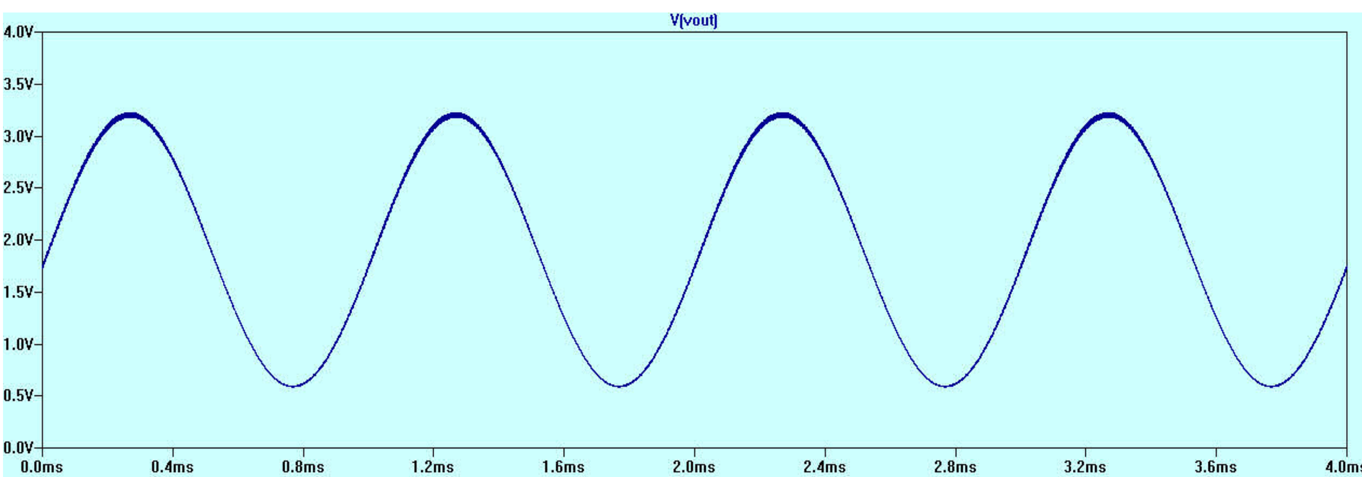

Figure 20 - The output from the demodulator in Figure 19

The output impedance is low enough to be applied to typical circuitry with a minimum of difficulty.

The distortion plot for that demodulator is shown in Figure 21.

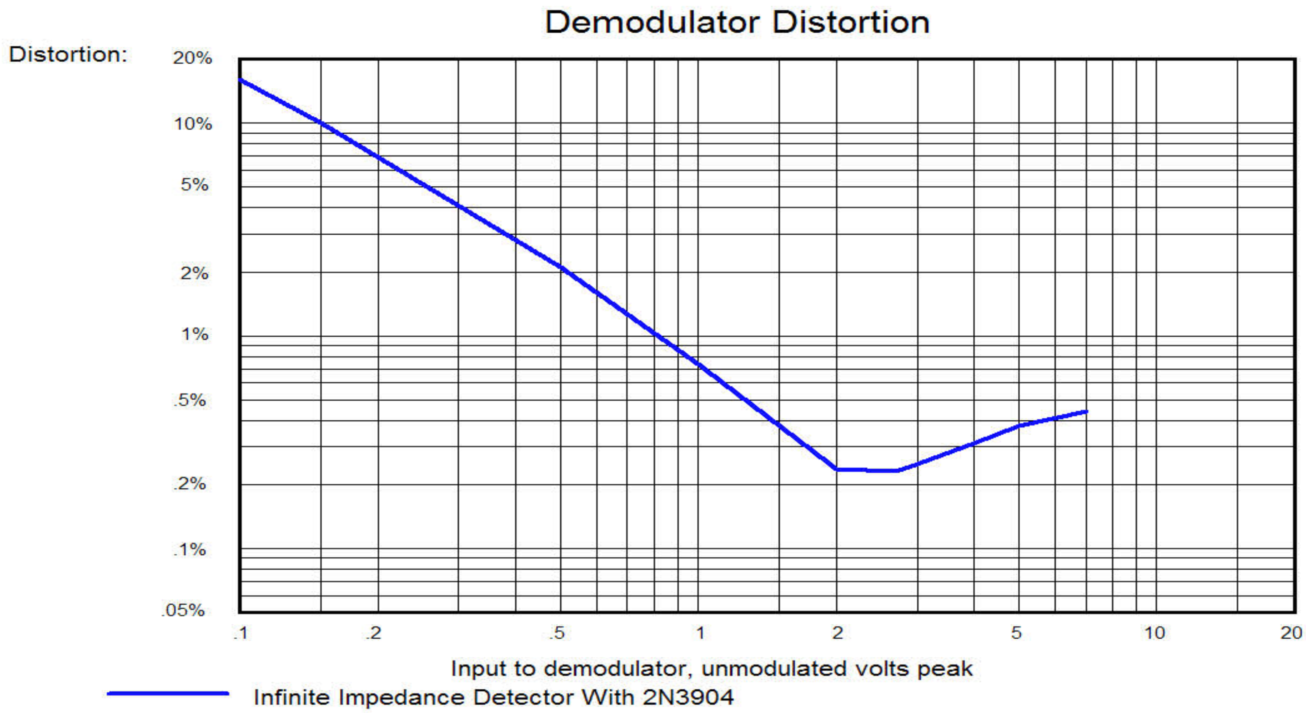

Figure 21 - Distortion plot for infinite-impedance demodulator using bipolar-transistor

The Vacuum-tube Infinite-impedance Demodulator

A vacuum tube version of the infinite impedance demodulator is shown in Figure 22.

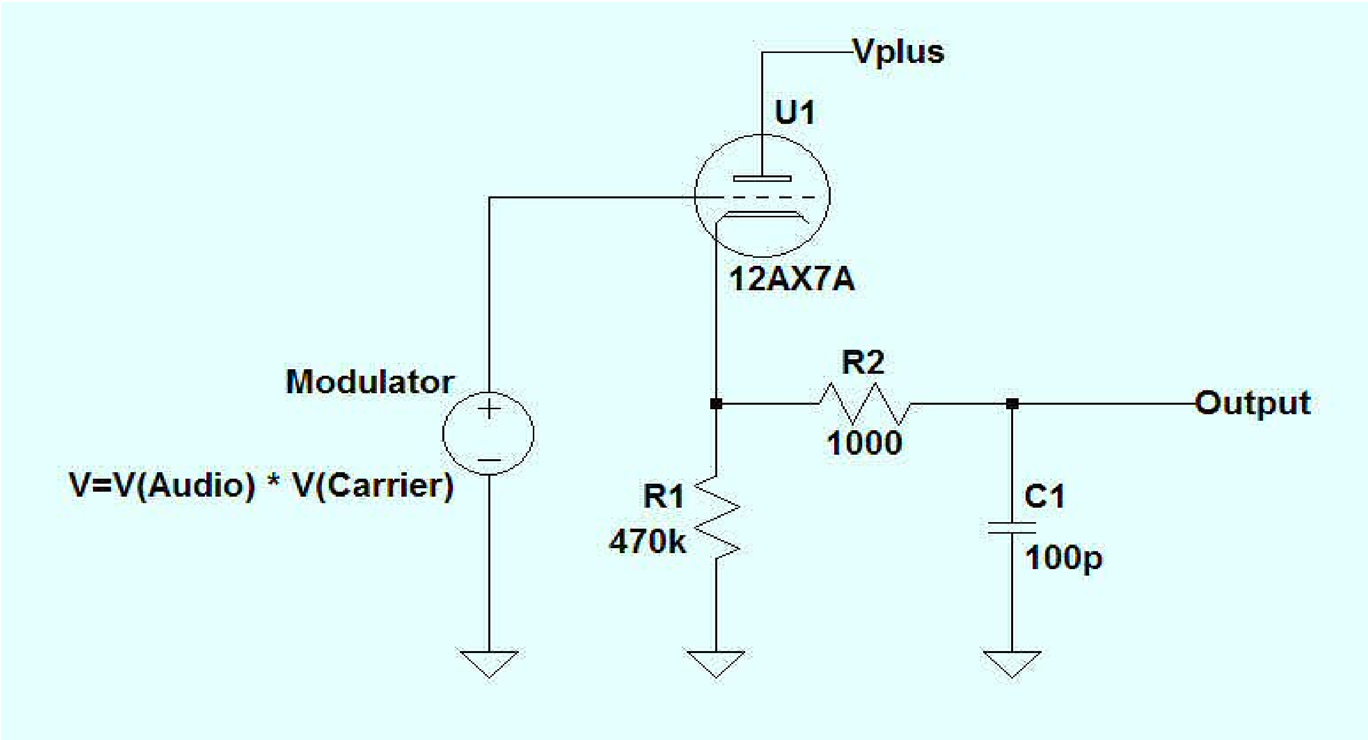

Figure 22 - Vacuum-tube infinite-impedance demodulator

Spice runs revealed that the high-mu tubes produced distinctly better distortion figures than the lower-mu tubes used in the early days of this circuit. As with the transistor version of this circuit, a resistor has been placed between the cathode and the load capacitor C1 to preclude oscillation.

The output waveform of this circuit is shown in Figure 23.

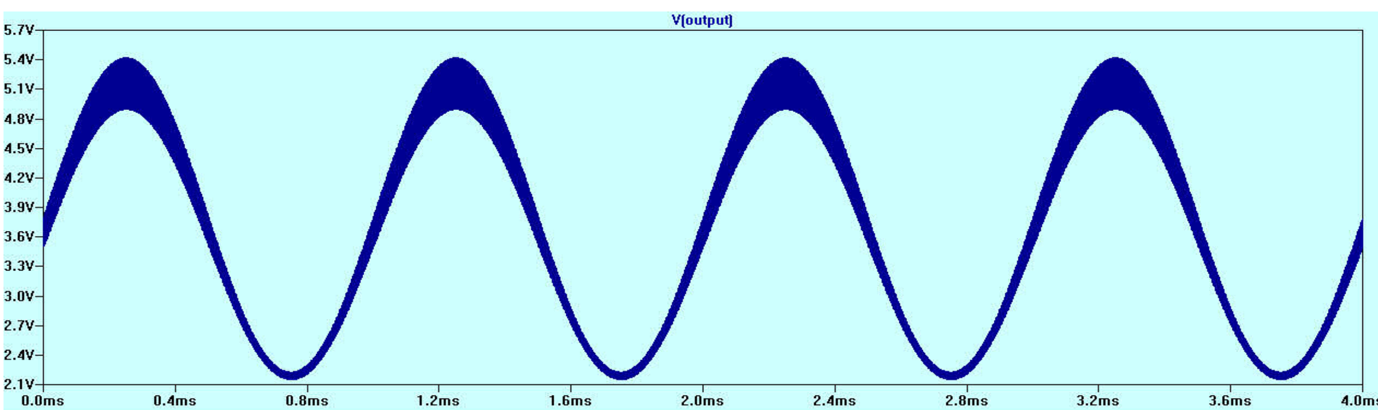

Figure 23 - The output of the circuit in Figure 22.

This circuit should be loaded lightly; it is still of a high-impedance nature.

The distortion of this circuit is shown in Figure 24.

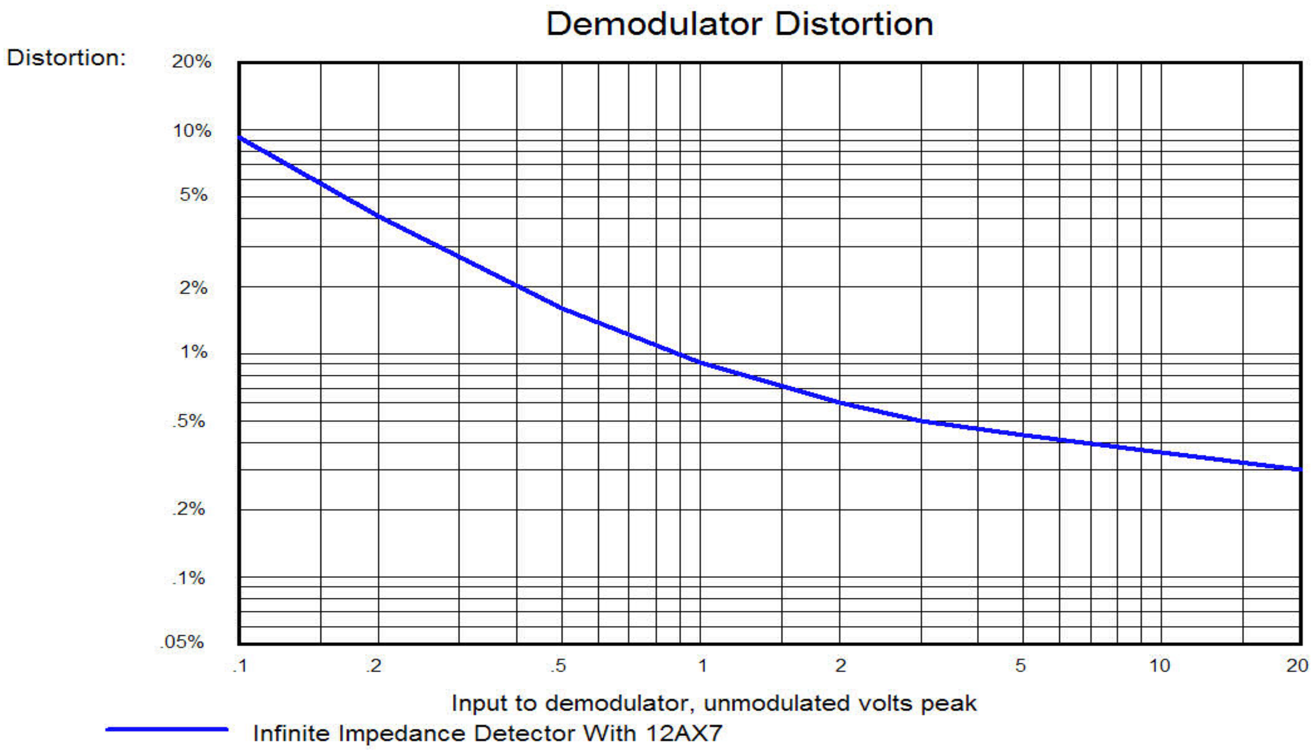

Figure 24 - Distortion of the vacuum-tube infinite-impedance demodulator

This circuit is shown without an AC load. The device is of a high impedance and should not be loaded.

The Selsted-Smith Demodulator

The original Selsted-Smith circuit is shown in Figure 25.

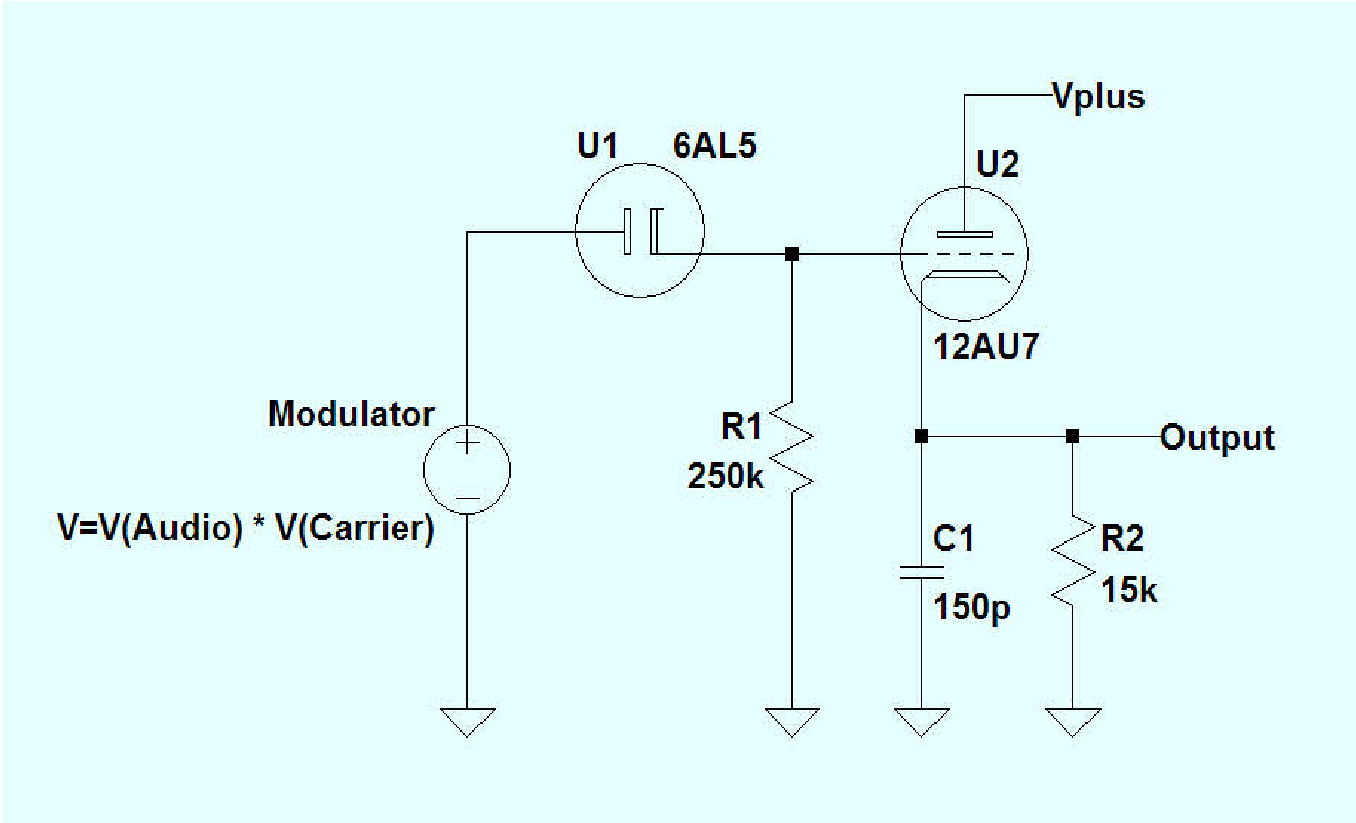

Figure 25 - The original Selsted-Smith demodulator

Exploration using Spice revealed that this circuit can be improved upon by increasing the value of R1, and by increasing the value of R2 (and reducing the value of C1) and by using a high-mu tube.

The signal delivered by this demodulator is shown in Figure 26.

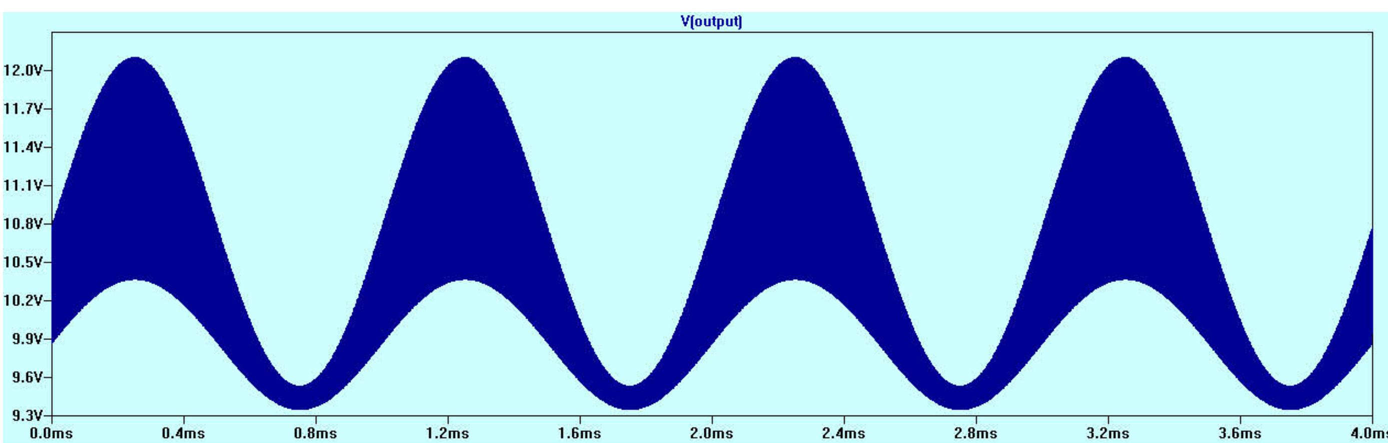

Figure 26 - Output from the original Sellsted-Smith demodulator

The changes alluded to will reduce the level of the residual RF as well as reducing distortion. Distortion of the original as a function of input signal level is shown in Figure 26.

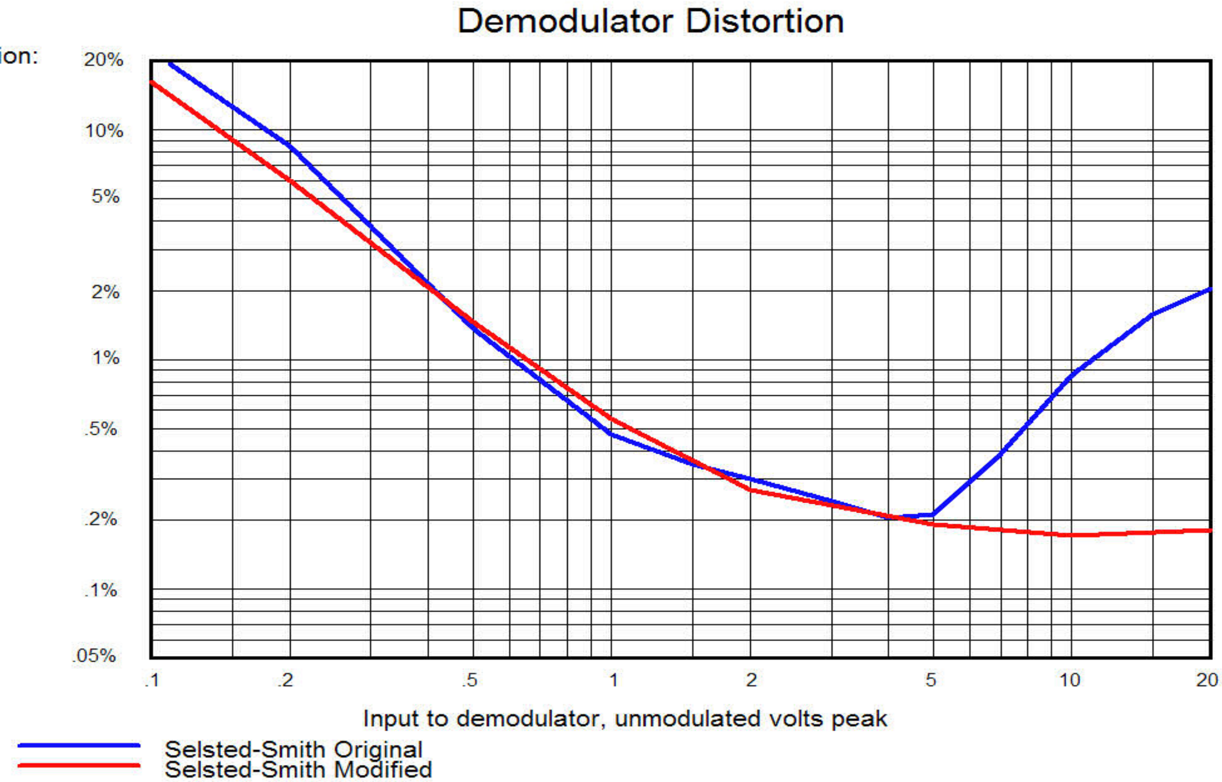

Figure 27 - Distortion of the original Sellsted-Smith demodulator (upper, in blue) and as modified (below, in red).

The "All American Five" Radio

A receiver commonly called the All American Five was designed in about 1945 or thereabouts. It was a successful and economical design that lasted essentially intact for two or more decades with little alteration. The most significant change was in the tube lineup, changing from large 8-pin octal tubes to 7-pin miniatures, probably to enable a smaller unit for the consumer. From our viewpoint this change was transparent.

The Spice-equivalent of the demodulator in this highly-regarded radio is shown in Figure 28.

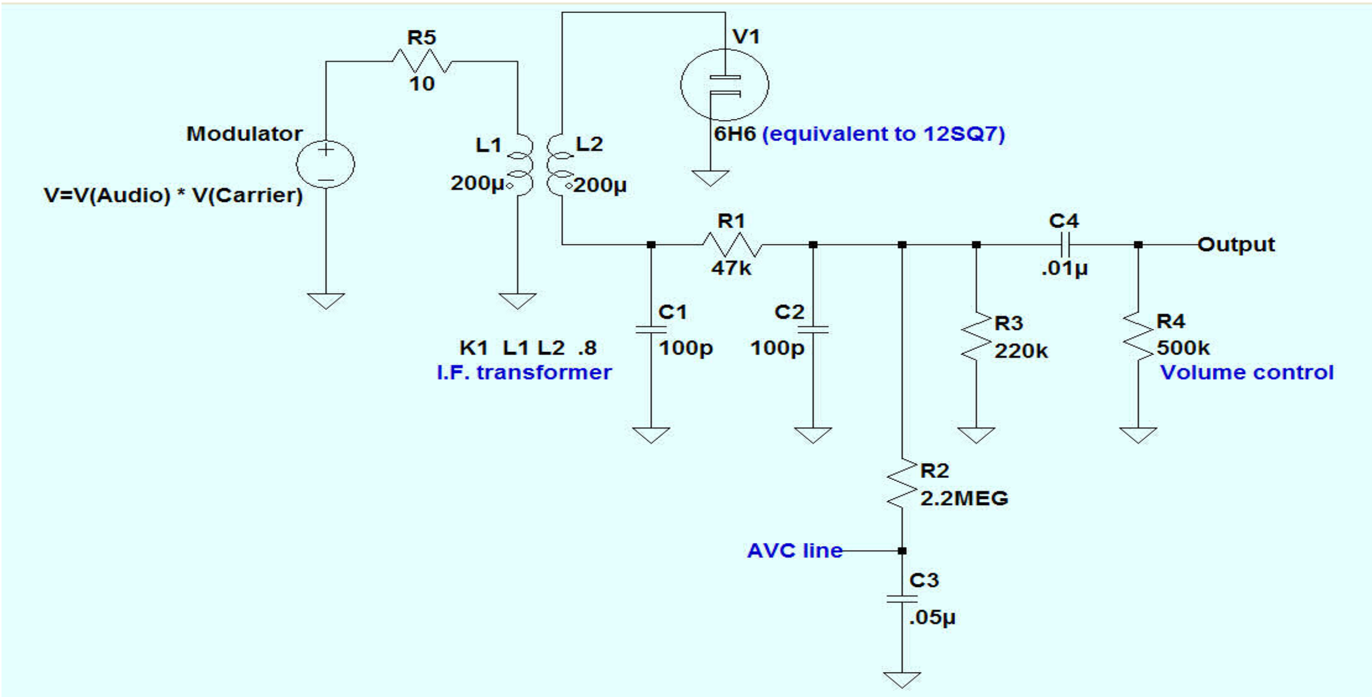

Figure 28 - The All American Five receiver's demodulator

When our standard signal is applied we see the output of the demodulator proper in Figure 29.

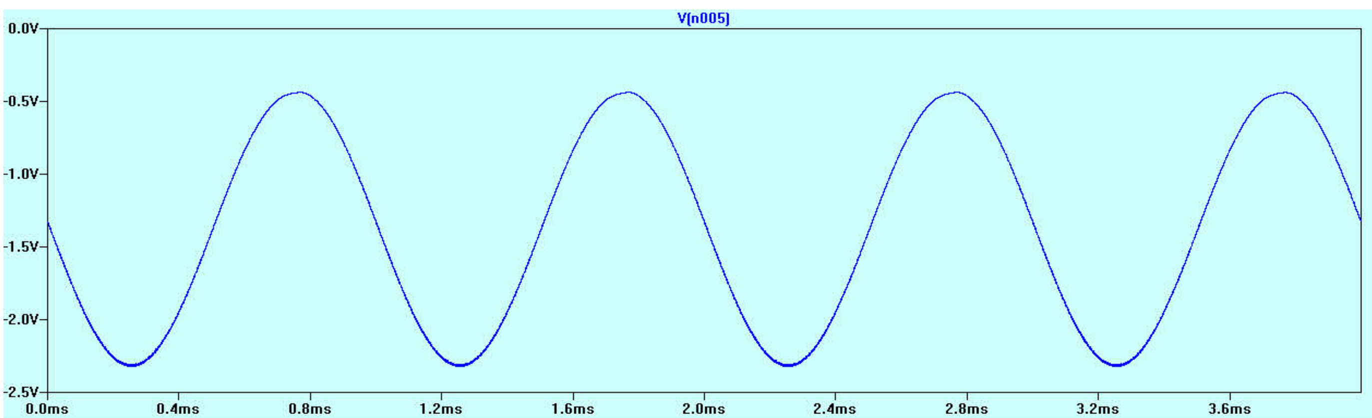

Figure 29 - Output from the All American Five demodulator

This is the waveform as seen at the junction of R1 and C2.

The distortion of this demodulator as a function of input signal level is shown in Figure 30.

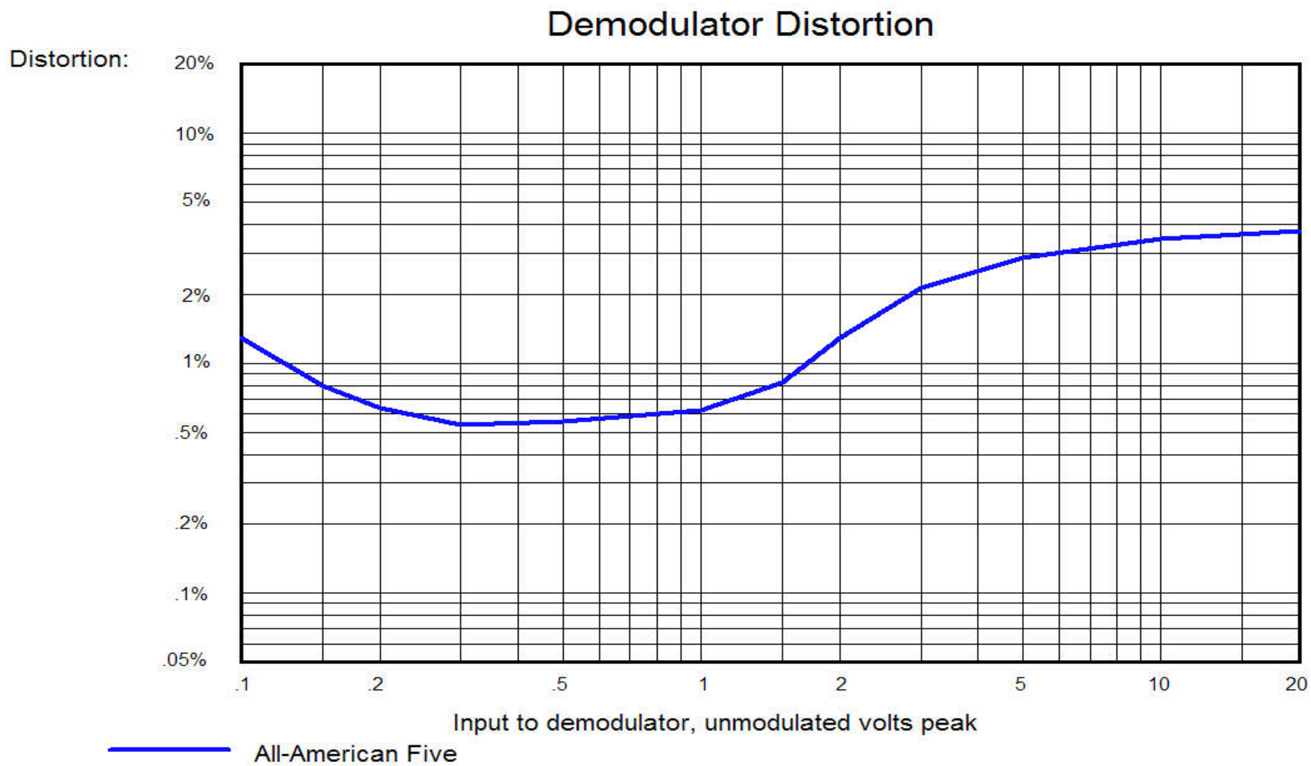

Figure 30 - Distortion as a function of signal level

The ARC-5 Military Radio

Designed in the late 30s and put into production about 1940 was a series of compact radio receivers intended to use in aircraft for short-range communications. The Spice equivalent of the demodulator circuitry in these radios is shown in Figure 31.

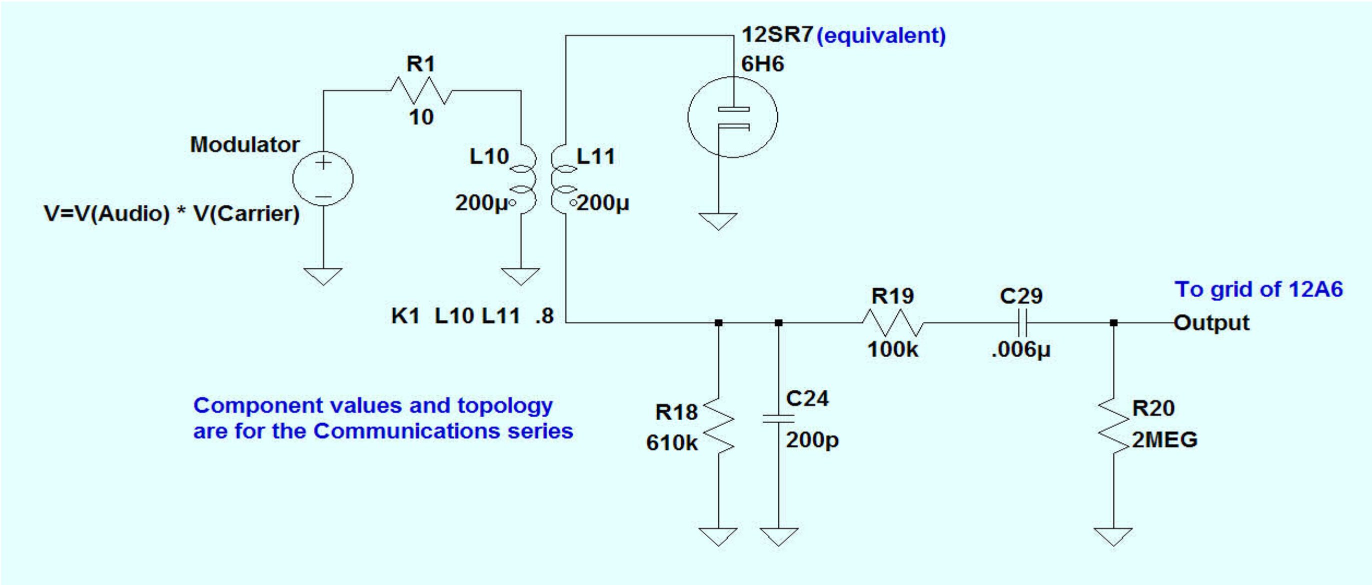

Figure 31 - The ARC-5 receiver's demodulator

The schematic of the ARC-5 receiver's demodulator as shown is for the Communications series (1.5 MHz and above). The lowest-frequency receiver, used only for navigation in the 200 to 500 kHz range, had very bad slew-rate distortion brought about because of the need to reduce the feedthrough of the 85 kHz I.F. signal.

The output of this demodulator is shown in Figure 32.

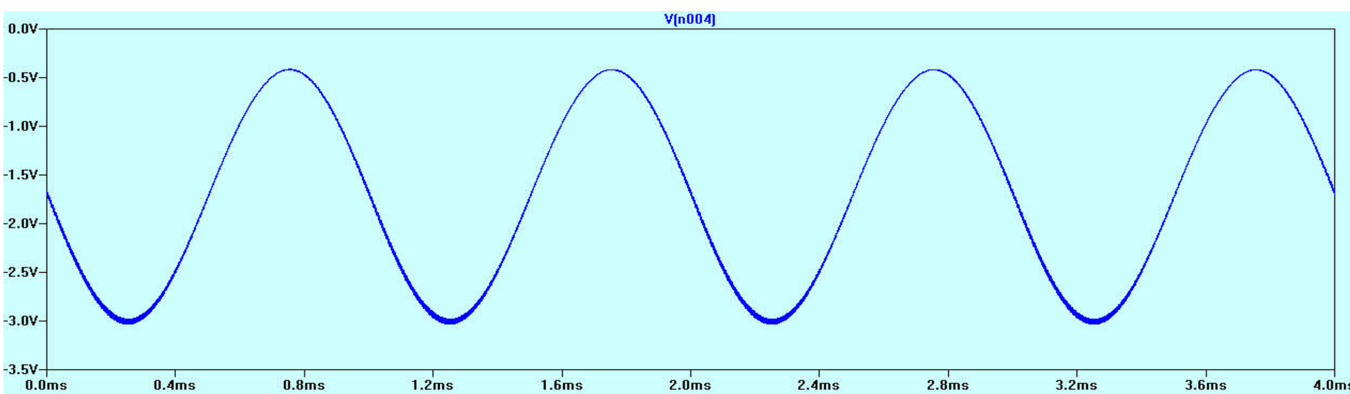

Figure 32 - The output of the ARC-5 receiver's demodulator

This is a nicely-behaved demodulator. The distortion as a function of input signal level is shown in Figure 33.

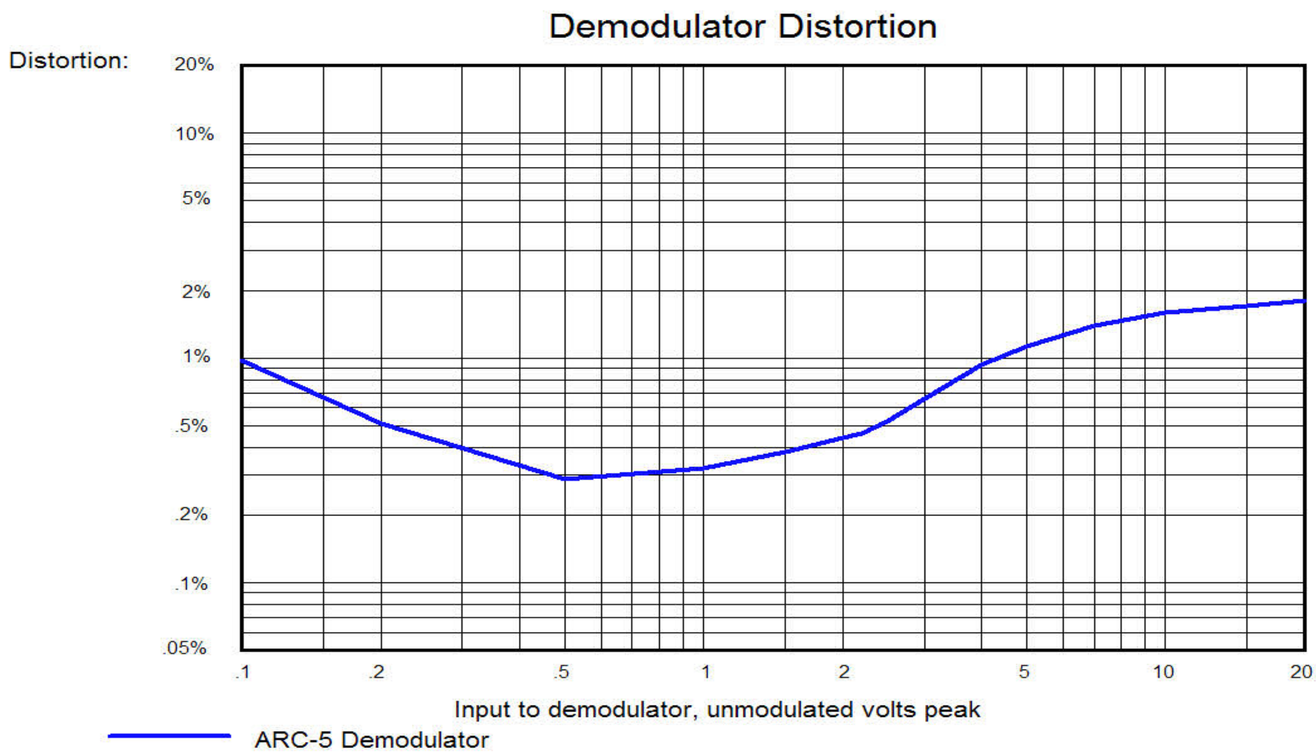

Figure 33 - Distortion of the ARC-5 demodulator

The AVC voltage is developed from a separate (earlier) point in the circuit. As a result it does not have any effect on audio linearity.

The BC-348 Military Radio

In a manner similar to the ARC-5, the BC-348 was designed in the late 30s but had design improvements for the intended application of long-range communications.

The Spice equivalent of this radio's demodulator is shown in Figure 34.

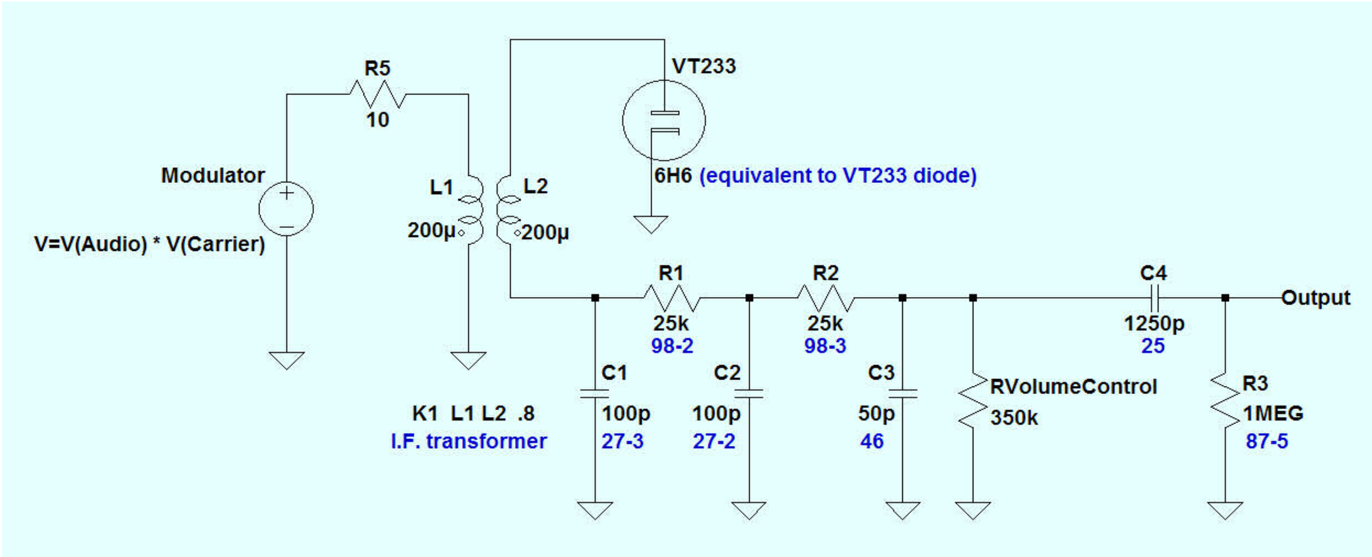

Figure 34 - The demodulator circuitry for the BC-348 radio

The well-behaved output from this demodulator is illustrated in Figure 35.

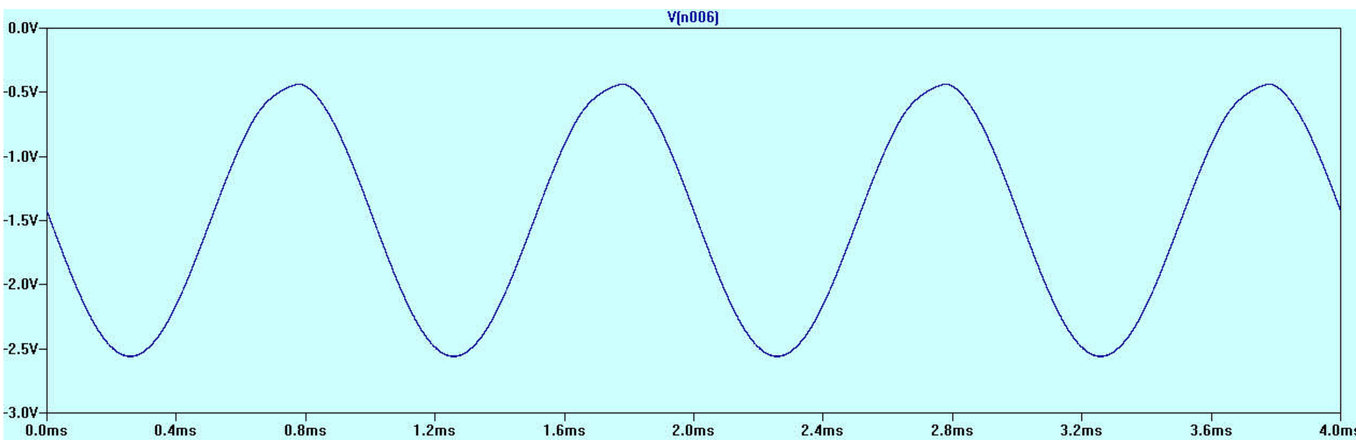

Figure 35 - Output of the BC-348 demodulator

A small amount of negative-peak (top of this waveform) distortion can be seen.

Figure 36 shows the distortion as a function of signal level.

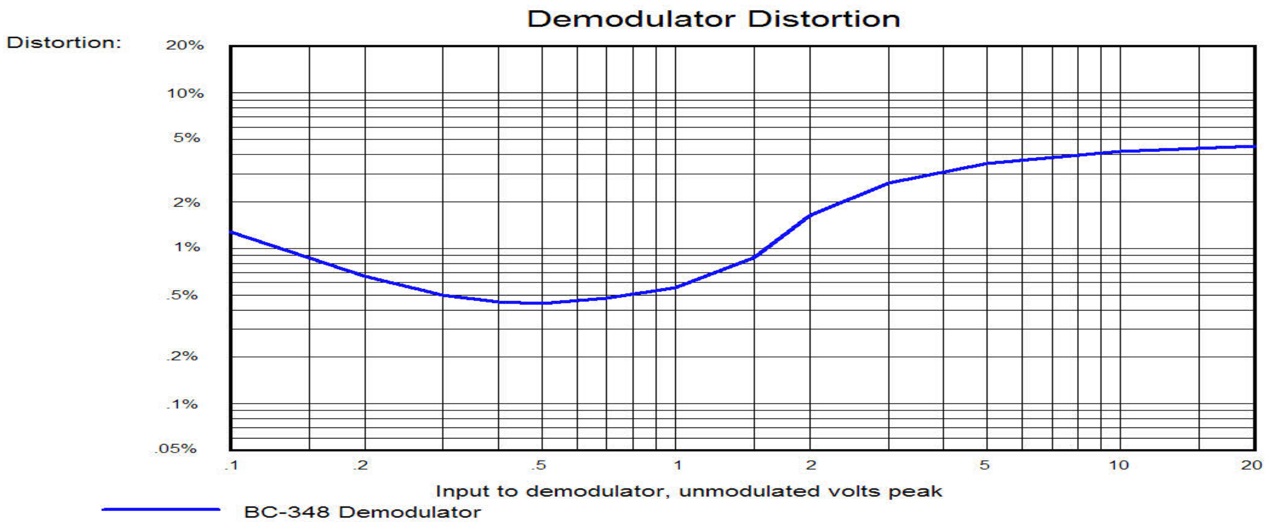

Figure 36 - Distortion versus signal level for the BC-348 demodulator

The Collins 51J4 Receiver

A popular all-band radio receiver was the Collins 51J4. The Spice equivalent for the demodulator in that receiver is shown in Figure 37; as shown the noise limiter is out of the circuit.

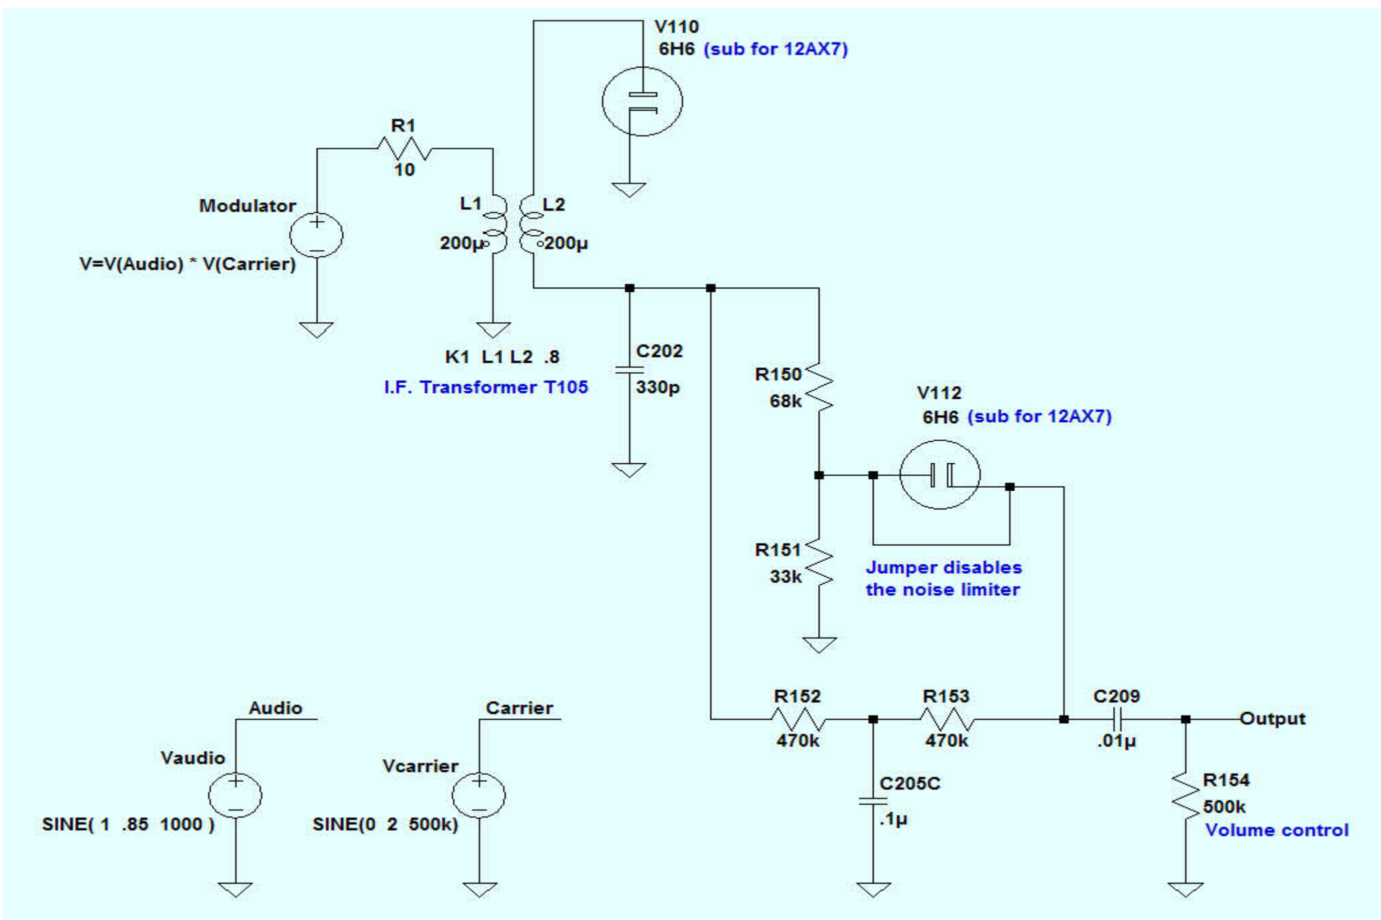

Figure 37 - The Collins 51J4 demodulator

The demodulator output (as seen at the junction of C202 and R150) is shown in Figure 37.

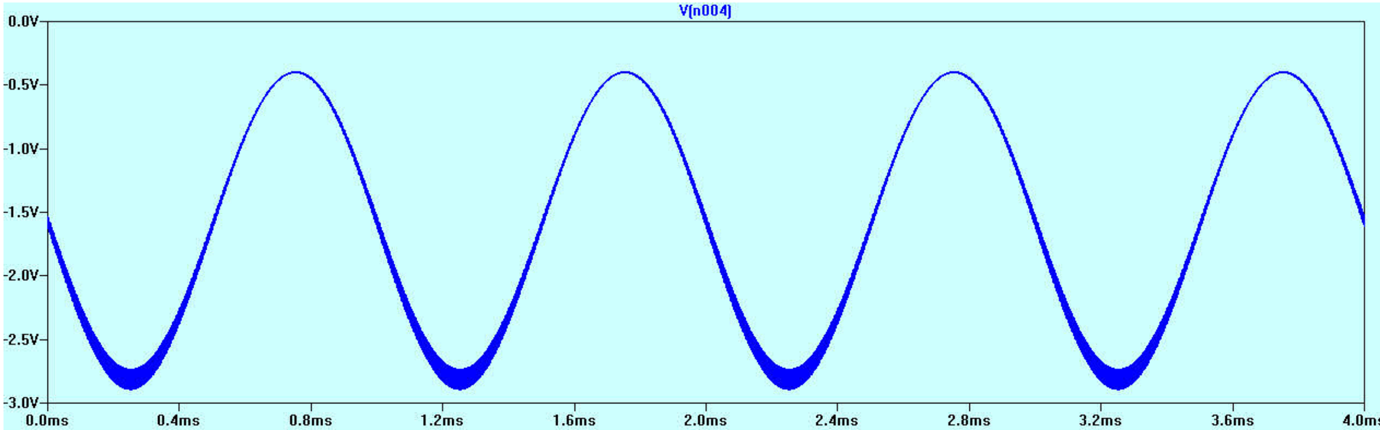

Figure 38 - The 51J4 demodulator output

The noise-limiter has been switched off because it has a surprisingly serious impact on distortion.

The signal-strength meter on the 51J4 (whose deflection is a function of AVC voltage) does not move with modulation because in that receiver a separate demodulator is used for AVC purposes. The design of the AVC demodulator is such that there is very little disparity between the AC and the DC loads.

Distortion of this demodulator as function of input signal strength is shown in Figure 39.

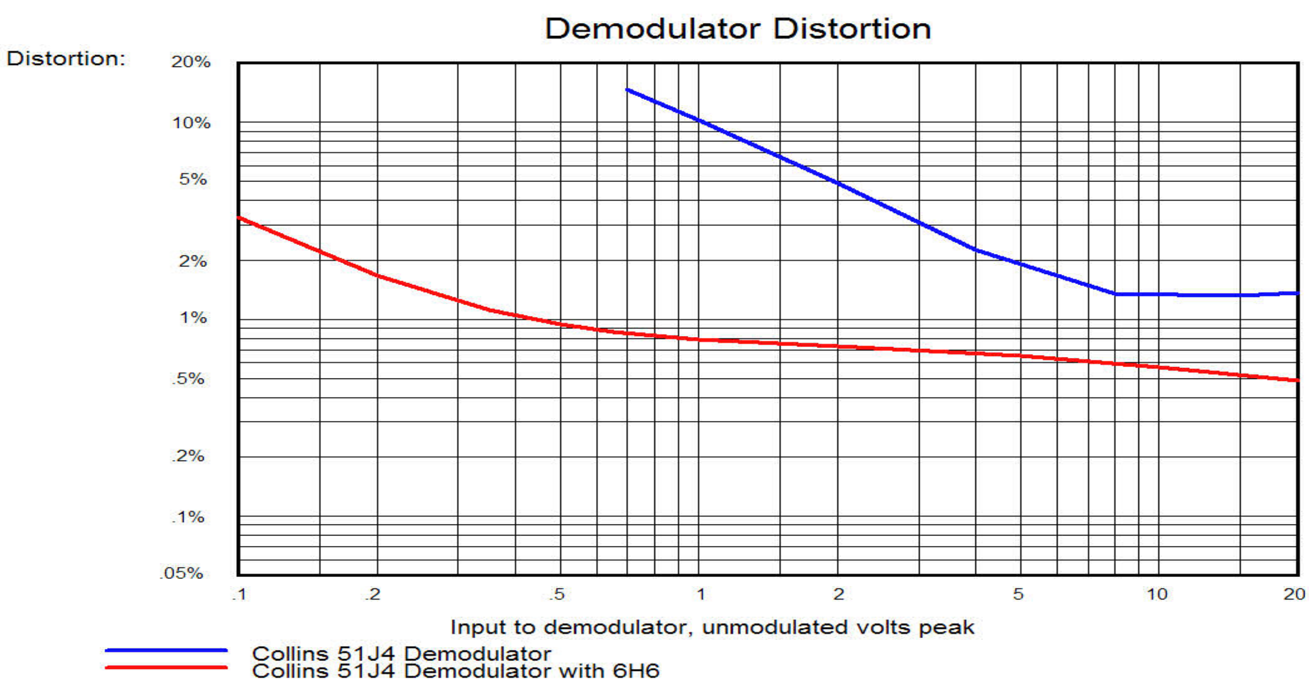

Figure 39 - Distortion of the Collins 51J4 demodulator

The upper (blue) trace is the distortion using the 12AX7 diode-connected for demodulation using Spice. The lower trace shows how the distortion is dramatically improved by the simple expedient of using a 6H6 tube instead. The blue trace is simply not believable; I suspect the 12AX7 model is not correct at the signal levels involved. (That model appears to be quite nicely-behaved at the usual voltage levels that are commonly encountered in audio amplifier design.)

The Hammarlund HQ129X receiver

Another popular all-band receiver was the HQ129X. The Spice-equivalent of this receiver's demodulator is shown in Figure 40.

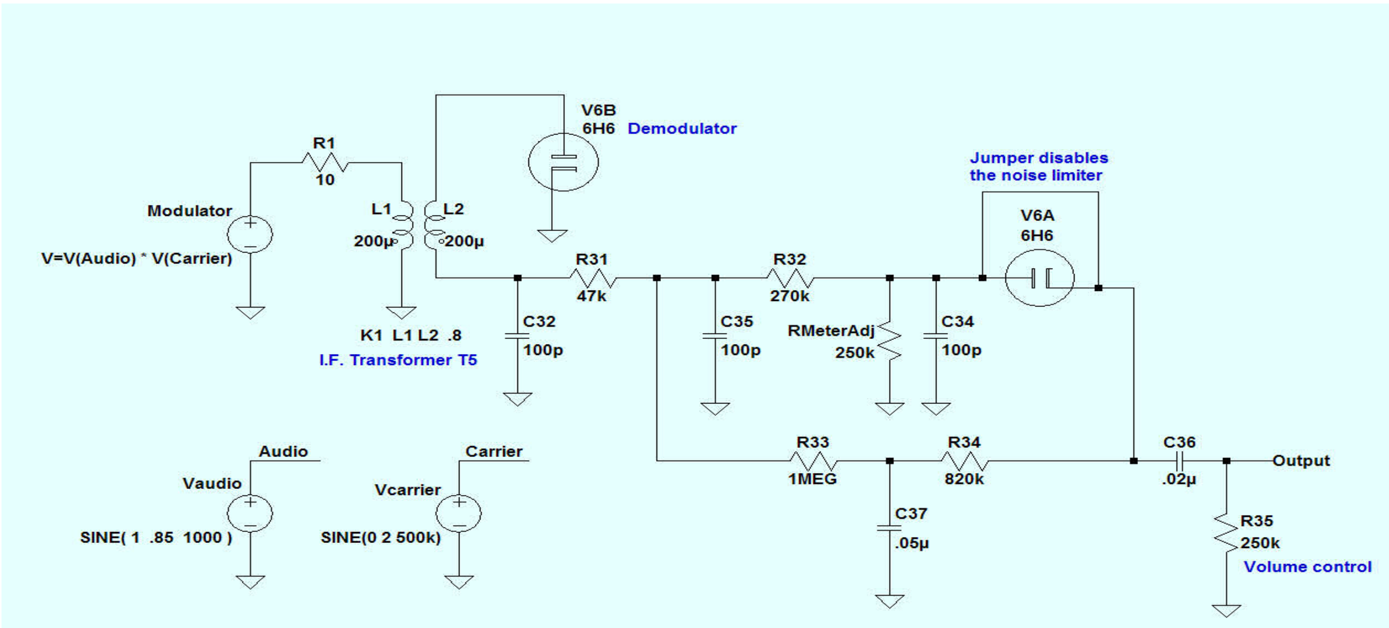

Figure 40 - The Hammarlund HQ129X demodulator

The output from this demodulator is shown in Figure 41.

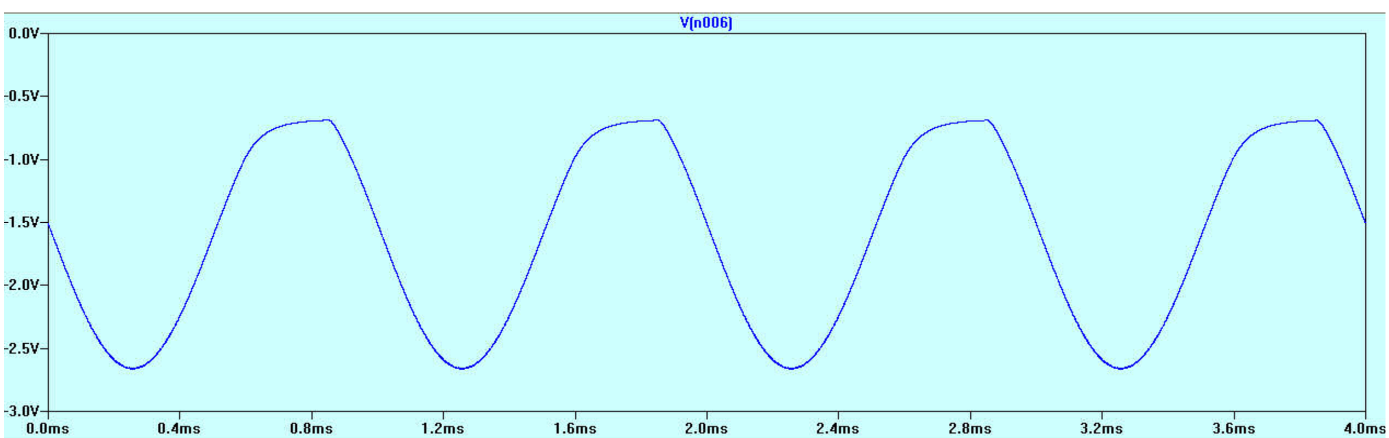

Figure 41 - HQ129X demodulator output

The RF feed through has been reduced to near-vanishing levels. The effect of the AC load on the demodulator is quite obvious. The noise limiter has been switched off.

The distortion plot is shown in Figure 42.

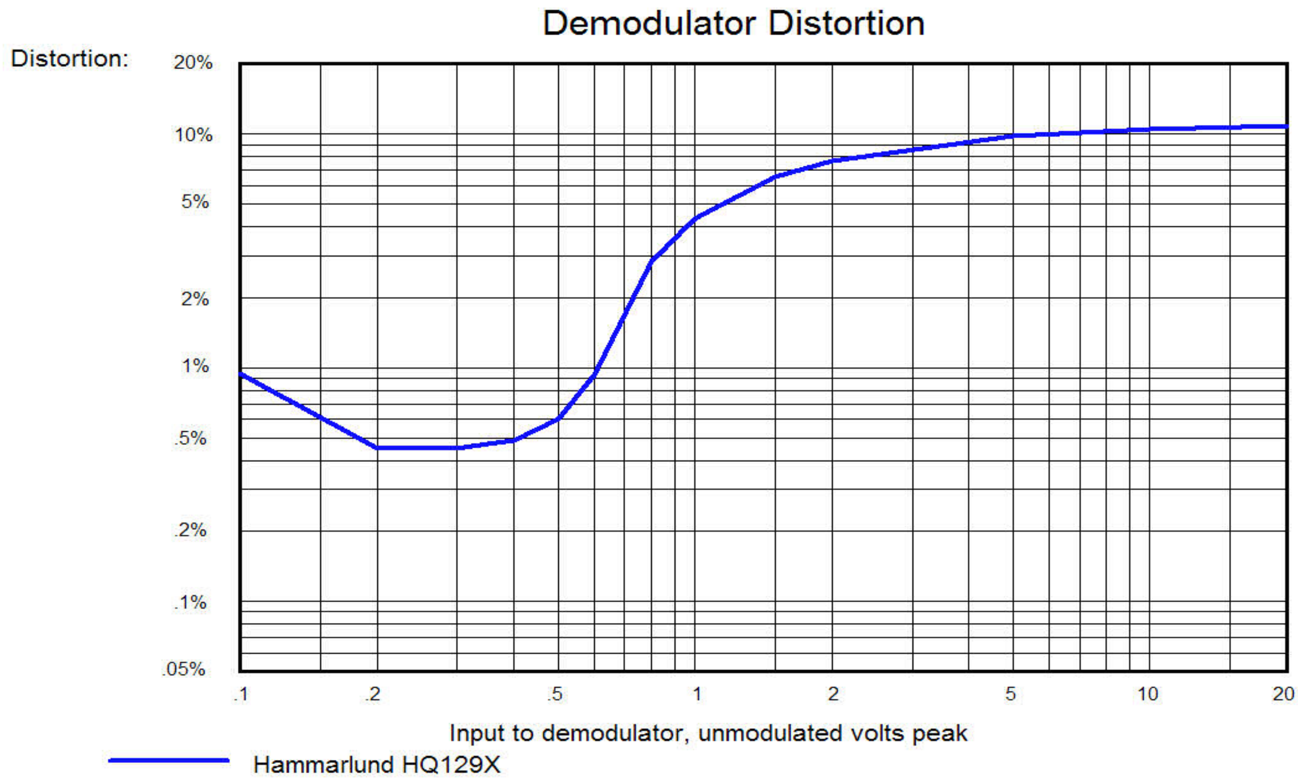

Figure 42 - Distortion plot for the Hammarlund HQ129X

General Radio 1931 Modulation Monitor

The General Radio Model 1931 modulation monitor was designed in 1935. It minimizes the problems seen above by using good design techniques along with high signal levels. This unit was designed to be attached to a radio transmitter and so it operates at much higher signal levels than those encountered in an ordinary radio receiver. The no-modulation signal level in this device is 20 volts as measured at the junction of inductor L5 and resistor R3 in Figure 43.

Shown in Figure 42 is the Spice equivalent circuit for this demodulator.

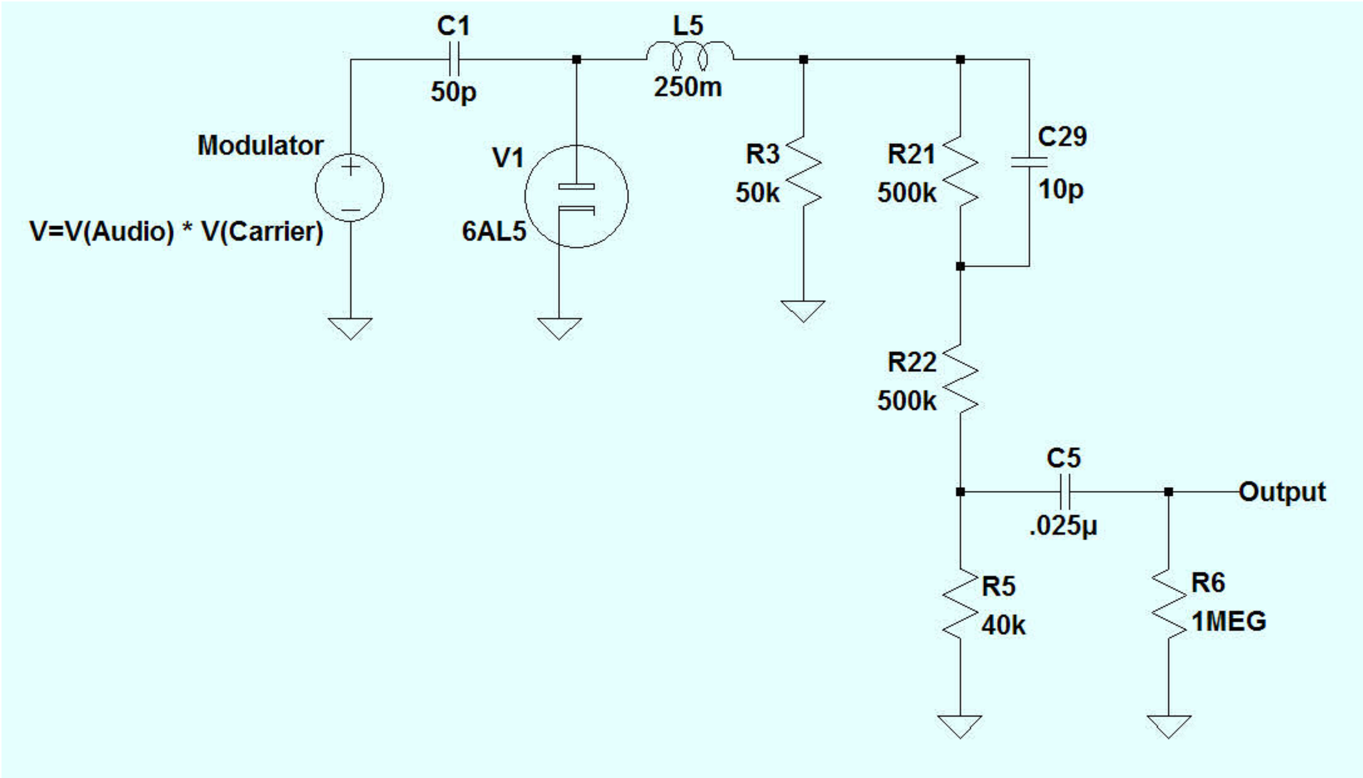

Figure 43 - The General Radio 1931A modulation monitor demodulator

The demodulated signal is quite clean, even at 100% modulation as seen in in Figure 43.

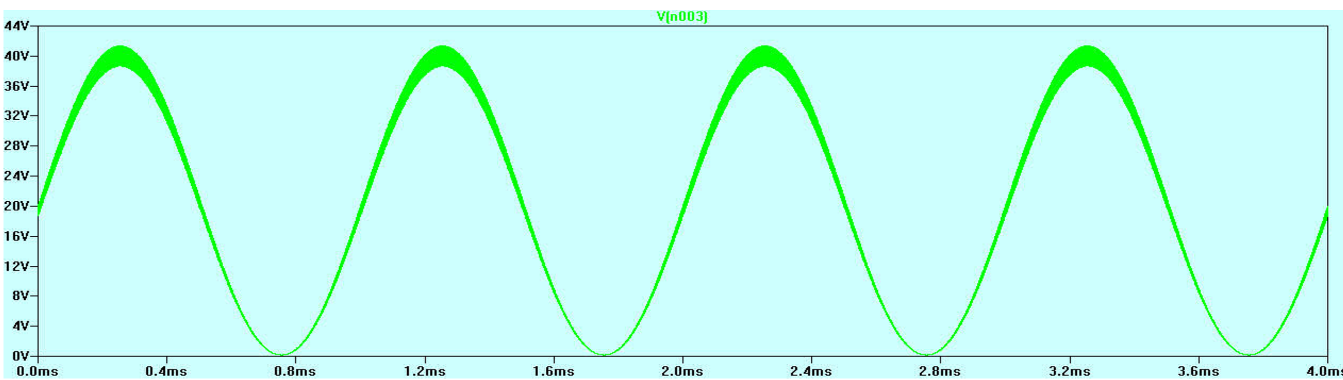

Figure 44 - Demodulated output from the General Radio 1931A modulation monitor

This waveform shows little visible distortion even at 100% modulation. (FFT distortion measurement shows about 0.1% THD.) The only filtering at this point is the inductor L5 operating as a simple lowpass filter in conjunction with resistor R3. The output from the demodulator proper is decoupled by resistor R5 operating with the sum of resistors R21 and R22. The residual RF appearing on this waveform is filtered later in the circuitry.

Belar AMM-2 Modulation Monitor

The Belar AMM-2 modulation monitor minimize the problems seen above by application of modern design techniques. The slew-rate problem is dramatically reduced by returning the exponential decay of the diode demodulator to a power supply voltage instead of to ground. The negative-peak clipping due to an AC load is eliminated by direct-coupling the output of the post-detection lowpass filter to a buffer amplifier. Extracted from the output of that buffer are the DC component to read carrier level and the AC component to read modulation. A Spice approximation, with a few libertties taken, is shown in Figure 45.

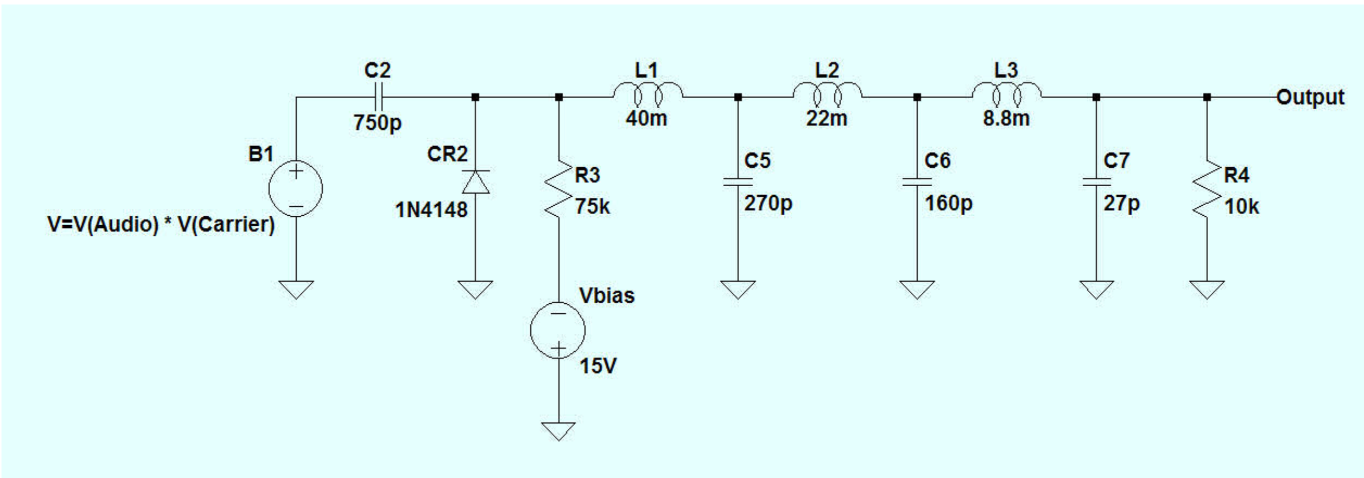

Figure 45 - An approximation of the AMM-2 demodulator

Components L1 through C7 form a linear-phase (non-overshooting on transients) lowpass filter to remove the residual carrier from the demodulator proper. R4 is the sole load on the filter; the filter output is direct-coupled to a high-input impedance buffer.   The output from the filter is shown in Figure 46.

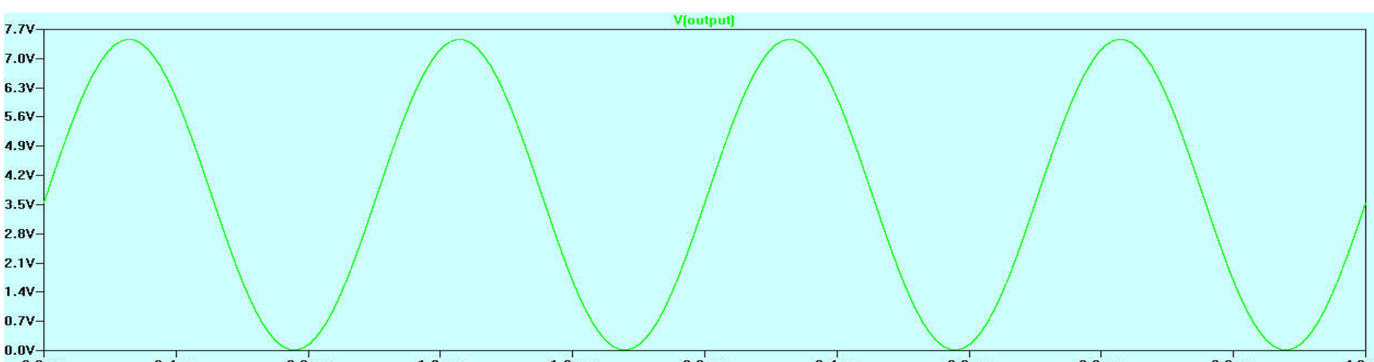

Figure 46 - Output from the AMM-2 demodulator

This is a well-behaved waveform even at 100% modulation as shown.

Both of those modulation monitors are designed to operate at only one signal level.

The Tonne Demodulator

A demodulator was developed whose prime attribute was very good stability with temperature changes and which turned out to have rather low distortion in the demodulated output. The schematic of this design is shown in Figure 47.

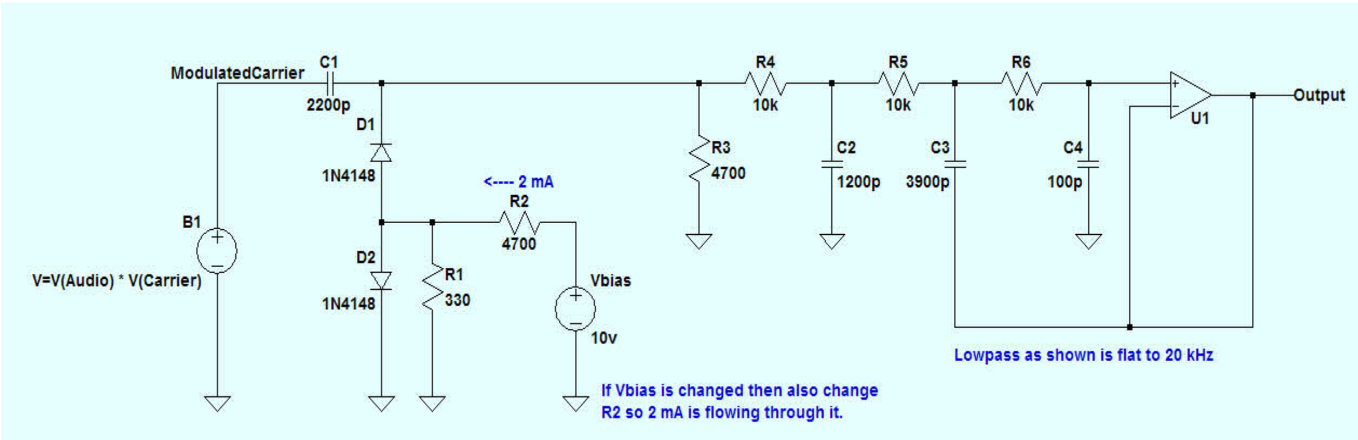

Figure 47. - The Tonne demodulator

The lowpass filter as shown is quite flat to 20 kHz and has the capability of driving a low-impedance load with no effect at all on the demodulator's performance.

The output from this circuit is shown in Figure 48.

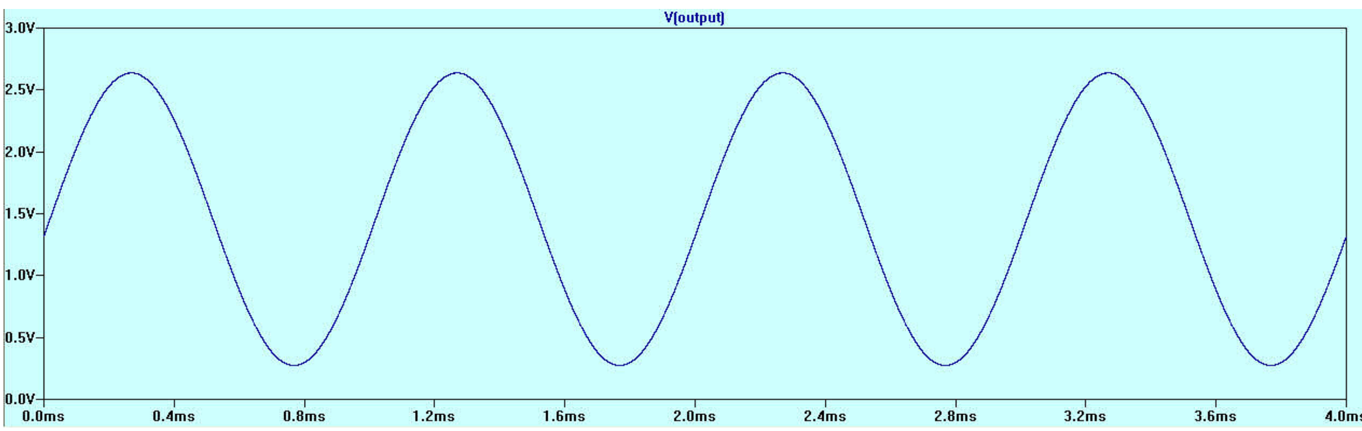

Figure 48 - Output of Tonne demodulator

The distortion from this demodulator as a function of signal level is shown in Figure 49.

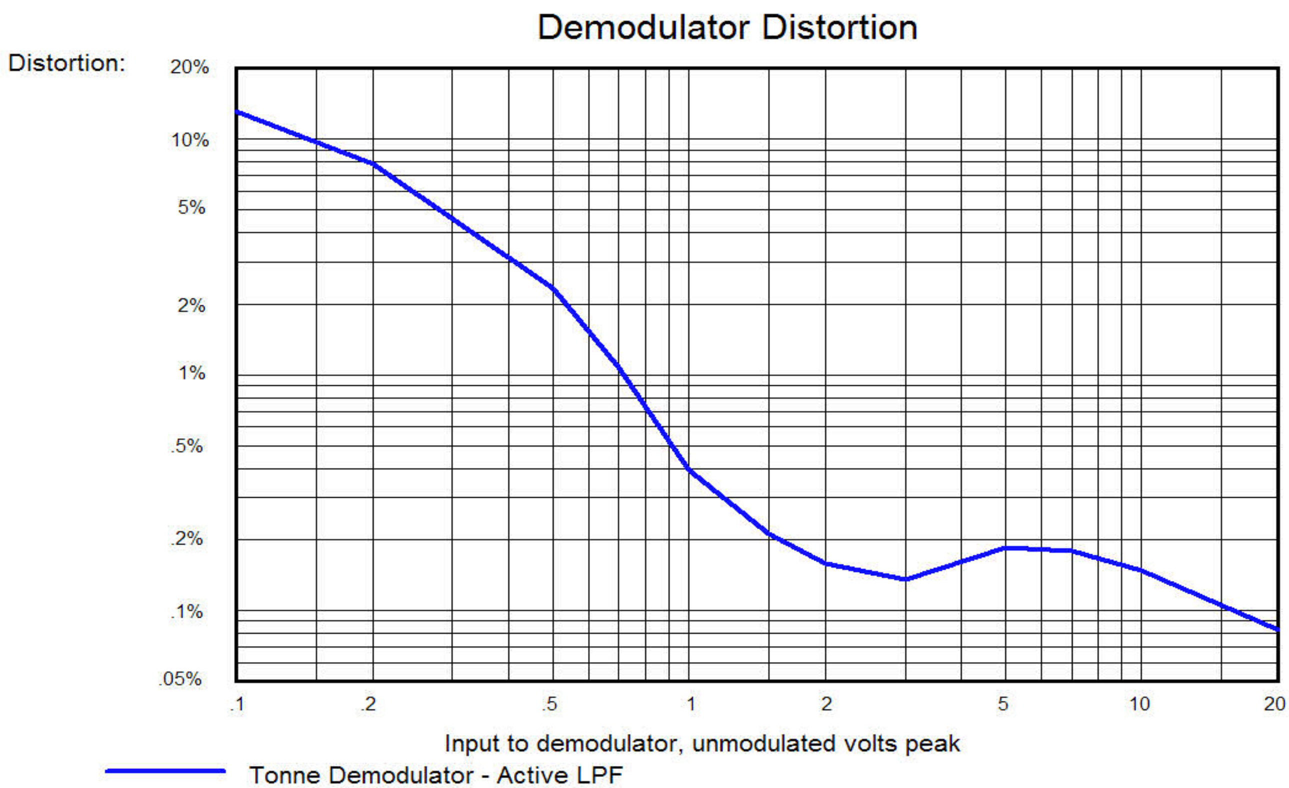

Figure 49 - Distortion as a function of signal level for the Tonne demodulator

The object of this writeup has been to highlight the problems that can be encountered with the ordinary envelope demodulator. If the reader is serious about designing a low-distortion demodulator for use in AM radio receivers then the information presented here should prove useful.

Comments, suggestions and feedback are of course welcome!

- Jim Tonne   W4ENE   Tonne Software   www.TonneSoftware.com