Op Amps

Starting an in-depth look at the designer's favourite chip, the operational amplifier.

PAUL CHAPPELL

In circuits built from discrete components, a large proportion of the design effort is absorbed in trying to minimize the imperfections and unpredictability of the components used. In building a low frequency gain stage with a bipolar transistor, for example, beginners quickly learn that a current bias circuit (Fig. 1a) is no good because the wide variation in current gain between different samples of transistors of the same type means the base bias resistor must be selected by trial and error for each individual device.

The potentiometer bias circuit (Fig. 1b) reduces the circuit's sensitivity to transistor gain; with suitable choice of resistor values the circuit can accommodate any transistor of a given type without modification. The price paid for this convenience is partly that the signal headroom is reduced because of the voltage across the emitter resistor, but mainly that the circuit as a whole has very much less gain than the transistor itself.

The technique of trading off gain (or giving it up altogether) in return for better performance in other respects is a very useful one. Figures 1c and 1d show two non-amplifier applications of the idea. The first is a simple (and not very satisfactory) Miller integrator, where the transistor is used to linearize the charging of a capacitor. The second uses a transistor to simulate a very large capacitor (roughly h(fe) x C), at least as far as charging is concerned.

The drawback of transistors is that individually they don't have a lot of gain to give up. Useful circuit building blocks are made not from one but several devices.

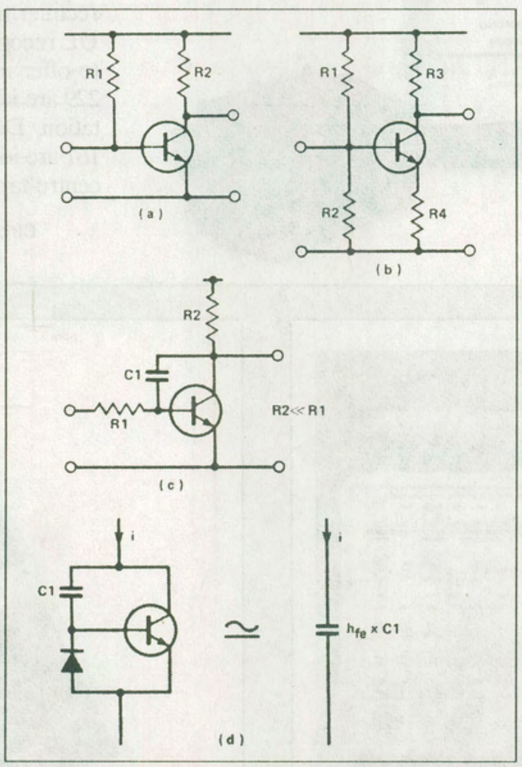

Fig. 1(a) Transistor current bias circuit. (b) Improved potentiometer bias circuit. (c) Simple integrator circuit. (d) Simulating a large capacitor.

They also suffer from being direct embodiments of a basic physical process and make about as much concession to practically as a piece of school laboratory equipment. The 0.7V base-emitter voltage exists not because designers want it that way, but because that's how transistors are. It's how the physics works.

A good case can be made for the view that op amps have more in common with discrete components than with other ICs. Used raw, without any associated passive components, they are totally unmanageable, but just connect a few external components and a very wide range of useful circuit building blocks can be made.

The most striking advantage of op amps over transistors is that the available gain is several orders of magnitude higher, allowing very precise control over the circuit's characteristics with a single device. What's more, the external connections are arranged for the convenience of the designer, not constrained by the demands of physics.

However, like discrete components, op amps have their own imperfections and idiosyncrasies. Technology has moved the decimal points a few places, but greater performance seems to lead inevitably to greater expectations. There is always somebody who can think of an application which would be possible if only the latest device were just that little bit better.

Coaxing the best performance from op amps is still the same mixture of art and science as for discrete components. In this series of articles I intend to cover the basic techniques and also to touch on some of the finer points of designing with these versatile devices.

Op amp ICs

With most families of ICs the pin configuration for each member has to be learned individually but with op amps it's easy. They come packaged in ones, twos or fours and the pin connections are almost always the same for any brand (Fig. 2a). High performance devices are usually packaged individually in the 8-pin DIL or T099 package and the spare pins may be used for offset cancellation or external compensation but the basic configuration of inputs, outputs and supply connections is usually adhered to.

Op amps are very tolerant of supply voltages. Most will operate with single supplies from below 10V to above 35V. Some (particularly those intended for battery operation) will run from as little as 4V. Voltages above 40V are rare but you can have that if you want it (and can afford it.).

Op amps circuits are usually run from split rail supplies (Fig. 2b, 2c). This is because the limits of the input and output voltages for correct operation fall short of the supply voltages, so the central 0V rail is a useful bias and reference point. (Note that although data sheets often give absolute maximum input voltages as being equal to the supply voltages, this rating shows the most the IC will suffer without damage, not the range in which it will operate properly).

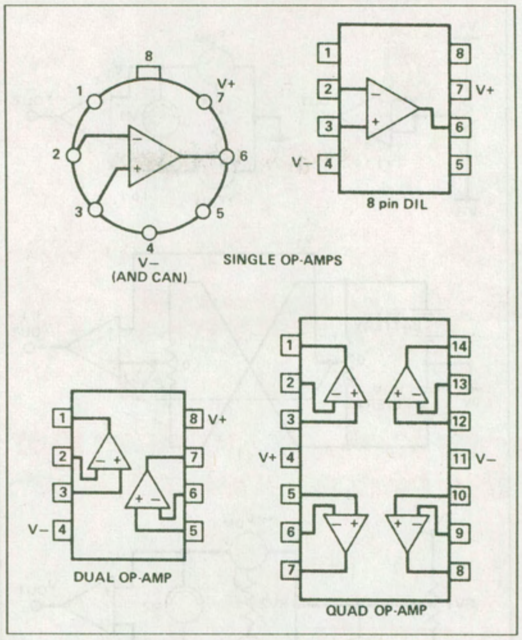

Fig. 2(a) Op amp pin connections.

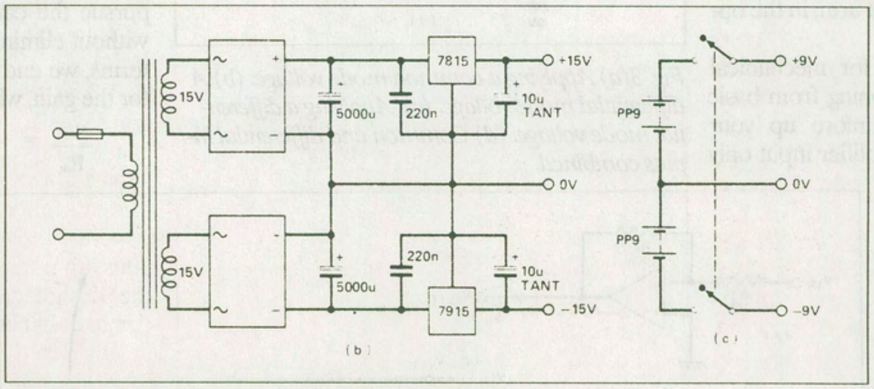

Fig. 2(b), (c) Split rail power supplies for op amp circuits.

Op amp Basics

Let's indulge in a flight of fancy for a moment. We've just received a sample of the very latest op amp. It has extremely low drift and offset, superb common mode rejection, very low noise, bias currents of 1nA — in short it's the kind of IC any manufacturer would love to produce. Unless I say otherwise, it's this $200 Rolls Royce of op amps that we'll be using in this article. Let's see how it behaves.

The IC has a voltage gain of 10^6. This means the output will be 10^6 x *difference* in the input voltages, taking into account which is the higher in voltage. If the +input is 1uV higher than the -input, the output will be at +1V. If the -input is 1uV above the +input, the output will be at -1V. If the two inputs are at the same voltage, the output will sit firmly at 0V.

Connect together the two inputs and vary their common voltage by means of a pot (Fig. 3a). The output will sit firmly at 0V regardless of the setting of the pot (it won't follow the input voltage) because there is no difference between the two input voltages. This shows that the IC has excellent (perfect, in fact.) *common mode rejection* — it ignores voltages common to both inputs.

Figure 3b shows the IC's response to a *differential mode* signal. You have to imagine here that if the pot is set at the centre of its rotation, both voltage sources are zero. If it is rotated clockwise, v1 gives a light positive voltage and v2 gives an equal negative one. If the pot is rotated anticlockwise, v1 will be negative and v2 the same amount positive. If you are happier looking at a more concrete circuit, Fig. 3c shows one that will do the trick.

The centre-zero microvolt meter M1 registers the differential mode voltage. The output of the op amp will be one million times the voltage shown on the meter.

Figure 3d combines a common mode voltage (set by RV1 and shown on M1) and a differential mode voltage (set by RV2 and registered on M2). When the two inputs are not at the same voltage, their common mode voltage is defined as the voltage exactly halfway in between, so vcm = 1/2(v(+) + v(-). The differential mode voltage is centered neatly on the common mode voltage.

Varying RV1 will have no effect on the output voltage, regardless of the setting of RV2. Varying RV2 will give an output exactly one million times the reading on M2, regardless of the setting of RV1. Just to make sure you've got the hang of it, if M1 shows - 6.22V, M2 shows +3.2uV, what is the common mode input voltage? The differential mode input voltage? The voltages at the + and - inputs? Most important of all, what is the output voltage? Answers at the bottom of the page.

Amplifiers

About the first thing anybody learns about op amps is that they can be used to make amplifiers with a gain precisely controlled by the values of a pair of resistors. Figure 4a shows one of the ways this can be done.

Having read somewhere that the gain is given by -R2/R1, my fantasy is that the circuit of Fig. 4a has a gain of -10. I'll try to prove it to you.

The first thing to notice is that if vin is 0V, the output will also be at 0V. If it tried to go just a teensy bit positive, the potential divider action of R1 and R2 would put a positive voltage on the - input, which would tend to push the output back towards 0V. If it tried to go negative, the resulting negative voltage on the - input would force it back up again.

In a way it's like one of those children's toys with weighted bases that always settle in the vertical position after being knocked. Any attempt to push to the op amp's output voltage by brute force results in a restoring force which will push the output back to OV as soon as you let go.

Suppose that vin is increased to 1V. If I'm right about the gain being -10, the output will now want to settle at -10V. If so the voltage at the -input will be 0V, exactly the same at the +input so the output must also be at 0V. What's gone wrong? It can't be at -10V and 0V at the same time.

The fault lies in my faith in rule-of-thumb calculations. The formula -R2/R1 for the gain is a very good approximation for most practical purposes, but it's not exact.

Suppose that Vout settled at just a little above -10V. This would allow just enough positive voltage on the – input to be amplified up a million times and maintain the output at this voltage. Once again, the weighted base action comes into play. Any attempt to shift the output from this voltage results in a restoring force (using the term loosely.) which tends to force it back again. If you find the idea of "just a little above -10V" too vague, don't worry. The calculations will be along in just a moment.

One way of looking at the circuit is to see it as a kind of voltage lever. The arms of the lever will be proportional to the resistor values and the pivot will be at the -input terminal of the amplifier.

Pushing down on the input (lowering the voltage) makes the output rise ten times as far. The pivot is just a little bit loose – it moves just one millionth of the distance of the output arm, in the opposite direction.

If you don't care for mechanical analogies, perhaps reasoning from basic electronic principles is more up your street. Assuming the amplifier input only takes 1nA of current and since we are dealing with tens and hundreds of uA flowing in R1 and R2, it's reasonable to say that for practical purposes all the current in R1 must also flow in R2. Now, if there is the same current flowing through two resistors then by Ohm's law the voltage across each will be proportional to its resistance.

In other words, if 1.2V is dropped across R1, and the very same current is flowing in R2, you can say without further ado (and without bothering to calculate the current) that the voltage dropped across R2 will be ten times as great: 12V.

Now, whatever voltage appears at the output of the amplifier, the voltage at the - input will only be one millionth as much. There's very little point in taking it into account at all. We might as well say that it stays at 0V. So the input voltage will be the voltage across R1, the voltage across R2 will be the output voltage and we've already worked out that this will be ten times as great (or -10), taking into account that it moves in the opposite direction). In other words, the circuit has a gain of -10.

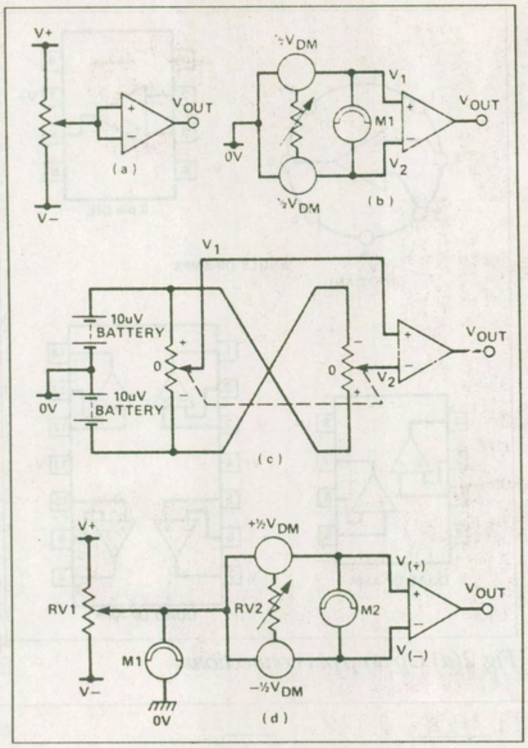

Fig. 3(a) Applying a common mode voltage. (b) A differential mode voltage. (c) Applying a differential mode voltage. (d) Common and differential inputs combined.

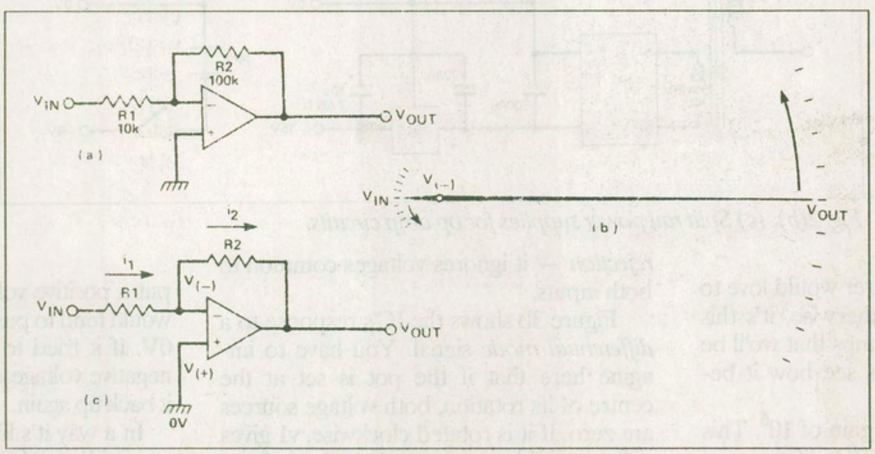

Fig. 4 (a) A basic inverting amplifier. (b) The lever analogy. (c) Inverting amplifier.

If all this business about ignoring little errors makes you feel uncomfortable, the only way to settle the matter is to do the calculations. Looking at Fig. 4c, by Ohm's law we can write:

$$i = \frac{V_{in} - V(-)}{R1} \quad \text{and} \quad i_2 = \frac{V(-) - V_{out}}{R2}$$

Now, if the op amp's input takes negligible current (I'll have to fudge this bit for the time being or things will get impossibly complicated. I'll come back to it later) then i(1) = i(2), so:

$$\frac{V_{in} - V(-)}{R1} = \frac{V(-) - V_{out}}{R2}$$

We also know that the gain of the op amp is 10(6), so v(-) = - 10(-6) Vout, giving:

$$\frac{V_{in}}{R1} + \frac{10^{-6}V_{out}}{R1} = -\frac{10^{-6}V_{out}}{R2} - \frac{V_{out}}{R2}$$

It's usual at this stage to point out that the terms involving 10(-6)Vout are very much smaller than either of the other two terms (can you spot a condition where one or other wouldn't be?) and so can be neglected, giving:

$$\frac{V_{in}}{R1} \simeq -\frac{V_{out}}{R2} \quad \text{or} \quad \frac{V_{out}}{V_{in}} \simeq \frac{R2}{R1}$$

which leads to the usual rule-of-thumb formula for the gain of - R2/R1. If we pursue the calculation to the bitter end without eliminating the two inconvenient terms, we end up with the exact formula for the gain, which is:

$$\frac{V_{out}}{R_{in}} = -\frac{R2}{R1 + 10^{-6}(R1+R2)}$$

Using this formula, the circuit of Fig. 4a, which I said would have a gain of -10, actually has a gain of -9.99989. So the rule of thumb in this case is not too far from the truth. In fact, in comparison with the 5% resistor tolerances likely to be used in a practical circuit, it's as close to perfect as you need.

My fudge factor, assuming the inputs take no current, doesn't affect the validity of the formula, although a proof of this and an investigation of just what effect it will have must wait for another time. There's more to these op amps than meets the eye.

Answers to Problems

The answers to the problems I posed earlier on, by the way, are: the common mode input voltage is -6.22V, the differential mode voltage is +2.3uV (I hope you got those.), the voltages of the + and - inputs are – 6.2199984 and -6.2200016V respectively, and the output will be +3.2V. ■