something useful when energised with AC, not DC. This is why the world adopted AC mains electricity in the first place –because it can be transformed.

The coil that we drive is called the **primary** and the other coil/s which we draw upon are called **secondaries**. Ideally, we would like the tightest possible coupling between the coils so very little energy is wasted. This is usually accomplished by winding them very close together on a single ferromagnetic core. Nearly all the flux will flow preferentially in the high permeability core material rather than in the low permeability air outside, so nearly all the flux will be ‘guided’ through each and every turn of wire.

If the sinusoidal mains wall voltage is applied across the primary of the transformer, an alternating current will flow which in turn produces a corresponding H-field in the core. The H-field, when multiplied by the permeability of the core material, produces a certain flux density or B-field in the core. The alternating flux flowing in the core in turn produces a back-EMF which opposes the applied wall voltage. Without this opposing EMF the primary current would be limited only by the wire resistance. In other words, the primary has *inductance* which limits the current to a steady value known as the **magnetizing current**. This is equal to the AC voltage divided by the inductive reactance:

$$I_{mag} = \frac{V_{pri}}{2\pi fL}$$

Magnetising current does nothing useful as far as transforming is concerned; it is an unwanted inefficiency. The transformer designer will minimise it by winding hundreds or thousands of primary turns to achieve a fairly high inductance, limiting this current to perhaps a few tens of milliamps in the sort of transformers we’re interested in.

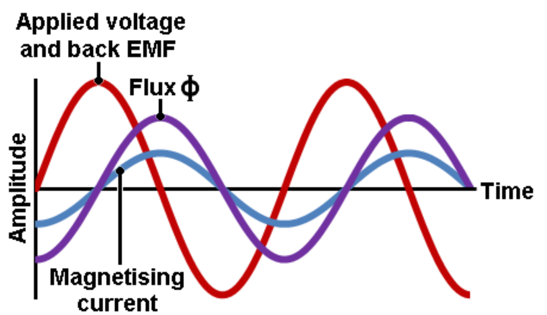

Fig. 1.6: Phase relationship between applied voltage, magnetising current, and magnetic flux, for the primary winding of an ideal transformer.

Fig. 1.6 shows the phase relationship between voltage, flux, and magnetising current in a transformer primary. The flux is in phase with the current since it is conjured by it. But the back-EMF (which exactly opposes the applied voltage and therefore lies directly on top of the same trace) is proportional to the *rate of change* of the flux, in accordance with faraday’s law. The flux is changing fastest at its zero crossings, so it is here that the EMF reaches its peaks. This is the physical explanation for why the voltage across a simple inductor always leads the current through it by 90°.

to the cathode internally (in which case they are sometimes referred to as 'directly heated' as a shorthand). This forces us to use a separate heater supply dedicated to the rectifier, which will float on top of the HT, thereby eliminating any possibility of straining the heater-cathode insulation. Many valve power transformers include a 5V winding for just this purpose.

Directly heated rectifiers are normally powered from a heater winding with a centre tap, as in fig. 2.33a. Taking the DC output from this centre-tap forces the DC ripple current to distribute more evenly through the filament, avoid a hotspot as well as cancelling residual heater hum on the HT. For indirectly-heated rectifiers where the cathode is internally connected to the heater, the DC output should be connected to the actual cathode pin as in fig. 2.33b.[4] Connecting it as in c. is wrong since it forces ripple current to flow through the heater and transformer winding.

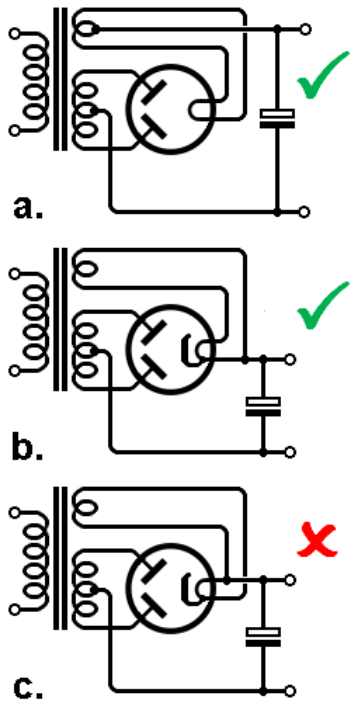

Fig. 2.33: Proper connection to valve rectifier filament / cathode.

2.6.4: Hot Switching and Inrush Current

Valve rectifiers should always be allowed to charge the reservoir naturally from cold, as the heater warms up (see section 5.3.4). 'Hot switching' refers to pre-heating the heater/cathode before switching on the anode current. This is the principle of most guitar-amplifier standby switches, for example. Depending on how such a switch is implemented this can mean the rectifier is fully warmed up, but the reservoir empty, when the switch is finally thrown. The valve will then suffer the massive inrush current into the reservoir capacitance. This is likely to cause momentary saturation, which can be very damaging to the cathode, and valve rectifier failure (arcing) in guitar amps is uncomfortably common for this reason. Hot switching of rectifier valves was expressly discouraged by valve manufacturers, as the RCA 5U4GB datasheet notes: "*Even occasional hot-switching with capacitor-input circuits permits the flow of plate current having magnitudes which can aversely affect tube life and reliability.*"

A related problem may also present when *opening* a standby switch, by inducing a ghastly flyback voltage across the transformer secondary, large enough to cause

[4] Delaney, W. J. (1949). Power Pack Problems, *Practical Wireless*, September, pp.353-4.

bobbing up and down on top of the AC waveform to create a new sinusoidal waveform at the C_1-D_1 junction, which is 'jacked up' by the voltage stored across C_1. The positive-going half cycles of this waveform allow C_2 to be charged through D_2, producing the final DC voltage ④, which is twice the original peak voltage. All the voltage multipliers in this chapter work on this 'ratchet' or 'bucket brigade' principle.

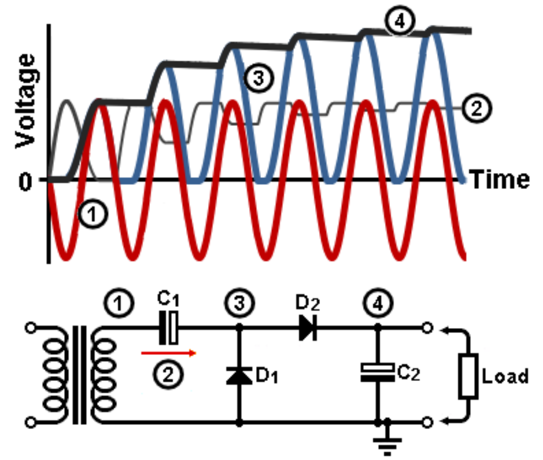

Fig. 3.2: Waveforms in the half-wave voltage doubler.

From the point of view of the transformer the circuit is full-wave since both halves of the AC waveform are used equally, which avoids transformer saturation. But C_2 is only charged up once per mains cycle, and since the load is connected across this capacitor, the circuit is only half-wave from the point of view of the load. In other words, the ripple voltage is the same as for a simple half-wave rectifier:

$$V_{ripple} = \frac{I_{dc}}{fC_2} \tag{3.1}$$

(The ripple voltage increases disproportionately if the circuit is extended to higher multiplication factors as described in section 3.3.)

When considering ripple current, we must look at the circuit from the point of view of the transformer, where it appears to be full-wave. In fact, the ripple current for this voltage doubler is identical to an ordinary bridge rectifier supplying the same DC load power and reservoir capacitor C_2. Notice that we stipulate the same load *power*, i.e. twice the load current or one *quarter* the load resistance, since our fictitious bridge rectifier would not double the voltage.

Thus, one way to design a voltage-doubler of this kind is first to imagine we need half the voltage and twice the load current that we are aiming for, but the *same* ripple voltage. We then design an ordinary bridge-rectifier that would satisfy these figures, using any and all of the methods already familiar to us from the previous chapter. Finally, we convert it into a voltage doubler configuration, leaving the transformer unchanged. The voltage doubler will then meet our actual voltage and current requirements. The main differences are that the ripple frequency will be 50/60Hz rather than 100/120Hz, and it will take longer to charge up after switch on.

increased LF distortion or crosstalk in a stereo amplifier, while more severe positive feedback can cause sustained LF oscillation called **motorboating**, or sub-1Hz oscillation called **breathing**.

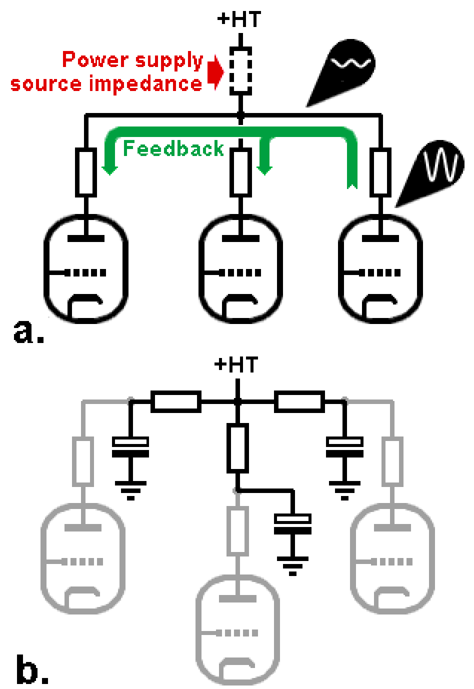

Fig. 4.6: **a:** Unwanted feedback due to a common impedance in the power supply. **b:** Decoupling filters isolate each stage from the rest while also providing smoothing and bypassing.

This problem is reduced by feeding each stage from its own RC or LC filter so that each stage is decoupled from the rest, and the lower the cut-off frequency the better. Any signals trying to creep along the power supply rail from other stages will then be attenuated by the local decoupling, thus reducing the loop gain around the amplifier through the power supply. Yes, this is similar in principle to smoothing, except we're now talking about audio-induced ripple rather than rectifier ripple. One massive capacitor in the attic could be thoroughly effective at smoothing rectifier ripple, yet be useless for decoupling since the long wire leading from it to the amplifier would represent a shared (common) impedance across which a program-induced voltage would appear.

Ideally, we would provide one filter for every amplifier stage, but the usual compromise of one capacitor for every two stages usually suffices for decoupling. If stereo channels share the same power supply then similar valves from opposite channels can share a filter, as in fig. 4.7, since positive feedback is not an issue with this arrangement (the problem is instead crosstalk, but as it should only exist at sub-audio frequencies it is less important). However, wide-bandwidth or digital circuits should also have ceramic or plastic bypass capacitors mounted very close to each amplifying stage or IC, in addition to general electrolytic bypassing, as even a little inductance can form a common impedance that leads to HF (rather than LF) instability. This is much less likely in an audio circuit, especially the valve kind –see the next section.

Whether to use a system of cascaded filters or individual decoupling filters all connected to a common star point (see fig. 4.9), is at the designer's discretion. For similar component values a cascaded arrangement provides better ripple smoothing rules either way. Note that the mains safety earth conductor must *never* be switched under any circumstances.

The switch must be rated for use on mains voltage, and its current rating should be generously larger than the operating current. Most readily-available mains switches are rated for several amps –far more than most valve amplifiers consume– so this should not be a problem. As a practical matter, an unofficial convention is to mount the switch so that, if something were to fall on it from above, it would tend to knock it into the 'off' or 'safe' position.

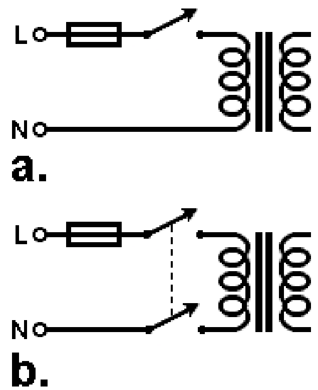

Fig. 5.7: **a:** The primary switch must be placed in series with the live wire. **b:** Double switching can be used for total mains isolation.

5.2.2: Secondary Switching

Switching of low-voltage circuits is simple and requires hardly any thought beyond the appropriate voltage and current ratings. Arcing can be suppressed to some extent by placing a small (100nF say) capacitor in parallel with the switch, as in fig. 5.8. When the switch opens, the capacitor presents a brief low-impedance path until it is fully charged, by which time the switch has (hopefully) already opened.

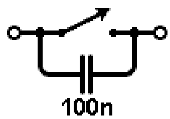

Fig. 5.8: A capacitor may be connected across a switch to suppress arcing, particularly when switching DC.

5.3: Standby Switching

Occasionally we may want to provide a separate switch for the HT supply. This commonly takes the form of a so-called standby switch, which allows the heaters to warm up before the HT is applied. Really it would be better to call it a 'pre-heating switch', as 'standby' has since become an accepted term for the power-saving feature found on many modern appliances. Nevertheless, the term 'standby switch' is firmly entrenched in guitar amp lore, somewhat unfortunately, as explained in the next section.

There are some valves that genuinely *d*o need a warm-up period before the HT is applied, but this will be explicitly stated on their data sheets. This includes mercury-arc rectifiers (the mercury must be vaporised before use), valves used in DC-coupled applications such as voltage regulators, and some special purpose valves with unusual getters that must be pre-heated. Such valves are not normally used in audio amplifiers, much less guitar amplifiers.

The trouble with implementing a standby (pre-heating) switch is that the HT voltage will often exceed the voltage ratings of readily available switches by a fair margin,

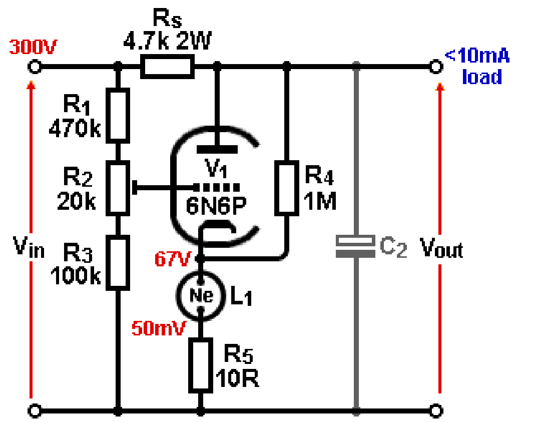

Fig. 6.21: Practical King stabiliser. Not recommended for general use.

The neon lamp provides a relatively fixed cathode voltage which, after everything else, conveniently allowed a 4.7kΩ dropping resistor to be used with 5mA shunt current as before. This is not the true optimum value (which will also depend on load current since that also affects the anode voltage) but it is close enough to prove the point. R_4 ensures the neon always strikes, as explained in 6.3.1 earlier. The purpose of R_5 is to allow the quiescent current to be monitored and set to 5mA, i.e., adjust the trimmer R_2 until the voltage across R_5 is 50mV. This can be a frustrating exercise; the circuit is *very* sensitive and will drift for at least five minutes as everything warms up.

Since R_s is not optimised, ripple reduction is not worth mentioning –output capacitor C_1 can do that job. The main feature of this circuit is its ability to stabilise DC. This turned out to be better than the author expected, considering the compromises made. Once fully warmed up, it can cope with an input voltage variation from 280V (drop out) to 320V (10mA maximum neon current) while maintaining the DC output voltage within ±1.5%. Note that the heater voltage was supplied from the same transformer and so followed the same line variation. Not bad for one modest triode. Really though, the circuit is interesting as an academic exercise. It is too sensitive to be useful outside the laboratory, so we shall leave it in 1923 and move on to another 1920s development: negative feedback.

6.6: Feedback Shunt Valve Regulators

Feedforward shunt stabilisers provide line regulation but not load regulation. With feed*back* we get load regulation *and* some line regulation. Negative feedback brings with it all the well-known advantages of lower output impedance and greater independence from component variation and ageing. This means it can eliminate the need for trimming. Feed*forward* circuits can deliver better line regulation with the same parts, but the attendant difficulties in making this happen are usually enough to make anyone choose the feedback route instead, trading some line regulation in exchange for load regulation and less swearing. It is also possible to use a little of both, as we shall see later with mixed-mode regulators.

voltage, though more stages could be added. Using a HEF40106B with a 15V supply, the idle supply current was 1.4mA and the unloaded output voltage was found to be 47V. The output voltage sagged down to 41V with a 10kΩ load, and the supply current increased to 13mA –an overall efficiency of 84%.

7.6: HT Voltage Reducers

Sometimes we do not need a regulated voltage supply, but we do want to be able to adjust the voltage over a range. This might be for reasons of experimentation, or for power reduction in a guitar amp, as covered in chapter 10.

7.6.1: Voltage Follower

A voltage follower offers a simple way to control a DC voltage, illustrated in fig. 7.23. This concept is of course related to the capacitor multiplier and Zener follower (section 7.7.1). The pot provides an adjustable reference voltage to the base of the follower/pass-device which acts as a simple buffer. It dumbly copies the pot voltage and provides a low output impedance, isolating the pot from the load. The output voltage is simply an attenuated version of the input voltage, including any ripple, which is attenuated by the same proportion. When using a transistor the output voltage will be slightly lower than the pot wiper voltage, whereas in the case of a valve it will be slightly higher, owing to bias voltage. A MOSFET is the logical choice for a high voltage environment. Note that if the pot is turned fully up the base will be shorted to the collector and the pass device will behave as a simple diode; this is a useful feature for power reduction in a guitar amp since it gives us access to normal, unadulterated operation when the pot is at maximum. The pass device may be called upon to dissipate a lot of power depending on how much the voltage has been reduced. Assuming the load resistance is constant, the worst-case dissipation occurs when the output voltage is set to half the maximum, at which point dissipation in the pass device will be equal to dissipation in the load.

The pass device does not necessarily need to work in the pure DC environment as in fig. 7.24a, but can be placed before the reservoir capacitor, effectively creating a

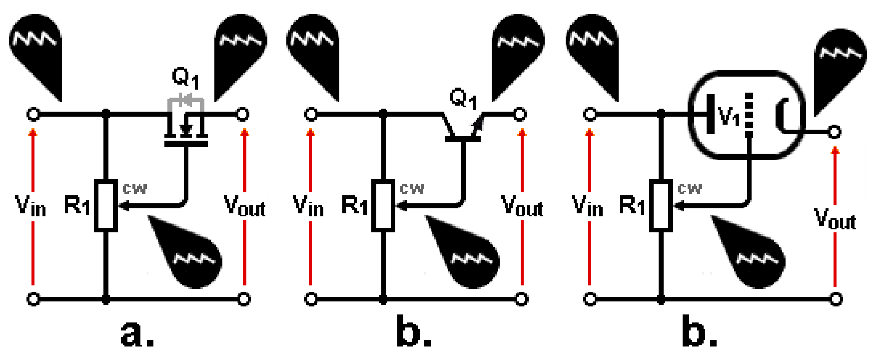

Fig. 7.23: Unregulated voltage adjustment using a follower.

transformer to the heater winding too. The result is that the heater voltage will not be pure 50/60Hz but will contain all the high-frequency hash that rectifiers produce, plus any noise coming from the wall supply itself, of course.

Fig. 9.9 shows an oscillogram of the heater-voltage waveform from a transformer which also provides the HT in a small valve amp, and it is obviously not a very good sine wave. The top of the wave is clipped and shows the classic ‘hangover’ explained earlier in section 2.4.3. The high dV/dt of this hangover, plus any rectifier switching noise that may exist too, will easily couple into the audio circuit via valve interelectrode capacitances. The lower trace shows the signal picked up on the grid of the input valve (which had a 1MΩ grid leak) and transients coinciding with each hangover are clearly visible. It is these transients which often cause heater interference to sound buzzy, rather than the low hum we might otherwise expect. Heater balancing can suppress this differential-mode noise.

9.3.5: Electrical Heater Balancing

The heater supply must always have a reference to audio ground, which may be a direct connection or an elevating circuit (next section). This is equally true for AC or DC supplies. Leaving the heater supply floating will result in almighty hum due to primary-to-secondary transformer leakage current, and is a common beginner’s error. AC heater supplies should also be balanced to suppress the EM field. Not only does this reduce the magnitude of the voltage on each wire (e.g., each wire handles ±3.15V rather than one wire handling 6.3V and the other zero), but the opposing fields will tend to couple equal-but-opposite hum signals into the audio circuitry, which should cancel each other out.

If the transformer heater winding has a centre tap it can be grounded to create a balanced heater supply as in fig. 9.10a, and this is by far the most common approach found in vintage equipment. However, if a short develops between the anode and heater pin of a valve socket (on popular octal valves like the EL34 these pins are unfortunately adjacent) then it will short the anode directly to ground through the heater supply, which is the worst possible fault condition. It is therefore preferable to insert a small resistor (ideally fusible / flameproof) into the heater centre tap to serve as an emergency fuse, as in fig. 9.10b. This is also a good habit for beginners to adopt, as they often become fixated with grounding the centre-tap simply ‘because it is there’, even when some other ground reference is already being used. A burnt-out resistor is a much cheaper way to discover a grounding conflict, rather than a burnt-out transformer winding.

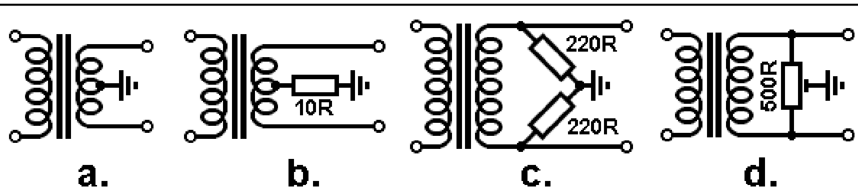

Fig. 9.10: Examples of electrical heater balancing.

section). With this approach the available drive voltage to the power valves reduces in proportion to everything else, so the master volume control can be moved to the input of the driver stage. Many conventional amplifiers already include a gain control in this position. On the other hand, as the power control is turned down the tone is likely to change more with this topology than with the previous one, but overall the results can still be satisfying. Again, the first dropping resistor to the preamp smoothing filters will need to be made larger than usual.

The most extreme option is to vary the supply voltage to the whole ampilfier, as in fig. 10.14. This gives results closest to using a variac, except the heater voltage is unaffected. With this method any existing gain/volume controls should suffice. At very low-voltage settings the preamp distortion can suffer, approaching a dull fuzz, so this topology is probably the least appealing. On the other hand, it is the least invasive when retrofitted into an existing amp.

10.3.5: Design Considerations

The most obvious way to vary the HT voltage is by way of a voltage regulator or follower, such as those already discussed in chapter 7. Fig. 10.15 shows the simplest approach, reproduced from section 7.6.1. The voltage is varied by a pot which feeds a MOSFET source-follower that does the hard work of passing the load current and dissipating excess power.

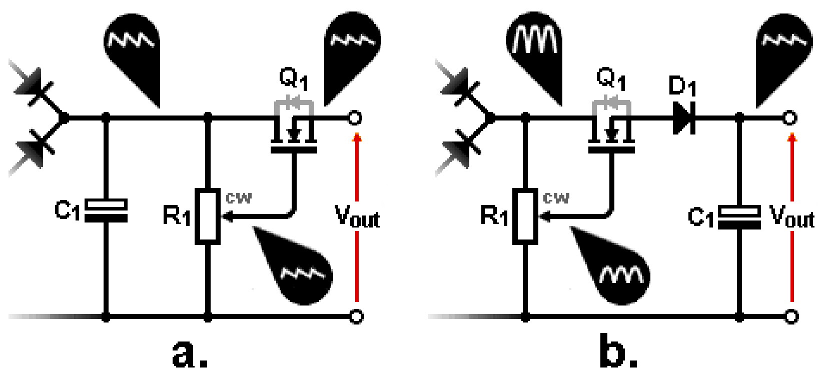

Fig. 10.15: **a:** Adjustable DC source follower; **b:** Source follower as a controlled rectifier reduces dissipation in R_1.

Assuming the load presented by the amplifier circuit is resistive –which is more-or-less true since we will be varying the screen and bias voltages proportionately– then since $P = V^2/R$ the load power will vary with the square of the HT voltage reduction, and the audio output power and SPL will follow suit. In other words, halving the voltage should reduce audio power to one quarter of its initial value, which is a loudness reduction of one third. In practice the power tends to drop at a slightly faster rate, as we saw from fig. 10.10, which works in our favour. Ideally though, for the smoothest ‘feel’ of loudness-versus-rotation we would like the power to drop by