Introduction to current-drive audio amplifiers for audio-frequency induction-loop ('hearing loop') systems (Rev. 1)

© J. M. Woodgate B.Sc.(Eng.),C.Eng., MIET, MIEEE, FAES

Abstract: Some basic electrical theory is first summarized, to aid consistency of understanding. The differences between conventional voltage-drive amplifiers and current-drive amplifiers are explained using ideal amplifiers. Examples are then given of the use of freely-available integrated audio-amplifier parts in current-drive configurations. The paper is intended to be as accessible as possible, without requiring advanced technical knowledge.

1 History

Most of electronics is based on applying an input voltage to some circuit and getting an output voltage that is bigger, smaller or altered in a desirable way. The idea of apply an input current to something probably surfaced in the early days of 'electronic' television (camera tubes and cathode-ray tubes, not scanning discs). Much use was made of electromagnet coils for focusing and deflecting electron beams. These need feeding with a carefully-controlled current, even though their resistance changes a lot as they heat up. If a constant voltage were applied, then, the current would decrease as the coil heats up. What is needed is a *constant-current supply*, which provides the necessary current even though the coil resistance changes. Luckily, a pentode valve/tube provided a constant current, proportional to the input voltage, without any special circuitry beyond keeping the screen (g2) voltage constant.

While is it always uncertain who first did anything, it was Leon Pieters of Ampetronic who took the 'constant-current' concept from television technology and applied it to hearing loops. In this case, it's not needed to combat the effect of heating of the loop wire but the effect of the loop inductance on the frequency response of the magnetic field generated by the loop.

2 Basics

We can't get away without understanding a bit of basic electronics, because without it the words we have to use won't mean anything. So, we begin with a *battery*, which converts internal chemical energy into electricity. It has a *voltage* of 9 volts (after the Italian scientist Volta). Voltage is the electrical pressure that drives the electricity round a *circuit*, which is a network, including at least one closed loop, of wires and *components.*

A battery and a traditional filament lamp are connected in a single loop, forming a *series circuit.* Electric *current* flows from the battery to the lamp and back to the battery. There actually is something that flows – electrons, which are sub-atomic particles carrying *negative electric charge*. It's only negative because Ben Franklin, very long ago, had to choose which sort of charge to call *positive* and he (knowing nothing about electrons) made the second-best decision. This flow is called *electric current*, and because electrons have negative charge, the current conventionally flows the other way, from the *positive terminal* of the battery, through the wires and the lamps, back to the *negative terminal* of the battery. Obviously, there must always be a closed loop otherwise no current can flow.

The filament wire in the lamp is very thin indeed and the current flow makes it white hot, or nearly so, giving out light. (But it also gives out far more heat, which is why filament lamps are being replaced by energy-saving lamps, which produce far less heat). The amount of current flow, measured in *amps* (after the French scientist Ampère), is determined by the voltage and the *resistance* of the filament, which is measured in *ohms* (after the German scientist Ohm). The connecting wires have resistance, too, but we make sure they are thick enough to have negligible resistance.

We have arrived at *Ohm's Law,* one of the two most fundamental equations in electricity:

$$\textit{Current} = \textit{Voltage}/\textit{Resistance}$$

The '/' means 'divided by'. We often use symbols rather than words, an in symbols the equation is:

$$I = V/R$$

It may seem strange that the symbol for current is 'I', but the reason is in the dim past. We can also write the equation in two other ways:

$$R = V/I \text{ and } V = IR$$

'IR' means '*I* times *R'.*

3 Alternating Current and Direct Current

The battery produces *direct current* – the current always flows from the positive terminal, through the circuit, to the negative terminal. This is the sort of current that Edison liked but his rival Tesla preferred *alternating current*, in which the current direction regularly reverses. For example, public electricity supplies deliver alternating current which reverses 100 or 120 times a second. Two reversals are called a *cycle* and 1 divided by the time for two reversals is called the *frequency*, which is measured in *Hertz* (after the German scientist Hertz). Alternating current has the advantage that it can be sent over long distances at very high voltages (to minimize power loss in the resistance of the wires) and the voltage reduced locally by *transformers* to much lower values suitable for use in industry and homes.

In a circuit carrying alternating current there may be one or two more 'resistance-like' effects, but they are not quite the same as resistance because they involve storing energy when the current is flowing one way and releasing it when the current reverses. Resistance doesn't involve any storage. For us, the more significant effect is called *inductance,* where energy is stored in a magnetic field created by the current flow. All wires, loops and coils have inductance, which is measured in *henrys* (after the American scientist Henry). In a loop system we want the magnetic field but we do not want the current-opposing effect of the inductance to stop the system working well.

The opposing effect of inductance is called *inductive reactance* and is proportional to the frequency of the current, with a factor of 2π included (π, pi, is the ratio of the circumference of a circle to its radius, approximately 3.14):

$$X_L = 2\pi fL$$

X_L is the inductive reactance (measured in ohms)

f is the frequency in Hertz

L is the inductance in henrys

The other type of opposition is called *capacitance,* measured in *farads* (after the British scientist Faraday). The opposing effect is called *capacitive reactance* and is *inversely* proportional to the frequency, with that 2π factor again:

$$X_C = 1/(2\pi fC)$$

X_C is the capacitive reactance (measured in ohms)

$2\pi f$ has a name of its own – *angular frequency* – and symbol ω (lower-case omega). It is measured in radians/second, but don't bother about that.

C is the capacitance in farads. However, a 1 farad capacitor is a rare beast; most practical capacitors have values in microfarads μF (one millionth), nanofarads nF (one billionth) and picofarads pF (one trillionth).

Resistors, inductors and capacitors are *components,* and you can buy them at Radio Shack (and other fine stores, of course). Although resistance, inductive and capacitive reactance are all measured in ohms, we can't add them up directly to find the total effect. This is because reactances introduce *phase-shift* between the alternating voltage and current. Inductive reactance delays the current, while capacitive reactance delays the voltage (usually described as the current 'leading' the voltage, which looks a bit strange).

Think of one cycle of AC voltage or current (two changes of direction) as two marching paces: left, right. Inductors and capacitors make the current break step with the voltage, while marching at the same pace (the frequency). For reasons embedded in the underlying mathematics, phase-shift is measured as an angle, such as 45°.

If we have a *resistor,* an *inductor* and a *capacitor* in series with a *generator* producing an alternating voltage, the total opposing effect is:

$$Z = \sqrt{(R^2 + X_L^2 - X_C^2)}$$

Z is a new creature, called *impedance*. It is still measured in ohms. This means that we have a modified Ohm's Law for AC circuits, which can be written in any of the three ways:

$$I = V/Z \text{ and } Z = V/I \text{ and } V = IZ$$

You may notice that a particular value of f makes X_L and X_C equal, so they cancel out and we are left with

just *R*. This effect is called *resonance* and it's very important in electronics in general but not in hearing loop systems.

4 Ideal generators and amplifiers

We have to start with ideal things because if we start with real things the explanations get too complicated. Instead, we find out what ideal things do and then make any necessary corrections for real things; often they are too small to bother with.

An *ideal voltage generator* maintains its voltage constant, whether the current required by a circuit is 1 microamp (one millionth, written 1 μA: μ is 'mu') or 1 megamp (1 million, written 1 MA). This is because it has *zero internal impedance*. An *ideal current generator* maintains a constant current in a circuit, even if that means that the voltage across its terminals is 1 microvolt (1 μV) or 1 megavolt (1 MV). This is because it has *infinite* internal impedance. Internal impedance, or *source impedance* is a very important concept and we use it a lot.

An *ideal amplifier* has an output terminal and at least one input terminal; the type we are most concerned with has two input terminals. The amplifier is powered by magic, so that we can deal with the essential power supplies of a real amplifier as a separate subject. The output voltage of the amplifier is a larger (amplified) copy of the input voltage. With a two-input ('differential') amplifier, one input produces an *inverted* copy of itself at the output (so that the output goes negative when the input goes positive), while the other input produces a right-way-up copy. Not surprisingly, these are termed *inverting input* and *non-inverting input.* Electronics people often say *minus input* and *plus input.* Neither of the inputs draw current from any generator they are connected to; they have infinite internal impedance. The output is an ideal voltage generator – it has zero internal impedance.

The ratio of the output voltage to the input voltage is the *gain*, which may be from 1 to millions. An ideal amplifier with infinite gain is called an *ideal operational amplifier* or *ideal op-amp.* You can buy real integrated circuit (IC, 'chip') op-amps (with gains from about 50000 to several million) at the parts store; some are very low-cost, others definitely are not. For hearing loop technology, there is no need to use costly op-amps.

Some audio amplifier integrated circuits are specialized op-amps, while others do not have all the properties of an op-amp but are very similar. For example, they may have two inputs but the internal impedance of each is, say, 50 kilohms (50 kΩ; k means 'x1000' and 'Ω' is upper-case omega) instead of being infinite.

5 Graphics

We need graphics to show and discuss circuits. These are called circuit diagrams or schematics. To begin with, they will be about ideal components and not necessarily practical circuits that can be built. But they make explanations so much easier than without them. I am using the schematic drawing part of the powerful simulation program LTspice (http://www.linear.com/designtools/software/), provided free *pro bono* by Linear Technology Corporation, because it's very easy to use. But the graphics embedded in this paper are not in LTspice's native format, so you don't need the program yet (maybe later!).

First of all, we need to look at the symbols used in schematics, a selection of which is shown in Figure 1

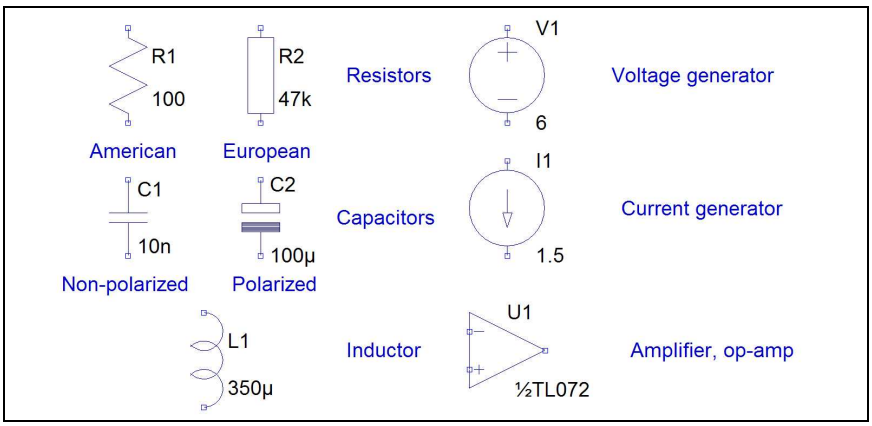

Figure 1 Symbols used in schematics

Components are numbered in sequence (R1, R2..) for reference and the value of each is also shown , with or without the unit (100 for a resistor must be 100 ohms). 'k' means 1000, 'μ' means 1 millionth, 'n'

means 1 billionth. We saw 'MA' earlier, meaning 1 million amps, so 'M' means 1 million, BUT 'm' means one thousandth. There are many more of these 'metric prefixes', which you can find listed on the Web, e.g. at:

http://astro.unl.edu/classaction/tables/intro/si_prefixes.html

but we don't use all of them in electronics.

Generators are also numbered and the voltage or current stated as well. Most generators produce either DC or AC, especially in the sort of example schematics we are soon to study.

6 Audio frequencies and waveforms

There is a sort of mantra that humans can hear sounds of frequencies between 20 Hz and 20 kHz (20 kilohertz, 20 000 Hz). But some are more audible than others, even for people with 'normal' hearing. Unless they are very loud, sound with frequencies less than about 500 Hz are progressively less audible as the frequency decreases. There is a similar effect above 4 kHz. How do these frequencies in Hertz relate to musical tones? The oboe's 'A' note to which the orchestra tunes is 440 Hz. 20 Hz is below the range of most instruments, while 4.186 kHz is often the highest note on a piano. In audio electronics, we have a standard measuring frequency of 1 kHz. So frequency is closely related to musical pitch.

So far, we have simply described AC in terms of regularly reversing current. But the current (or voltage) can vary in an infinity of ways between reversals. This variation is what gives sounds their different qualities. For reasons embedded in the underlying mathematics, the 'simplest' variation is called a 'sine wave', and all sounds, voltages or currents can be built up by adding sine waves of different frequency and amplitude (size) There is an electronic instrument called an 'oscilloscope' that allows us to see on a screen the variations in amplitude over time, called the 'waveform'; this is very familiar from films with any sort of 'scientific' context. A human whistler, ocarinas and some flutes produce nearly pure sine waves, i.e. just a single frequency. But musically, sine waves are bland and boring.

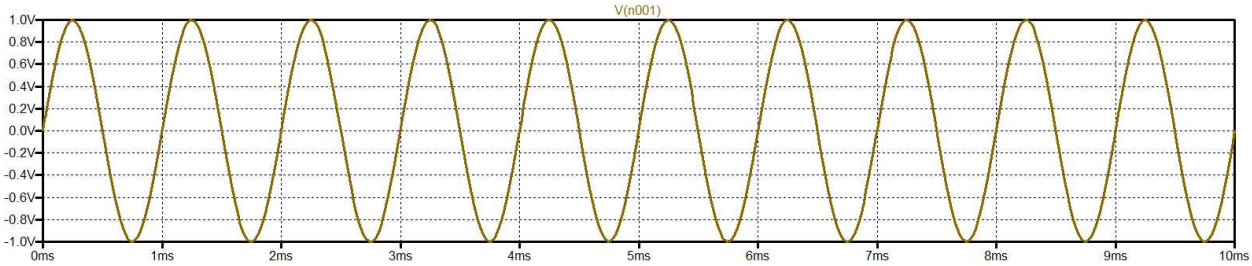

Figure 2 Ten cycles of a 1 kHz sine wave with peak voltage 1 V ('ms' = milliseconds)

We use this sort of graphic a lot in electronics. The horizontal scale (x-axis) is time in this case, and it's often frequency in other cases. The vertical scale (y-axis) is volts, but it might be something else in another case. Don't worry what x and y mean: they are just parts of the labels 'x-axis' and y-axis'.

Speech and music waveforms are very much more wriggly, as shown in Figure 3. They can be separated into many sine-wave components of different, constant frequencies and time-varying amplitudes.

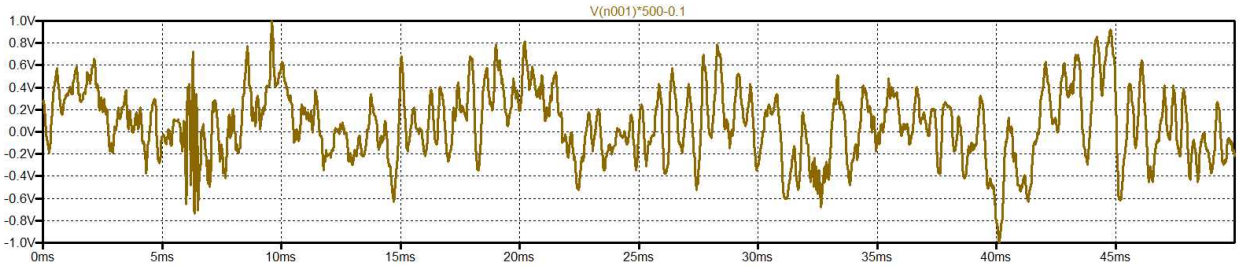

Figure 3 50 ms of a male voice recording with peak voltage 1 V

7 Voltages and currents with different waveforms

Both the sine wave and the speech signals have a peak voltage of 1 V, but the sine wave reaches that every 0.5ms, whereas the speech reaches it very briefly at an interval of 41 ms, and another 50 ms slice of the same recording might not reach it at all. So we can't really compare them on the basis of peak voltage (or current). What works mathematically is to compare them on the basis of their heating effect. When current flows through a resistance, electrical power is converted to heat. There is an equation, Joule's (after the British engineer) Law, which we can write, using *W* for power:

Power = Volts x Amps or W = VI or I = W/V or V = W/I

and then, if we combine Joule's Law and Ohm's Law, we get:

$$W = V^2/R, \text{ or } W = I^2R$$

and we can switch round which symbol is on the left as for Ohm's Law. Note that V and I appear squared. After some mathematics based on those equations, the necessary voltage or current we require is found to be the *square root of the average of the voltage or current squared*, normally averaged over at least one cycle This is called the *'Root-Mean-Square'* (RMS) voltage or current. Note that there is no such thing as 'RMS power' no matter how many times you see it in specifications. It can be calculated, but then so can your age divided by your height – they don't mean anything useful. And what is wrongly described as 'RMS power' is actually average power.

Meters can be designed to measure RMS voltages and currents; there are often called 'true-RMS' meters. Other meters use approximations to RMS and are not accurate on speech, music and noise signals.

8 Amplifiers (at last!)

In discussing how circuits work, voltages and currents other than the power supplies are called 'signals'. In explanations, we usually don't show the power supplies, because, if correctly designed, they just make the circuit work, without affecting *how* it works at all, except by limiting the voltage, current or power that the circuit can produce or absorb.

So, here we have a simple amplifier schematic/circuit diagram (Figure 4). It isn't quite a simple as possible; it includes two resistors and two capacitors. They are there because without them the circuit with 'ideal' components would be too ideal to be helpful.

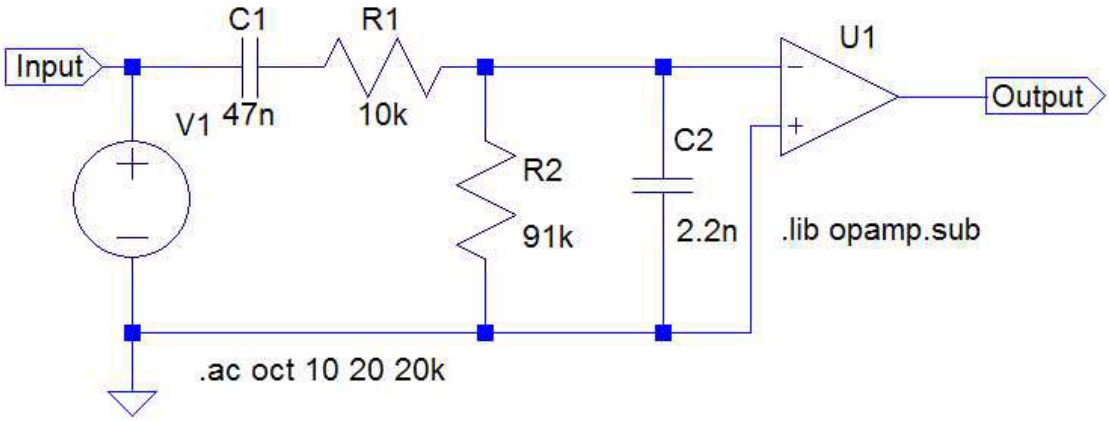

Figure 4 A nearly-basic amplifier

The input signal comes from the generator V1, which produces a sine wave voltage at 1 V peak (not RMS), which, when we ask LTspice to simulate the circuit, is swept in frequency from 20 Hz to 20 kHz, while LTspice calculates the output voltage at each frequency. Disregard the + and - signs: they are only needed if we are looking at phase shift. The amplifier U1 has a gain of 10 times, but the other parts cause the overall gain – the ratio of output to input – to be less. U1 is some sort of integrated circuit or module that we can't get inside; we can only use its connections to the outside world as the manufacturer advises.

C1 is there in a practical circuit to block DC coming from whatever is connected to Input. R1 and C2 are there because we don't want the amplifier to be busy amplifying stuff at frequencies higher than we can hear. R2 is necessary because each input must have a path to the zero-voltage wire at the bottom. C1 and R2 cause the amplifier's output to decrease at low frequencies. R1 and R2 cause the voltage at the minus input of the amplifier to be just over 0.9 V peak at middle frequencies, where the capacitors have almost no effect. The small triangle symbol is usually described as 'ground' or 'earth', but that is misleading and gets complicated. It really tells LTspice 'This is the *zero-voltage reference point* from which you shall measure all voltages'. Unlike real wires, LTspice wires have no resistance or inductance - they are ideal wires. The magic spells 'lib.opamp.sub' and '.ac oct 10 20 20k' are instructions to LTspice to tell us what the circuit does. You may notice that the input signal goes to the minus input, so the output waveform is upside-down compared to the input waveform. For speech and music, this has little effect, except a small subjective change in some male voices. This sort of amplifier is an *inverting amplifier.*

When we run LTspice, it tells us the *frequency response* of the whole schematic from Input to Output, as shown in Figure 5.

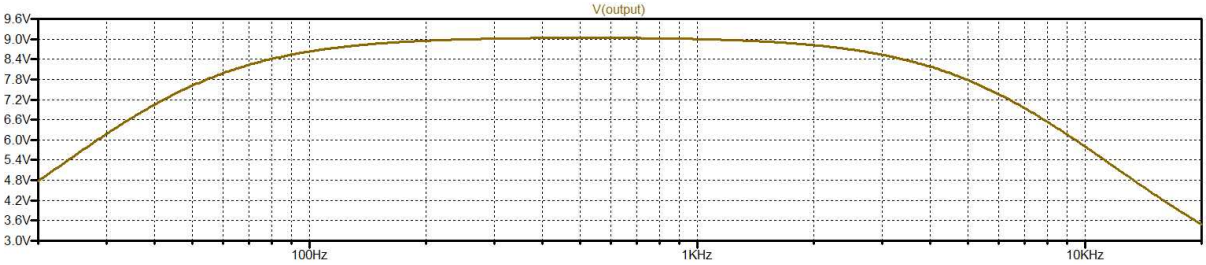

Figure 5 Frequency response in terms of output volts

We see that at middle frequencies, the output reaches 9 V, which is what we would expect with a signal at the minus input of 0.9 V and an amplifier gain of 10.

9 The dreaded decibel

We don't often use frequency response graphics with the response measured in volts. For a large number of reasons, we prefer to use decibels on the y-axis. For example, in a cascade of amplifiers the voltage gains multiply but the decibel values simply add. A voltage of *v* volts is expressed in decibels referred to 1 V as:

$$L_V = 20\log(v/1)$$

The same formula works for currents. There must always be a reference voltage or current, stated or implied, unless we are looking at the difference in decibels between two voltages or currents (the ratio of the voltages or currents themselves), in which case the reference voltage or current cancels out.

The frequency response of our amplifier with a decibel scale for the output (the *output level)* is shown in Figure 6. LTSpice works in decibels referred to 1 V unless we play tricks with it.

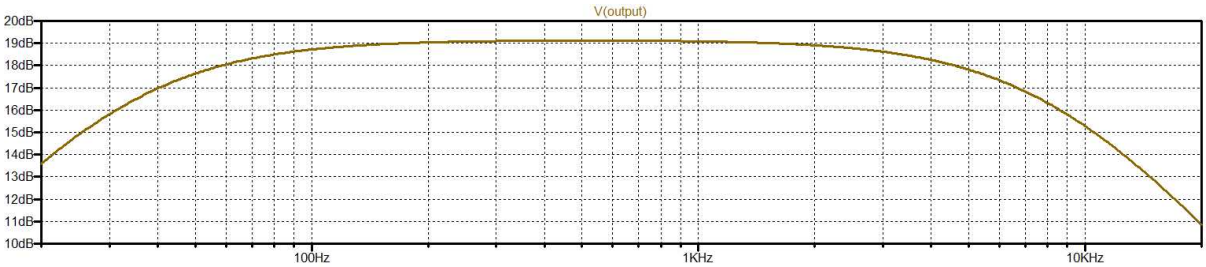

Figure 6 Frequency response in terms of output volts expressed in decibels referred to 1 V

Traditionally, we would concentrate on the frequencies where the output level has decreased by 3 dB, a ratio of 0.71 in voltage and 0.5 in power, so these frequencies are *half-power points.* In our case, the frequencies are 31 Hz and 7.5 kHz. But in studying hearing, the -10 dB frequencies are often more significant, as anything below that isn't sufficiently audible (unless it's unwanted noise, in which case it has to be much weaker). In fact, a reduction of 10 dB is often sensed as 'half as loud'. Also the shapes of the *roll-offs* at each end of the *frequency response curve* can be important. The roll-offs in Figure 6 are called 'first-order' because they are due to one resistor and one capacitor. Higher order roll-offs are due to more components and show as steeper curves.

10 Negative feedback

This is something market researchers hate but electronic engineers couldn't do without. It was invented by the American engineer Black, and at first wasn't taken very seriously by his peers. It consists of simply sending back (*feedback)* a fraction of the output signal of an amplifier so as to oppose the input signal (*oppose = 'negative'*). This reduces the gain but it improves frequency response and reduces distortion, which are good things. Given enough of it, the overall performance depends almost entirely on the components creating the feedback and not on the properties of the forward path through the amplifier. This works especially well with op-amps, because their forward gains (usually represented as the symbol A_{OL}, for 'open-loop') are typically 50 thousand to 1 million. If we reduce the gain of the whole circuit to 10 or even 1000, we are applying a lot of feedback.

Figure 7 shows a simple example of negative feedback, in which the components connected between the output and the minus input provide it. The amplifier U1 is now an op-amp with an open-loop (i.e. without

feedback) gain of 100 thousand.

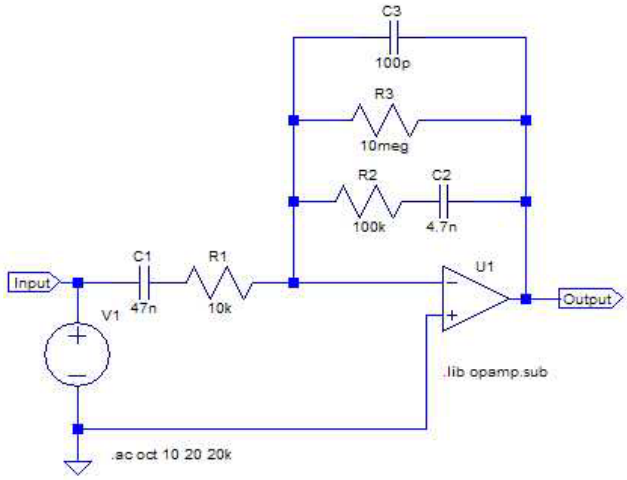

Figure 7 Amplifier with negative feedback

It isn't as simple a circuit as possible, because that would give unrealistic results. The main feedback path is R2 and C2. At middle and high frequencies, R2 and R1 act to set the gain from Input to Output to be a fraction less than 10 times, or 20 dB. It's a fraction less because of the effect of R3, 10 MΩ (megohms: Spice simulators use 'meg' instead of 'M' because they are case-insensitive and 'm' means 'milli'). R3 is necessary because it provides a DC path to the zero-voltage reference via the op-amp internals. C3 prevents the frequency response going up to high frequencies that we do not want to amplify. C2 has a special function – it cancels the roll-off effect of C1 down to very low frequencies, where its impedance becomes comparable with R1's 10 kΩ.

The overall result of this is shown in Figure 8. You can just see the effect of C2's cancellation effect in that extended low-frequency response and its limitation by R3 at the extreme left of the response curve.

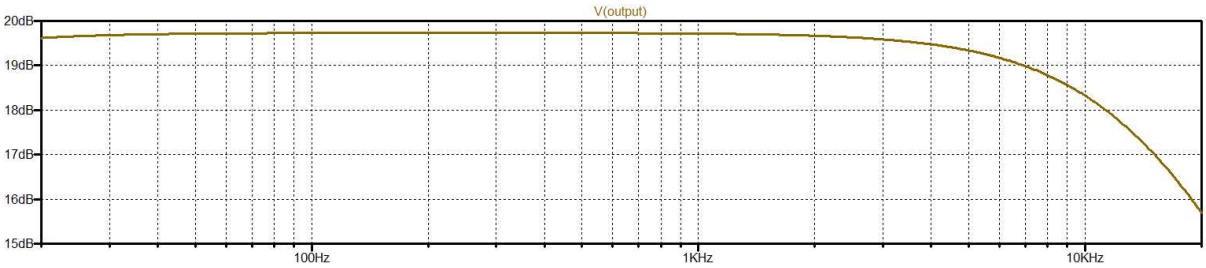

Figure 8 Frequency response of the circuit in Figure 7

The output voltage of an op-amp is equal to the open-loop gain (OLG) multiplied by the difference between the plus input and minus input voltages. Because the gain is so high, for any reasonable output voltage (up to, say, 20 V) the input voltage difference is minute and can be disregarded, In Figure 7, this means that the minus input is at zero-reference voltage, the same as the plus input. The input impedance at the minus input is infinite, so no current flows into it. The current flowing through C1 and R1 flows through the feedback network to the output, which we can also regard as current flowing the other way. Thus we have a *feedback current* that is proportional to the *output voltage,* and this configuration is known as *voltage feedback.*

We now come to a very important step on our way to current-drive amplifiers. We look at the effect of negative feedback on the *output source resistance* of an amplifier. In Figure 9, R4 represents the source resistance, which is actually inside a real op-amp that you can buy; 100 ohms is a typical value. Our schematic now has a new magic spell which allows us to do two simulation runs, one with R5 = 100 Ω (brown curve) and one with it set at 1 MΩ (red curve). You can see that without the feedback, R4 and R5 being both 100 ohms, the output voltage would drop from just under 10 V to just under 5 V. (if you are not sure, the same current flows through both, so applying Ohm's Law...). What happens with the feedback applied is so dramatic that in order to show it in Figure 10 I have abandoned decibels and asked LTspice for the actual voltages, which show the very small difference much better.

LTspice has a feature that allows very precise measurements on graphics. At 200 Hz, for example, the output voltage with the (negligible) 1 MΩ load resistance (R5) is 9.69562 V, while with R5 = 100 ohms it is

9.69459 V. Since the same current flows through the source resistance and the 100 ohms, their resistances are proportional to the voltages across them, so we can write:

$$OSR = 100 \times (9.69562 - 9.69459)/9.69459 = 0.0106\ \Omega$$

where *OSR* is the apparent output source resistance with feedback. It's gone down from 100 Ω without feedback to just over a hundredth of an ohm.

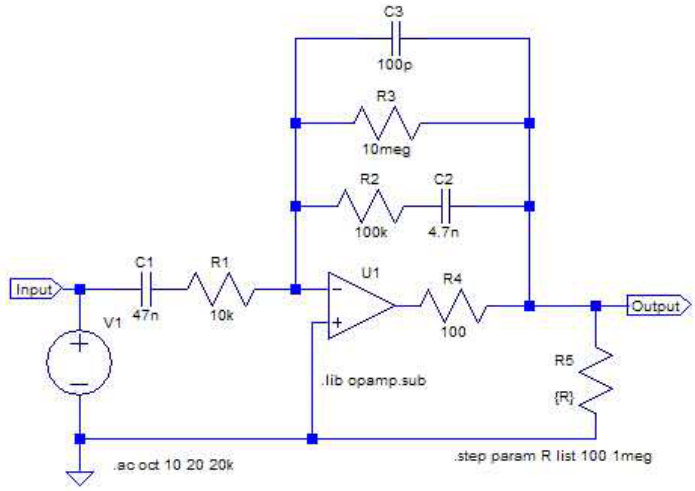

Figure 9 Amplifier with 100 Ω output source resistance

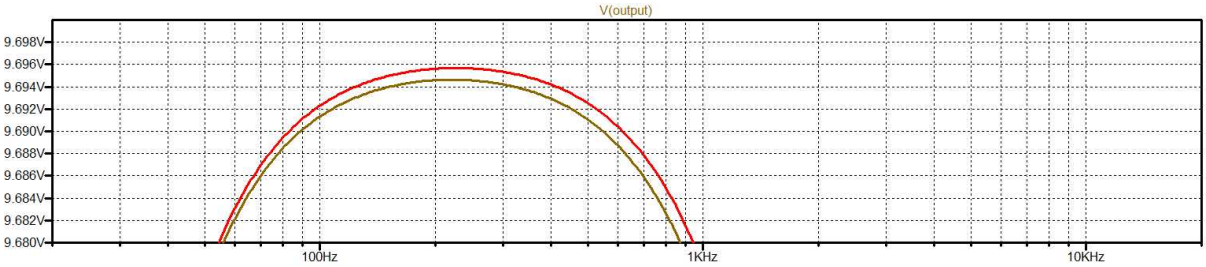

Figure 10 Effect of feedback on the apparent output source resistance

In fact, there is a fundamental principle at work. The op-amp's open-loop gain is 100 thousand, while the circuit has a gain of (nearly) 10. So the output source impedance is reduced in the same ratio, i.e. by a factor of 10 thousand.

11 Current feedback

Figures 7 and 9 have feedback derived from the output voltage. We can also derive feedback from the output current – and that makes (when correctly done) a ***current-drive amplifier!*** Figure 11 shows an example using an op-amp. The signal is applied to the plus input, leaving the minus input free for the feedback connection, which is taken from R5, which is in series with the load resistance R4. The voltage across R5 is proportional to the current through R4 and R5; no current flows into the minus input.

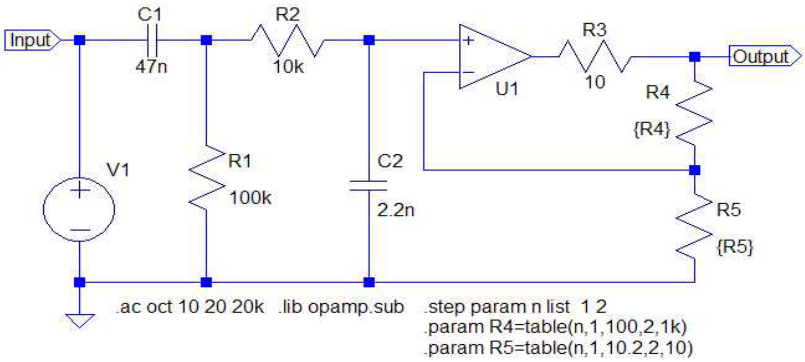

Figure 11 Amplifier with current feedback

Figure 12 shows what happens when the load resistance R4 is 100 Ω or 1kΩ. The output current changes only by a very small amount indeed. In fact, I had to use a special magic spell so that LTspice would

actually show two curves in Figure 12. This behaviour is exactly what we want for driving hearing loops, because their impedances vary with frequency.

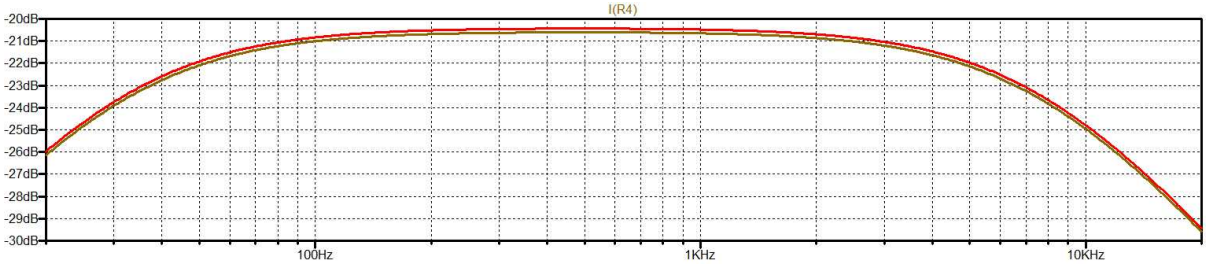

Figure 12 Frequency response of the Figure 11 schematic with two values of R4

12 Impedances of hearing loops

A hearing loop is just a loop of wire. The wire has electrical resistance, and because the AC current in it generates a magnetic field, it has inductance, which can be defined as the rate of change of magnetic flux with current. If it were laid out as a 'hairpin', with the 'go' and 'return' legs very close together, the inductance would be very small, because the magnetic fields from the two legs almost cancel out. It wouldn't be a very good hearing loop! Opened out into a useful loop encircling an area, its inductance increases. A guide value is 1.6 μH (microhenrys) per metre of wire. Suppose we have a loop of wire with an area of 1 mm^2 and the wire is 48 m long (it is a square, 12 m or 40 feet on each side). Its inductance is 48 x 1.6 = 77 μH. At 100 Hz, this has a reactance of 2π x 100 x 77 μ = 0.048 Ω, which is negligible compared with the resistance, which is 0.74 Ω. But at 5 kHz, the reactance is 50 times as much, i.e. 2.4 Ω. If we drove the loop with a constant voltage, the current would be considerably reduced at 5 kHz, causing poor intelligibility. But if we use a current-drive amplifier, that doesn't happen. Figure 13 shows our amplifier now with a loop as load and Figure 14 shows the resulting frequency response.

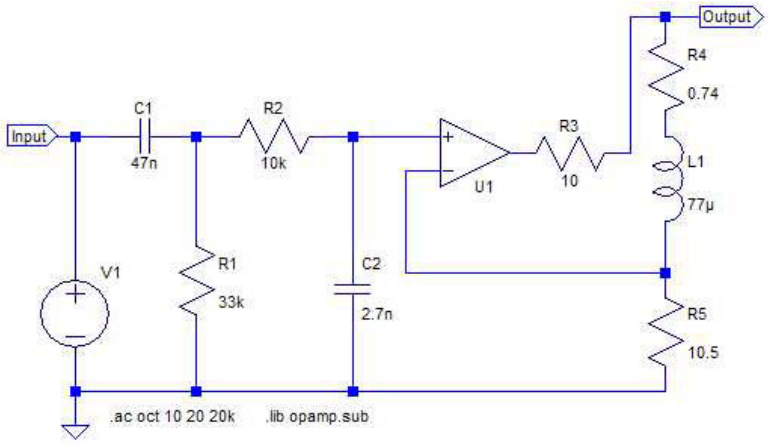

Figure 13 Current-feedback amplifier with loop load (R4 and L1)

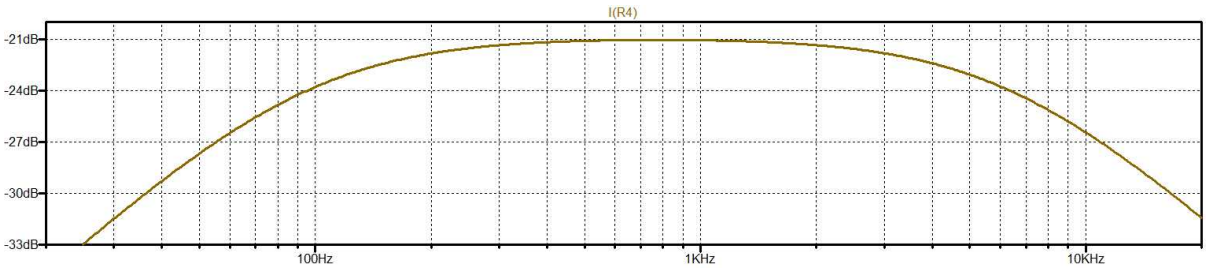

Figure 14 Frequency response of the loop current in the Figure 13 schematic

I have tweaked the values of the 'realism' components, C1, C2, R1 and R2, so that the frequency response of the loop current is -3 dB at 100 Hz and 5 kHz, as specified in IEC 60118-4. But the loop current is -21 dB referred to 1 Amp, i.e. 90 mA, which isn't going to produce 400 mA/m anywhere useful in the loop. Also, the value of R5 is much higher than is used in real amplifiers, because in this example the loop current is quite low.

For this 12 m square loop, to get 400mA/m field strength at the centre of the loop, we need a loop current given by:

$$I = \pi A \times 400/(2\sqrt{2} \times 1000)$$

A is the length of the side of the square, 12 m in this case, so I = 5.3 A. In practice, we need more, 5.7 A because our listeners will not have their ears at floor level and the reduction of field strength at the average 'listening height' of 1.45 m is not negligible. But we will set that aside for the time being. What is clear is that to get a loop current of over 5 A we will have to use an audio power IC instead of a simple op-amp. But there is a way we can reduce the current requirement, at the expense of needing more output volts. The specifications of many audio ICs (chips) allow this, and there will be more about it in some examples later.

13 Current requirements of loops

Loop system designers make frequent use of a graph or a table of current values. The current required depends on the width of the loop, its length (greater than, or equal to, the width) and the distance between the hearing aids and the plane of the loop, the listening height. For the most-used graph or table, we take the listening height to be 1.45 m, up from the floor or down from the ceiling. Also, we call length/width the 'aspect ratio'.

Here are the graph and table as calculated with Excel. (It is very difficult to get the complicated expression for the current into Excel without errors.) 'Width' is the shorter side of the loop (unless it's a square, with q = 1) and 'q' is the aspect ratio.

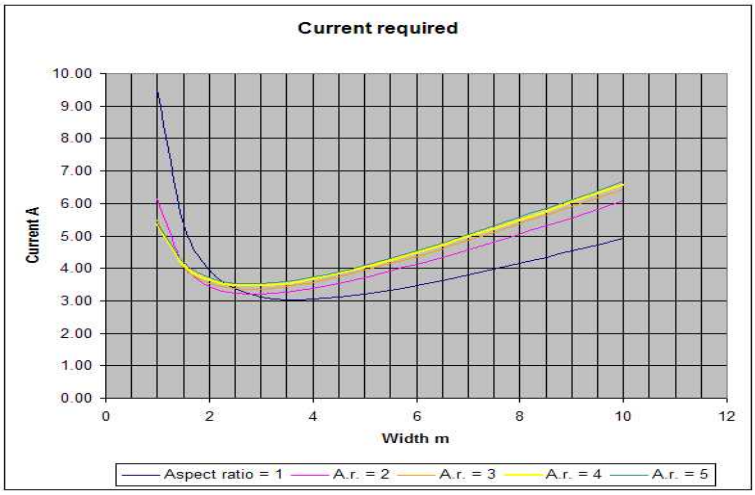

Figure 15 Current requirements for loops

q	1	2	3	4	5
Width m					
1	9.54	6.16	5.49	5.38	5.41
1.5	5.35	4.09	4.02	4.10	4.18
2	3.95	3.44	3.53	3.64	3.71
2.5	3.37	3.23	3.37	3.48	3.54
3	3.12	3.19	3.36	3.46	3.52
3.5	3.04	3.25	3.44	3.54	3.60
4	3.05	3.37	3.57	3.67	3.73
4.5	3.11	3.53	3.74	3.84	3.90
5	3.21	3.71	3.94	4.04	4.09
5.5	3.33	3.91	4.15	4.26	4.31
6	3.48	4.12	4.37	4.49	4.54
6.5	3.63	4.34	4.61	4.73	4.79
7	3.80	4.57	4.86	4.98	5.04
7.5	3.97	4.81	5.11	5.24	5.30
8	4.15	5.06	5.37	5.50	5.56
8.5	4.34	5.30	5.63	5.77	5.83
9	4.53	5.56	5.90	6.04	6.11
9.5	4.72	5.81	6.17	6.32	6.39
10	4.92	6.07	6.44	6.59	6.67

Figure 16 Loop current table from Excel. q is the aspect ratio

It is easy to see that no loop requires less than about 3 A. The only way we can use amplifier chips that won't deliver 3 A is to use loops of more than 1 turn. Provided the turns are all bunched up close together (e.g. the cores of a multi-core cable), the approximate current required is divided by the number of turns (often a little more is found to be necessary; an extra turn might solve that) and the inductance is

approximately multiplied by the square of the number of turns, this increasing the voltage required to drive the necessary current. This is often not a problem because the amplifier will produce enough output voltage. However, the installation of a loop with many turns can be difficult unless a multi-core cable is used; it is better to use an amplifier chip that will deliver more current. To go further with this, we need to look at actual examples of current-drive amplifiers using popular amplifier chips.

14 Keeping the amplifier cool

When the amplifier IC is working, it gets warm, because it takes power from the power supply but delivers only some of it to the loop. In fact, it can only deliver power permanently to the resistive part of the loop; anything it offers the inductive part has to be re-absorbed during the next half-cycle. This effect mostly *increases* the heating in the IC and we can see why in Figure 17. We can't get inside the IC to obtain these curves, so they are from an amplifier using separate transistors.

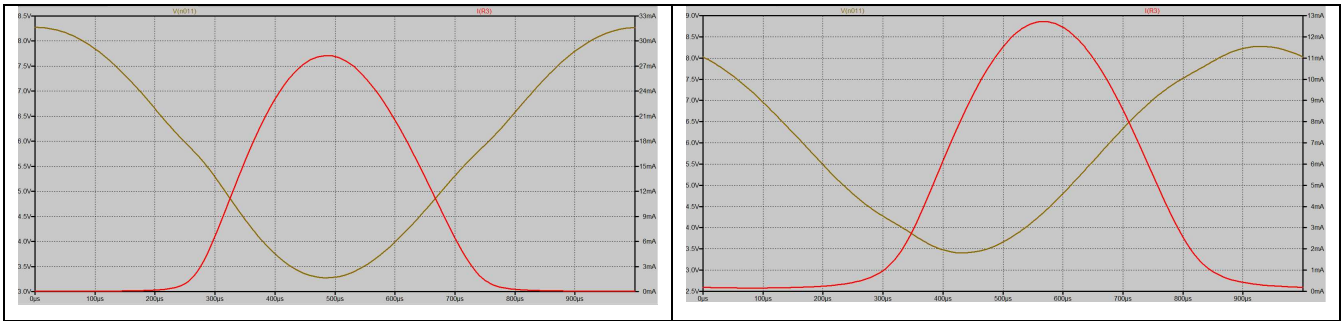

a) Voltage across (brown) and current through (red) an output transistor, with resistive load	b) Voltage across (brown)and current through (red) an output transistor, with inductive load (loop)

Figure 17 Voltage and current waveforms inside the amplifier

With the resistive load, the current is low while the voltage is high, and vice versa, so the heating power (voltage x current) is low. With the inductive load, the current is shifted in phase relative to the voltage, so that high current occurs at the same time as higher voltages, so the heating power is increased. We have to get rid of this heat by fixing the amplifier IC to a heat sink.

It's possible to calculate this extra heating effect, which varies according to how hard the amplifier is driven. The worst case is testing the amplifier with sine wave test signals, which can produce many times as much heat as is produced when the amplifier is working with speech signals. Figure 18 gives a rather idealized guide to what happens. The 'Peak signal' curve applies to sine-wave testing at full output. The phase angle can't get to exactly 90°, because the loop always has some resistance, but at 5 kHz it can get to more than 80°.

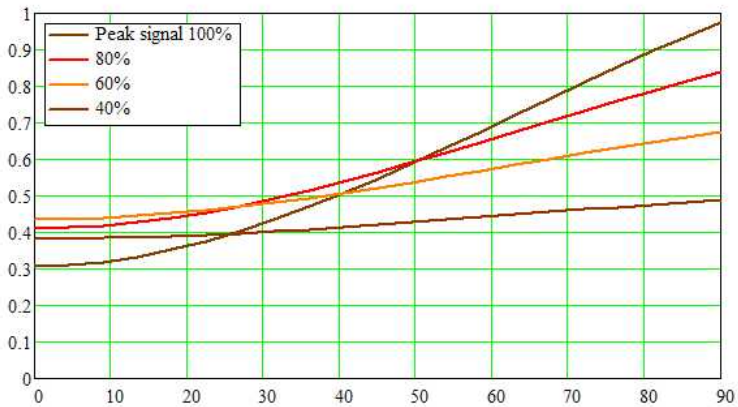

Figure 18 Heating effect in the amplifier varies with signal strength and loop impedance phase angle

The x-axis is the load phase angle, $\tan^{-1}(2\pi fL/R)$ and the y-axis is the heating power in terms of the supply voltage x the load current.

In practice, it's often more time-consuming to calculate a heat sink design than to measure temperatures on an actual amplifier. This is not the place to go into heat sink design, but there are some principles:

- thick metal is better than thin;
- soft metal is better than hard (aluminium is much better than steel and soft and quarter-hard aluminium are much better than harder grades);
- vertical is better than horizontal ('hot air rises', after all);
- black is better than shiny (radiation loss *is* significant, and black anodized aluminium is widely used for heat sinks).

15 Ideal and real amplifiers

An ideal amplifier has infinite input resistance and zero output resistance. Real amplifiers have finite resistances in both cases. The finite input resistance is not a significant problem. But when we make a current-drive amplifier circuit, we want the output source resistance R_s (which is produced by the circuit configuration; it is not a property of the amplifier chip, whose finite output resistance does not significantly affect the operation of the circuit) to be much higher that the loop impedance, so that the loop current does not depend significantly on the loop impedance, which varies with frequency:

$$Z_{loop} = \sqrt{\{(2\pi f L_{loop})^2 + R^2_{loop}\}}$$

f is the frequency in Hertz and the other symbols should be self-explanatory. Z_{loop} is highest at the highest frequency we are interested in, 5 kHz. Some mathematics shows us that if R_s is at least 5 times Z_{loop} at 5 kHz, the loop current is reduced by less than 0.2 dB, which is negligible. If it is 2 times instead of 5 times, the loss at 5 kHz is nearly 1 dB, which may make it difficult to meet the requirements of IEC/EN 60118-4.

Referring back to Figure 11, some more mathematics shows that the value of R_s is very nearly equal to the value of R5 multiplied by the amplifier gain, which is the ratio of the output voltage to the difference between the + input voltage and the – input voltage. (For an op-amp, this is very small indeed but for audio amplifier chips it can be anything from 10 to 400.) It has been alleged that the R_s of some loop amplifiers was not high enough in the past, leading to poor high-frequency response.

16 Circuits with real amplifiers

16.1 LM386

This device is small, inexpensive and works well, BUT its data sheet seems to be rather conservative in its output current specification (calculated from the output power values). The highest calculated value is 300 mA, and I tried to get a more definite value from the manufacturer, without success. It will produce a lot more output current without getting too hot, using test signals similar to speech and music, but these higher currents may cause longer-term damage, so we should not go there.

Figure 15 shows that no area-coverage loop requires less than 3 A loop current, so to use the LM386 we would need a loop of at least 10 turns, which is pretty impracticable. It would be so much simpler to use an amplifier with a higher output current. So what about a smaller loop, such as a neck loop? It turns out that there is still a problem. Remember, we need current drive for area-coverage loops because the inductive reactance of a loop at 5 kHz is bigger (a lot bigger) than the resistance. But *this does not always apply to practicable small loops.* We often find that, except for a single turn of very thick wire requiring a large current (bad for battery life even if the LM386 or a similar device can supply it), as we add more turns the inductive reactance never gets much larger than the resistance, in spite of it being roughly proportional to the square of the number of turns, because we have to reduce the thickness of the wire also in proportion to the number of turns, otherwise the loop would become very thick and stiff.

The solution is to use a different way of getting a constant loop current, which has the considerable advantage that the resulting loop works with any reasonably standard headphone output, supplied by a conventional voltage-drive amplifier. But the loop, which needs around 15 turns (a miniature 15-core cable with the cores connected in series), *is still current-driven*, because it has a resistor of 14 Ω in series, so it also looks like a headphone to the source of signals. That 14 Ω completely swamps the impedance of the loop itself, keeping the current through it constant. My neck loop has an inductance when spread into a 230 mm diameter circle of 144 µH, which is 4.5 ohms at 5 kHz, and it has 3.9 Ω resistance, making the impedance 6 Ω, so the '5 times rule' isn't met, but it works pretty well; the response at 5 kHz is down 0.27 dB, still negligible.

Note that this is by no means 'energy efficient': the majority of the battery power just warms up the 16 Ω resistor, but we simply don't have amplifier chips that can produce about 1 A output current with a supply voltage of about 30 mV to suit a single-turn loop. The Universe is not that accommodating to neck loops

and looped hats. In any case, the loop requires only 60 mA for 400 mA/m, so the maximum power would be 0.06^2 x 19.9 = 71.6 mW and the long-term (60 seconds) average power is 4.5 mW. Not much global warming potential there!

16.2 LM380

This device looks a bit more promising, because it can deliver 0.9 A. However, for reasons hidden in its internal electronics, it will only work correctly in current-feedback mode at much less than its full potential. The problem is the tendency, mentioned in the data sheet, for a burst of high-frequency oscillation to occur during part of the signal cycle. Unless you looked at the output with a wide-band oscilloscope, you would not see this, but it is not acceptable. The LM380 can accept a supply voltage as high as 22 V, but there is no point in using a higher voltage than necessary as it just heats up the device. It is thus much better suited to the 'series resistor' alternative approach to current drive explained above. Note that an 8 Ω series resistor needs to have a power rating of 8 x 0.9^2 = 6.48 W or more if the full output current of 0.9 A is needed, but with speech signals, the average power is only one-sixteenth or so of that.

There isn't anything special about the value of 8 Ω for the series resistor; one could use a higher value, which requires a higher supply voltage, so there doesn't seem to be much point.

16.3 TDA2003

This device is particularly suited for current-driving smallish area-coverage loops. It is designed to drive low impedance loads, down to 1.6 Ω, while being intended to work from a single supply voltage of up to 18 V, which is enough for what we want to do. Its peak current limit is 3.5 A but the non-repetitive limit, which is what applies to speech signals, is 4.5 A, giving an RMS value of 3.2 A. This is not enough for anything but a 3.5 m square single-turn loop, but if we use two turns, effectively giving us 6.4 A to play with, Figure 15 shows that we could drive quite large loops, more than 10 m square, 8 m by 16 m or 7 m by 35 m. In fact, it would be unusual for such large areas to be satisfactorily served by single peripheral loops; they would most likely require arrays of smaller loops, to reduce overspill and/or combat metal loss.

The device (there is also a TDA2003A version with the same specification) is listed as 'obsolete' but it seems to be still freely available.

To take a specific example, the characteristics of a typical 8 m square loop are:

Wire area: 1 mm^2

Current required for 400 mA/m at 1.45 m above centre: 2.1 A per turn

For 2 turns: Resistance: 1 Ω Inductance 225 μH approximately: it depends on how close the two conductors are. For best results use twin figure-8 cable or twisted pair, with the two cores connected in series.

Peak voltage required at 1.6 kHz: 7.3 V

Why 1.6 kHz? Because speech doesn't include all frequencies at the same strength. The curve of strength against frequency is called the spectrum, and it slopes off above 1.6 kHz, so while we need to have a constant loop current, at least at all frequencies from 200 Hz to 2.5 kHz (remember the current can and should be less at 100 Hz and 5 kHz), the system will never be called upon to supply maximum current at frequencies above 1.6 kHz. This is very helpful; by keeping the loop voltage down (it would have to be 23 V at 5 kHz), the heating of the TDA2003 and the requirements for the power supply are both reduced. To get 7.3 V peak, we need a supply voltage of rather more than double, i.e. 17 V.

With these figures we can build an amplifier and also use Spice simulation to investigate far more than we can do in a reasonable time by measurement. This device is not like an op-amp; the + and - inputs are quite different electrically. The + input (pin 1) has a resistance of about 100 kΩ to pin 3 (ground) while the – input (pin 2) 'sees' the emitter of a transistor and so has a quite low resistance to ground.

Figure 19 shows the circuit diagram of a TDA2003 amplifier in the form of an LTspice schematic, similar to that in Figure 13, flanked by graphs of loop current (top) and voltage (below). The loop is represented by L1 and R1. The TDA2003 Spice model was kindly made for me by Jim Thompson. The brown curves are for a conventional voltage-drive configuration and the red curves show, particularly, the flat frequency response of the loop current. LTspice requires a number of 'magic spells' on the diagram to tell it what to do. For the current-drive amplifier, R8 is very small (compared with R7) and R6 is 200 mΩ, giving an effective current-source resistance of 20 Ω. For the voltage-drive amplifier, R6 is made very small (compared with R1 and R4) and R8 is made 100 kΩ to compensate for the gain increase due to the removal of the current feedback.

It is easy to see that, as expected, the 20 Ω current source keeps the loop current constant up to at least 10 kHz, whereas without it, it begins to fall off above 400 Hz. But what happens at low frequencies? With voltage-drive, the current falls off below 300 Hz, whereas with current drive it is constant down to 10 Hz. (To keep things simple, I didn't put any extra 'realism' parts in this circuit; it really is flat down to 10 Hz.) The reason that current-drive improves the response is due to C4. With voltage drive, C4 sees a resistive load of close to 1 Ω (the loop resistance), and its impedance begins to rise noticeably below about 300 Hz, reaching 1 Ω at 160 Hz and increasing further. But in current-drive mode, the capacitor 'sees' 21 Ω, so the response does not begin to fall off even slightly until 15 Hz, which is almost off the end of the graph.

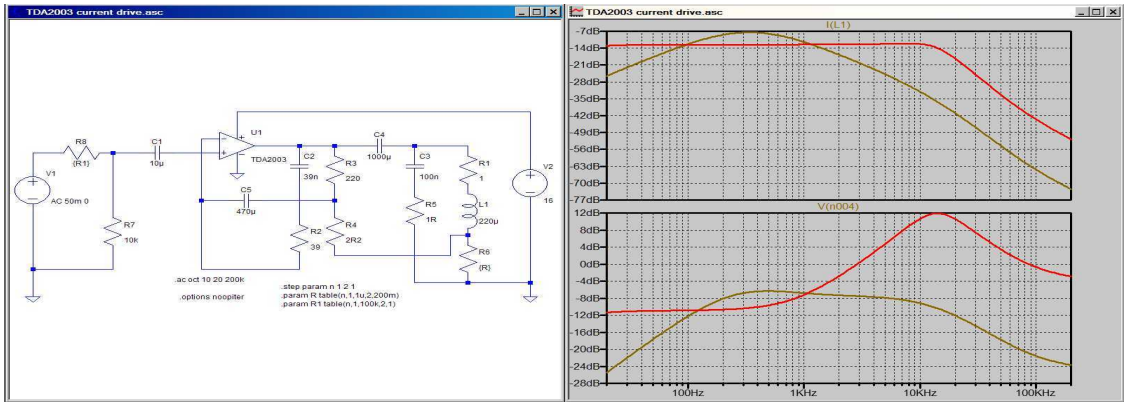

Figure 19 TDA2003 voltage- and current-drive amplifiers, showing the effects on the frequency response of the loop current and loop voltage

16.4 TDA2030

This device will work in current-drive mode but its maximum output current is internally limited to 3.5 A peak, which means 2.5 A RMS, and that is less than the 3 A minimum current required for *any* single-turn area-coverage loop. Also, the manufacturer now lists it as 'not recommended for new designs'. It will not do anything the TDA2003 can't.

There was once a TDA2030A version, rated for ±22 V supplies instead of ±18 V, which would probably be of little advantage for a loop amplifier, but the data sheet shows how to use such a device (plain or A version) to drive a pair of high-power discrete transistors. This is going well beyond 'Introduction...' so I'll say no more about it.

16.5 LM1875

This is a similar device to the TDA2030, but with a higher peak current limit of 4 A. It doesn't seem to offer any advantage over the TDA2003.

17 Measurements on a TDA2003 amplifier

These measurements relate to an actual amplifier using the circuit shown in Figure 19. The results with voltage drive are shown in Figure 20 a) and with current drive in Figure 20 b).

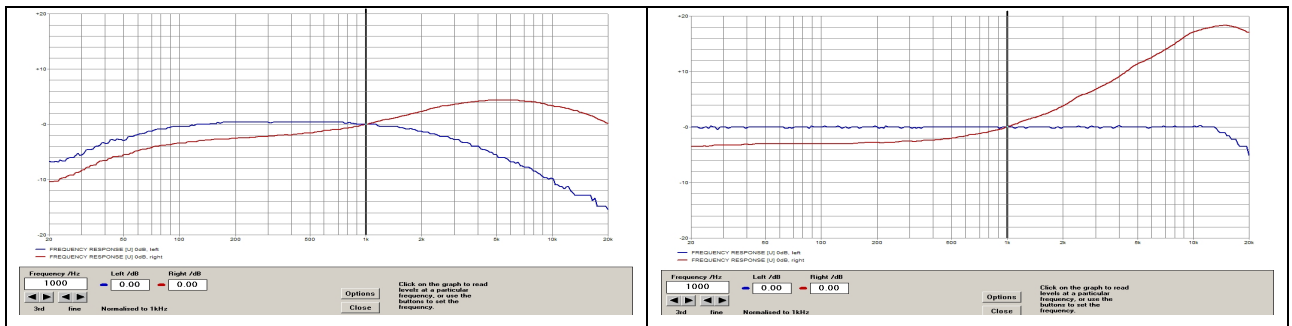

a) voltage drive; the blue line is loop current and the red line is loop voltage

b) current drive; the blue line is loop current and the red line is loop voltage

Figure 20 Measurements on a TDA2003 amplifier

The output voltage rises at high frequencies even with voltage drive because the output source

impedance of the amplifier plus the internal resistance of the capacitor C4 is not negligible compared with the 1 Ω resistance of the loop. In fact, since the response rises about 6 dB, we can say that the resistance in question is also about 1 Ω.

18 A complete loop system amplifier

The paper concentrates on the 'current-drive' part of the complete loop system amplifier, but a complete amplifier includes two more main building blocks and four more features. Figure 21 shows the essential parts.

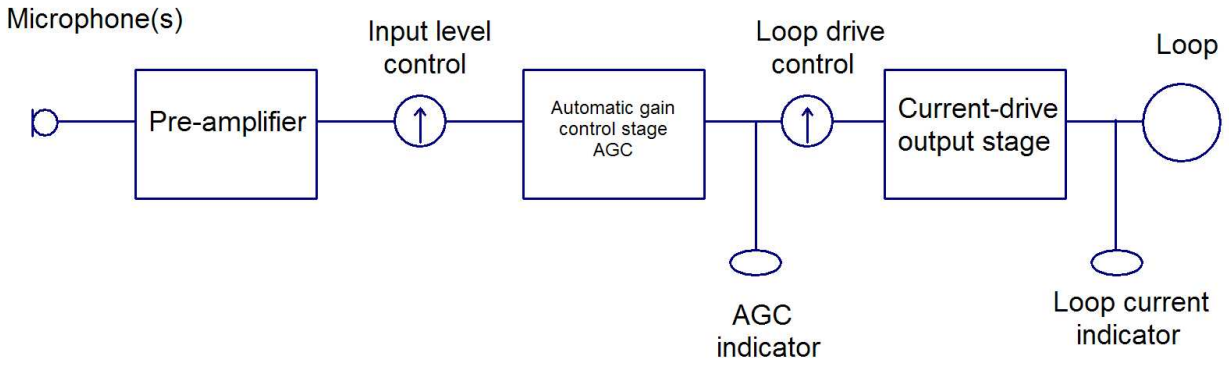

Figure 21 Block diagram of a complete loop system amplifier

The preamplifier accept inputs from microphones and perhaps other sources (DVD, TV) and adjusts their signal levels to be roughly equal. The input level control then sets the level into the AGC stage so that the AGC is in operation, as shown by the AGC indicator. The AGC needs a strong enough input signal to work. The output from the AGC stage is a (very nearly) constant level from all inputs. The loop drive control then sets the input level to the current-drive output stage to give the correct loop current, as shown by the loop current indicator. The installer should set the sensitivity of this indicator so as to get the correct loop current. If it's not set, the current may be too high and the telecoil sound therefore too loud. The loop current indicator also shows a fault (no indication) if the loop wire has been broken.

19 Acknowledgements

Thanks are particularly due to Jim Thompson of Analog Innovations, http://www.analog-innovations.com/ for developing the TDA2003 Spice model. All errors are mine.

20 Comments and questions

Please submit by email, quoting the text concerned and the line number(s). Please make your question as clear as possible.

21 Redistribution and extractions