Chapter 1: Fundamentals of Amplification

This chapter deals with the design and analysis of the basic triode gain stage, which is the main building block of a valve preamp. Some of this material is general and applies as much to hifi circuits as to guitar amps, but some of it is very particular to guitar preamps, where distortion is not merely tolerated but actively sought and manipulated. This distortion may take many forms and is determined by many co-dependent circuit elements. It very rare that a single component is solely responsible a particular aspect of the tone, so achieving a particular sound will normally require an understanding of the preamp (and the rest of the system for that matter) as a whole. This 'global' appreciation of the system marks the difference between the competent designer and the amateur circuit bender.

The most common triode type found in guitar amps is the ECC83 / 12AX7, while the ECC81 / 12AT7, ECC82 / 12AU7 and 12AY7 occasionally make an appearance. The consistent use of the same type of valves is partly historical, since many modern amplifiers are derivations (or merely copies!) of a few classic Fender amps. There are hundreds of other valves types worth experimenting with, but commercial designs are bound to use the same, readily available valve types, if only to satisfy consumer expectations. However, the ECC83 does have some special properties which make it ideal for use in overdriven designs –which will become apparent throughout this book– so it is likely to remain the triode of choice for guitar amps indefinitely, and much of this book will focus on its use.

For readers who are not familiar with the many designation numbers it is worth mentioning that the ECC83, ECC803, CV4004, M8137, 12AX7, 7025 and 6681 are all the same valve. The different numbers indicate either different manufacturers or special quality versions, but they all have the same electrical characteristics as far as guitar amps are concerned, and all can be used in the same circuit. Additional letters (such as 12AX7**A**) need not concern us; they were once used to indicate a controlled heater warm-up time or some other feature particular to that iteration of the valve, but with current-production preamp valves any extra letters are usually just a gimmick.

Some versions are reputed to have a particular tonal character, and much has been written about the apparent superiority of, say, the Mullard 'long plate' ECC83, or the RCA 'black plate' 12AX7. These subjective differences are not a consideration for the circuit designer and will not be mentioned here again. So-called 'tube rolling' and 'cork sniffing' is fun, but is left to the discretion of the reader. Real tonal control comes from the choice of topology, manipulation of overdrive characteristics, voicing, and from a complete understanding of the circuit's functionality, not from the particular manufacturer or vintage of the components used. This book with furnish the reader with this essential understanding, and the author will do his best to make it a painless experience.

1.1: Basic Theory of Valves

A thermionic valve or vacuum tube contains two or more electrodes suspended inside a glass bulb. Any air inside the bulb is thoroughly removed during manufacture, leaving only a vacuum. The first electrode is called the **cathode*** and normally consists of nickel† tube coated on the outside with a chemical paste that is formulated to have a very low **work function**, that is, it gives off electrons very easily when heated. Inside the cathode tube is the heater, which is a fine tungsten wire. Operation could hardly be simpler; pass current through the heater so that it gets hot and the cathode will in turn also get hot. However, since the heater is entirely unconnected from the cathode and plays no special part in the audio circuit, it is nearly always omitted from circuit diagrams and is not considered to be a ‘working’ electrode.

When the cathode reaches its normal working temperature of about 1050 kelvin or 777°C, electrons boil off it and drift around near its surface, forming an electronic cloud called the **space charge**. Without some other influence the space charge would build up to the point where it is so dense that it repels any further electrons from boiling off the cathode, which by itself is not very useful.

1.2: Valve Diodes

To make a useful valve another electrode must be introduced, called the **anode**‡. Since we now have two electrodes the resulting device is called a diode, and is illustrated in fig. 1.1. If the anode is made positive (electrode voltages are always measured relative to the cathode) it will attract electrons from the space charge. These electrons accelerate towards the anode and crash into it. In fact, if the anode is at +100V then the electrons will be travelling at about 5927km per second when they impact!

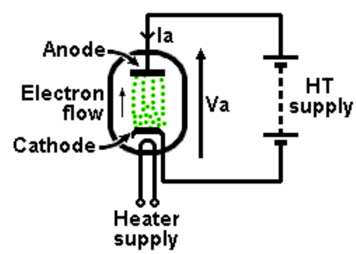

Fig 1.1: Current and voltage in a simple diode circuit.

New electrons continue to boil off the cathode to keep the space charge topped up, so altogether we have a net transferral of *negative* charge from cathode to anode. If we reckon current as flowing from cathode to anode then it must also be negative, but since most people prefer to work with positive numbers it is easier to say that a positive current flows from anode to cathode, as this amounts to exactly the same

* From the Greek *kata hodos* meaning ‘down way’.
† Nickel is used mainly because is has high resistivity which makes is easy to weld.
‡ Originally spelled ‘anhode’, from the Greek *ana hodos* meaning ‘up way’.

thing.* If the anode voltage is negative, electrons will not be attracted to it and no current will flow; hence the valve is a one-way device or diode.

The current flowing from anode to cathode is called the anode current, **Ia**, and the voltage measured between anode and cathode is called the anode voltage, **Va** (or more accurately **Vak**). The actual supply voltage is called the **HT** in this book. This is a historical term that stands for *high tension*, which can be taken to mean *high voltage*. In America the notation **B+** is often used to mean the same thing. In modern transistor circuits the positive supply voltage is variously referred to as V_{CC}, V_{DD} or simply V+.

If the HT is gradually increased then current in the diode will also increase, so for every value of anode voltage (Va) there is a corresponding value of anode current (Ia). If these points are plotted on a graph we will obtain something like that in fig. 1.2, which is called the **static anode characteristic**, or sometimes just the **I/V characteristic**. The characteristic is curved near the bottom when the anode voltage is low because here the anode has only a weak influence over the electrons,

---

* This sometimes causes consternation among students who argue that conventional current 'goes in the wrong direction', but this arises due to a misunderstanding of what current actually *is.* The confusion is understandable because we often describe current as being a flow of electrons, and immediately form a convenient mental picture of what is happening inside conductors. But while this explanation is good enough for everyday conversation, it is not the whole truth.

Strictly speaking, current does not have *any* direction in the usual sense because it is not a physical 'stuff' that moves around a circuit. Current is more properly defined as the *rate of change of electric flux*: $i = dQ/dt$, and is, therefore, an entirely mathematical concept. Current is a scalar quantity; it has no direction, only sign (positive or negative).

The misunderstanding is further compounded by the limitations of the English language which was never designed to cope with the peculiarities of quantum physics. The word *current* actually means the same thing as *flow*, and does not imply a particular direction, yet it is standard practice to describe circuits using phrases like *"current flows from A to B"*, when technically this is nonsense. It is a bit like saying *"the marathon runs from Charlton Way to The Mall"* when what we actually mean is *"the people run from Charlton Way to The Mall, but the marathon stays in one place: the place between Charlton Way and The Mall"*. So really we should say "*the current that exists between A and B is positive*", but usually it is more important to convey the functioning of a circuit in an intuitive way, rather than to describe accurately what is actually going on inside the wires. So although we say *"current flows from anode to cathode"*, or *"current is a flow of electrons"*, it should not be taken too literally. Fans of science fiction may be pleased to learn that what really happens inside a valve, or indeed any electric circuit, is *a disturbance in the flux*!

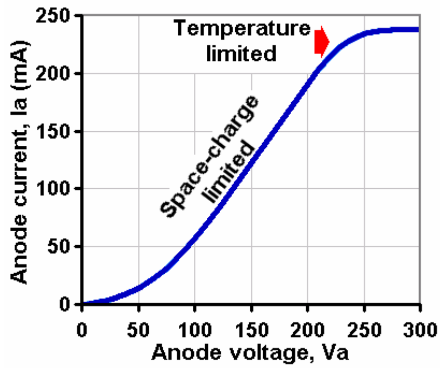

Fig 1.2: Static anode characteristic of a fictitious valve diode.

compared to the repulsion of the space charge. As the voltage increases the anode becomes more effective at drawing electrons out of the space charge, so they become faster and more numerous, so the curve bends upwards. If the anode voltage is made really high then it will attract all the electrons that the cathode is capable of emitting, so the curve levels off, and this is called **saturation**. Since this saturation region depends on the emission of the cathode which in turn depends on its temperature, the valve may be said to be **temperature limited**. The lower part of the graph is then termed **space-charge limited**. Valves are not normally designed to be operated at saturation as this can be damaging to the cathode's surface, so audio circuits always work in the space-charge limited region.

1.3: Triodes

By interposing a third electrode in between the anode and cathode we can exert some additional control over the anode current. This third electrode is called the **control grid** and is normally made from a spiral of fine molybdenum wire that encircles the cathode, held in place by a pair of copper support rods. Since we now have three electrodes (the heater doesn't count, remember) the device is called a triode. If a positive voltage is applied to the grid, electrons in the space charge will be drawn towards its electric field, but because the grid is very fine most of them will fly straight through the gaps and be captured by the electric field of the anode.

If the grid is made negative it will repel electrons, encouraging them to stay in the space charge, so anode current is restricted. If it is made very negative then the anode current can be choked off altogether, even though the anode may be hundreds of volts more positive. This condition is called **cut off**, and the grid voltage required to put the valve into cut off is sometimes called the **grid base**. Because the grid is very much closer to the cathode than the anode, a change in grid voltage has a more powerful effect on anode current than a similar change in anode voltage. This is related to a property called μ, which is discussed in section 1.5.

Thanks to the grid, for every value of anode voltage there is now a whole range of possible anode currents, depending on the grid voltage. The static anode characteristics of a triode can therefore be drawn as a family of curves, each corresponding to a different grid voltage, so they are called the **grid curves**. Fig. 1.3 shows the characteristics of an ECC83 / 12AX7 triode. Notice that the saturation region doesn't even appear on this graph as it is far above the working range.

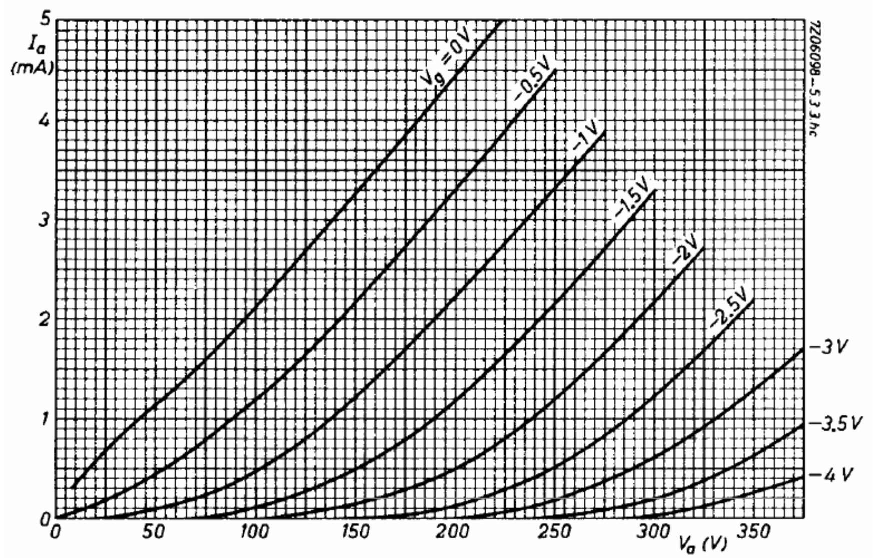

Fig 1.3: Static anode characteristic of the ECC83 / 12AX7. Taken from the Philips data sheet.

Another way to present exactly the same information is shown in fig. 1.4 where the curves correspond to different anode voltages, with grid voltage now on the horizontal axis. This graph is called the **mutual characteristics**. The static anode characteristic is generally the more useful form, however, but given either graph it is possible to draw the other.

There are three important characteristics of any triode, which are called (somewhat optimistically) the **valve constants**. These are anode resistance, amplification factor and transconductance. These constants are useful for understanding and analysing circuit behaviour, and for comparing the relative merits of different kinds of valve. All three can be derived from the static anode characteristics or mutual characteristics graphs, as shown in the following sections.

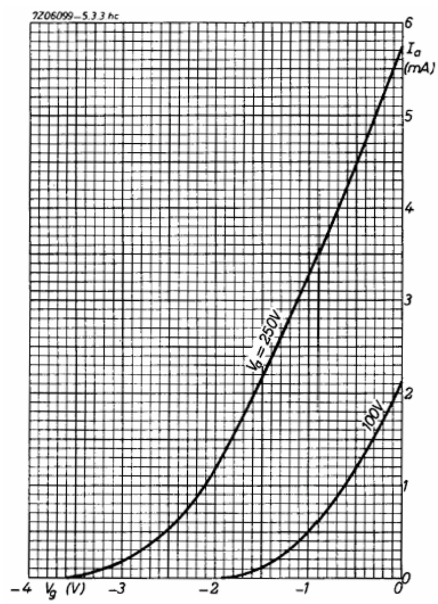

Fig 1.4: Mutual characteristic of the ECC83 / 12AX7. Taken from the Philips data sheet.

1.4: Anode resistance, r_a

Fig. 1.5 shows the static anode characteristics of the ECC83 again. At point A on the graph we have 225V across the valve and 1.6mA anode current flowing in it, so it could be imagined to be equivalent to a $225V / 1.6mA = 141k\Omega$ resistor. However, at point B it looks like a $150V / 0.5mA = 300k\Omega$ resistor, so clearly this device is nothing like a real fixed resistor; it is **non-linear**. Any of these values of resistance –taken at spot points on the graph– are called **beam resistance**. Beam resistance is only a DC value and is virtually useless to the circuit designer.

There is a much more useful way to figure the resistance of a valve. If we choose a point on the curve and then look at a small portion of the curve either side of that point it actually looks fairly straight (i.e., linear). For example, at point A, changing the anode voltage by 50Vcauses the anode current to change by about 0.9mA. Over this small section of the graph the resistance appears to be about $50V / 0.9mA = 55.6k\Omega$. This is the resistance that the valve presents to small variations in anode voltage (albeit superimposed on a steady DC voltage), so it is an AC property and is called the internal **anode resistance**. It is given the symbol $\mathbf{r_a}$, or $\mathbf{r_p}$ in America, and corresponds to the inverse of the slope of the grid curve. The steeper the slope, the lower this resistance, and wherever the slope of the curves is the same, r_a must also be the same.

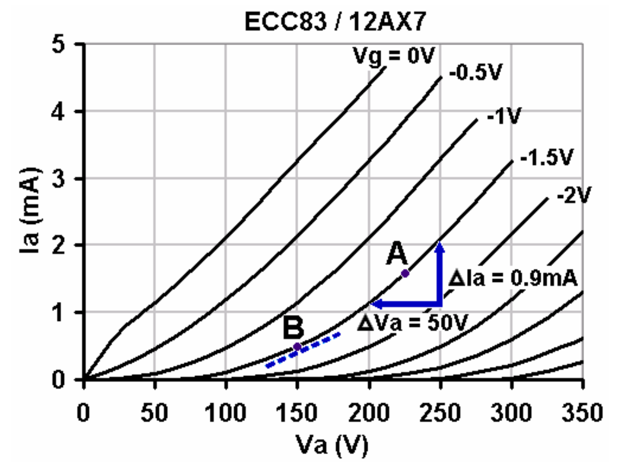

Fig 1.5: The AC anode resistance of a valve is the inverse of the gradient of the curve, or: $r_a = \Delta Va / \Delta Ia$

At point B the graph looks more curved than at A, but by drawing a tangent to the nearest grid curve (shown dashed) we can estimate the anode resistance to be about 83kΩ. This is greater than the previous figure because the curve is less steep here, but the values are at least more similar than the beam resistances we calculated earlier. Anode resistance is, therefore, a 'constant', roughly speaking.

1.5: Amplification Factor, μ

Looking at the static anode characteristic it is obvious that the grid curves have more-or-less the same shape as one another. Making the grid more negative

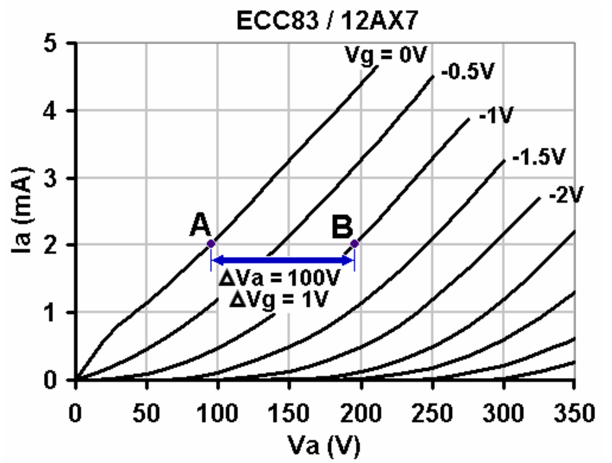

Fig 1.6: The amplification factor of a valve is the ratio of a change in anode voltage caused by a change in grid voltage, for a fixed anode current, or: $\mu = \Delta Va / \Delta Vg$

simply shifts the curve to the right, which reduces the anode current for a given anode voltage. However, the initial anode current could be restored by increasing the anode voltage.

For example, in fig. 1.6 an anode current of 2mA is achieved at an anode voltage of 90V when the grid is at 0V (point A). Reducing the grid voltage to –1V would reduce the current to 0.4mA, but it could be restored to 2mA by increasing the anode voltage to 190V (point B). The distance from A to B is 100V on the Va scale but only –1V on the Vg scale. The ratio of these two is called the **amplification factor** of the valve, and is given the symbol **μ** (mu). Another way to describe μ is to say that it is the measure of how much more effective the grid is than the anode at controlling anode current. In this case $\mu = 100V / 1V = 100$ (the minus sign is ignored). This figure is also the maximum possible voltage gain that the triode can achieve in any circuit.

We can see that μ is related to the horizontal separation of the grid curves, and over most of the graph the separation is about the same, so μ is another valve constant; in fact it is the *most* constant of the three constants! What's more, μ does not degrade much as the valve ages, unlike the other constants, which makes it a very useful figure for circuit designers.

1.6: Transconductance, g_m

The first two valve constants were derived by holding grid voltage constant (to find r_a) or by holding anode current constant (to find μ). The final important parameter is found by holding anode voltage constant. From fig. 1.7 we can see that if Va is 150V then changing Vg from –1V to 0V causes anode current to increase by about 2mA. This control that the grid voltage has over anode current is called the **transconductance**, or in older books the **mutual conductance**, and is given the symbol $\mathbf{g_m}$. Transconductance is related to the vertical separation of the grid curves, which is again moderately constant over much of the graph.*

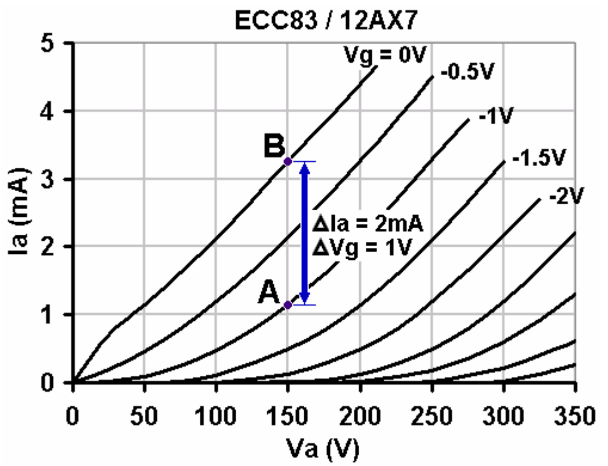

Fig 1.7: The transconductance of a valve is the ratio of a change in anode current caused by a change in grid voltage, with a fixed anode voltage, or: $g_m = \Delta Ia / \Delta Vg$

In this case the transconductance is $2mA / 1V = 2mA/V$. The units indicate that we get a 2mA change in anode current per volt change in grid voltage. The modern SI unit of conductance is the siemens (S) where 1S = 1A/V. Some old American texts may use the mho (ohm spelled backwards) which is equivalent to the siemens. In this case we could therefore write the transconductance alternatively as 2mA/V, 2mS (milli-siemens), 2000µmho (micro-mho), or 2000µA/V. However, the author prefers to stick to milliamps-per-volt in order to distinguish between *trans*conductance, which is a property of active components like valves, and pure conductance, which is simply the inverse of resistance. These are also the units used on most valve data sheets.

At any given point on the anode characteristics graph the three valve constants are related by van der Bijl's equation:[1]

$$\mu = g_m \times r_a \qquad (1.1)$$

So only g_m and μ need be found from the graph (these being the easiest to acquire accurate values for) while r_a can be calculated. As a valve ages, r_a tends to increase while g_m decreases. μ therefore remains more-or-less unchanged, falling only slightly.

* Incidentally, the transconductance also corresponds to the *gradient* of the mutual characteristics graph. When older texts talk about the 'slope' of a valve they are referring to its transconductance, not its anode resistance.

[1] van der Bijl, H. J. (1919) Theory and Operating Characteristics of the Thermionic Amplifier. *Proceedings of the Institute of Radio Engineers.* **7** (2) (April). p.109. This equation is often unfairly attributed to Barkhausen.

1.7: Amplification

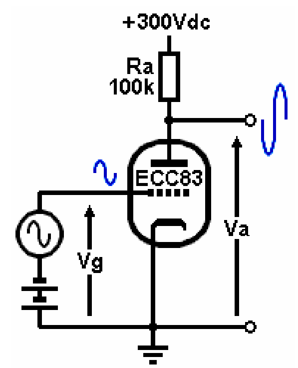

Fig 1.8: Basic voltage amplifier circuit.

From the anode characteristic graphs shown earlier it should be obvious that if we apply an audio signal voltage to the grid (i.e., vary the grid voltage) then current in the valve is bound to vary too. We can put this current variation to good use by putting a resistance in series with the valve so that a corresponding audio voltage is generated across the resistor, which is also called the **load**. Fig. 1.8 shows this in essence. Ra is the **anode resistor** (not to be confused with the anode resist*ance* r_a!) and provides a DC current feed from the HT, as well as serving as the load. The exact choice of this component is dealt with later. A battery is shown in series with the signal generator so that the AC audio voltage at the grid is superimposed on a small negative DC voltage; more on this in section 1.9. The cathode is grounded, fixing it at zero volts. Notice also that the cathode is connected to the input *and* output parts of the circuit, so it is 'common' to both. This circuit arrangement is therefore called a **common-cathode** amplifier.

An obvious question is: what is the voltage gain of this amplifier? If we knew the value of AC anode current caused by a given AC grid voltage then Ohm's law tells us that we could multiply this current by the value of Ra to find the AC output voltage (remember, these AC values are actually superimposed on steady DC values, but we can treat AC and DC parts separately). At first glance we might try to multiply the AC grid voltage by the transconductance to find the AC anode current: $i_a = v_g \times g_m$, but this will not work because g_m only applies when the anode voltage is fixed, whereas here it is changing. There are two approaches we can take: a graphical method or a more mathematical method. We will deal with the graphical method first, and the mathematical model later in this chapter.

1.8: The Load Line

The graphical method of analysing the circuit is to draw a **load line** on the static anode characteristics. This will shows us any and all of the possible voltages and currents in the circuit, and is the most powerful design tool in our armoury. Referring to fig. 1.8, imagine the valve is completely cut off so that no anode current flows. Hence there can be no voltage dropped across Ra, so the full HT reaches the anode. This point can therefore be plotted on the anode characteristics at:
Va = 300V , Ia = 0mA (point A in fig. 1.9).

Now suppose the valve is a short circuit (which is impossible in reality, but this needn't hinder our imagination). The full HT is now dropped across Ra with none across the valve. From Ohm's law the anode current would have to be

300V/100kΩ = 3mA. This point can also be plotted at: Va = 0V, Ia = 3mA (point B).

We could plot some intermediate values too, but because Ohm's law is a linear equation and the axes on the graph are linear too, there is no need; we can simply draw a straight line between A and B, completing the load line as shown in fig. 1.9. Examining the load line we see that it is intersected at various points by the grid curves. Each intersection shows us what Va and Ia will be for any given value of Vg, and these are the only combinations that are possible in this circuit. Load lines for other values of anode resistor could also be plotted for comparison (see section 1.15).

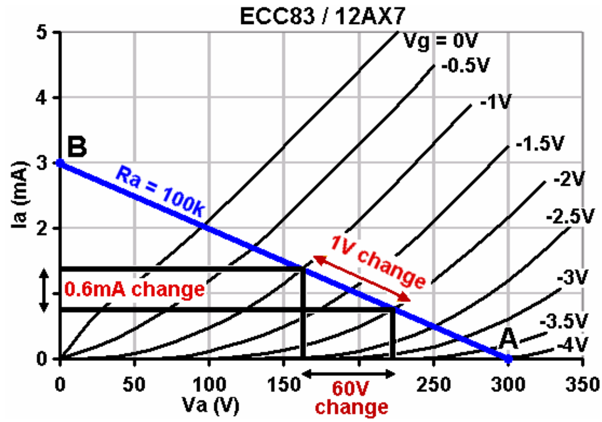

Fig. 1.9: A 100kΩ load line constructed on the ECC83 anode characteristics graph for a HT voltage of 300V. A small change in Vgk causes a change in Ia, which causes a large change in Va.

Having drawn the load line it is now possible to work out the voltage gain of the circuit. Suppose we made the grid two volts negative. We can deduce from the graph that 0.8mA would flow and the anode voltage would be 224V. If we then raised the grid voltage to one volt negative, current would increase to 1.4mA and the anode voltage would be pulled down to 164V. We have only changed the grid voltage by $2-1=1\text{V}$, but the anode voltage has changed by 224 – 164 = 60V. The gain is therefore:

$$A = \frac{v_{out}}{v_{in}} = \frac{60}{1} = 60$$

Note that this is less than the μ of the valve (about two thirds in fact, a ratio which will be explored in section 1.16). Also notice that when the grid voltage swings positive the anode voltage swings more negative, so the output signal will be up-side down or 180° out of phase compared with the input signal. This circuit is therefore referred to as **inverting**.

But there is something suspicious about this gain calculation. The grid curves are not evenly spaced along the load line, so if we did our calculation near the top where the curves are more stretched apart we would find a larger figure for gain than if we did it near the bottom where they are more bunched up. So which bit of the load line do we use? To answer this question we first need to understand how to bias the valve, and from there discover how the choice of bias influences the distortion generated by the circuit.

1.9: Biasing

Imagine a perfect sine wave. The wave has high peaks and low troughs. In between there must be some average level, so the wave could be said to be both rising above and falling below this average level. Now to speak in terms of voltage we have a sine wave that swings both positive and negative *relative* to some average voltage. This average voltage is, in a manner of speaking, our starting point, and it is up to us to decide where it should be on the load line in order that the valve can amplify both up-going and down-going parts of the signal, and this is called **biasing** the valve.

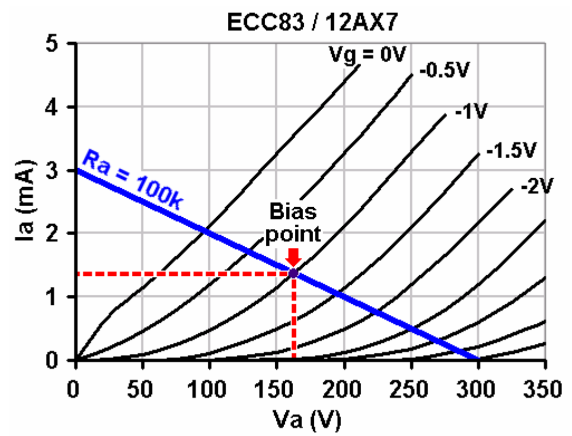

Fig. 1.10: A 100kΩ load line with Vg = –1V bias point indicated.

Suppose we fix the grid voltage at –1V and call this our **bias point**, as indicated in fig. 1.10. We are not yet inputting any signal, the valve is simply at rest. The valve is said to be in a state of **quiescence**, and we can see from the graph that the quiescent bias voltage of –1V causes a quiescent anode current of about 1.4mA and a quiescent anode voltage of about 163V. The bias voltage is the steady DC voltage onto which the audio voltage will be superimposed, and there are several ways to apply bias in practice.

A very direct method is shown in fig. 1.11. Here a bias battery is shown, but this could be substituted with a negative voltage supply from some other source. The negative voltage is applied to the grid via Rg which is the **grid-leak** resistor.* This resistor provides the voltage reference while preventing the audio signal from being shunted to ground via the bias voltage source, and is typically 1MΩ in value. **Coupling capacitor** C prevents the negative voltage on the grid from interfering with previous circuitry by blocking DC, so it is sometimes also called a **blocking capacitor**. C and Rg also form a high-pass filter so each will need to be suitably

* In modern parlance this might be called a **pull-down resistor** as it 'tethers' the grid to the required bias voltage, preventing it from collecting charge and floating up (or down) to some unknown potential.

sized to allow the desired audio frequencies to reach the grid (see chapter 4). Bias batteries are still sometimes used in valve hifi, and several years of life could be expected from a lithium battery since no DC current is actually required (current cannot flow in the grid while it is very negative). This type of biasing is called **grid bias** or **fixed bias** because the bias voltage is applied to the grid and (ideally) does not change during circuit operation.

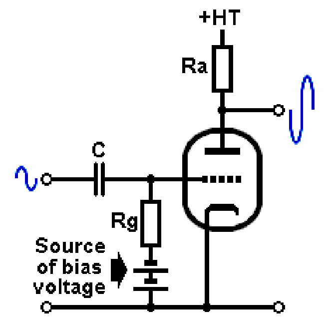

Fig. 1.11: Example of fixed bias. Coupling capacitor C prevents the bias voltage from interfering with previous circuitry.

Fig. 1.12 shows another form of biasing called **grid-leak bias** or **contact bias**. It was mentioned above that no current flows in the grid when it is very negative, but if the grid is only slightly negative –between 0Vand about –1V say– a few electrons leaving the space charge will succeed in hitting the grid, thereby generating a tiny current that flows (conventionally) into the grid, as indicated by the arrow. This **forward grid current** will generate a small negative voltage across the grid leak resistor that is sufficient to bias the valve by a small amount. Since the grid current is normally less than a microamp, Rg must be very large to generate any useful bias. This method is almost never used in modern circuits as it is not very predictable and results in excessive noise and distortion, although an exception will be met with the cascode in chapter 7.

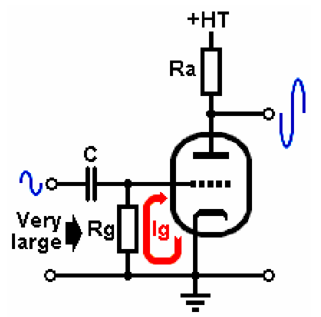

Fig. 1.12: Grid-leak bias takes advantage of the small but unavoidable forward grid current Ig.

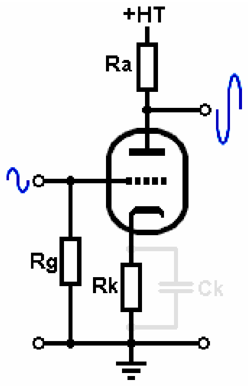

Fig. 1.13: Cathode biasing is the most common and convenient way to bias preamp valves.

Now let us reverse our way of thinking. Instead of making the grid negative with respect to the cathode, make the cathode *positive* with respect to the grid. This is a perfectly reasonable thing to do because the valve only reacts to the *difference* between the grid and cathode voltages; the *absolute* measured voltages are unimportant. It does not matter whether the grid is actually at a negative voltage while the cathode is at zero volts, or if the cathode is positive and the grid is at some lower voltage, the result is the same from the valve's point of view.

Since anode current flows in the valve we can put a resistor in series with the cathode and a voltage must develop across it, thereby raising the cathode voltage, so the valve provides its own bias. This method of biasing is shown in fig. 1.13 and is variously called **cathode bias**, **self bias** or **automatic bias**. Rk is called the **cathode resistor** or simply **bias resistor**. This form of biasing has several advantages:

- It is self adjusting. If the average current through the valve increases, so does the bias voltage, which will oppose the increase in anode current, and *vice versa*. This means that the bias will adjust itself naturally to changes in HT, valve aging, and variations in manufacturing tolerance. There is also less chance of the valve going into runaway and overdissipating in the event of some failure, which can be useful when dealing with power valves.
- The grid-leak resistor now pulls the grid down to ground (zero volts) so no coupling capacitor is required, at least, not unless we need to block some external DC voltage from coming *into* the circuit.
- Capacitor Ck (shown faint) can be added in parallel with Rk to control the gain, frequency response and linearity of the circuit. This is dealt with in section 1.18.

1.10: The Cathode Load Line

When designing a triode gain stage the usual procedure is to choose the triode, then chose a suitable anode load and draw the load line. Next, we choose the bias point and –if we are using cathode bias– find the necessary value of cathode resistor. The details of this whole process will emerge throughout this book, but for now let us concentrate on finding the cathode resistor. Fig. 1.14 shows a load line, repeated from fig. 1.9 earlier. Since we are now about to interfere with the cathode voltage we will no longer talk about the grid voltage Vg, but instead about the grid-to-cathode voltage Vgk, which is more correct.

If our chosen bias point is, say, $\mathrm{Vgk} = -1.5\mathrm{V}$ (point A) then this is the same as saying we want the cathode voltage to be raised to +1.5V. The graph indicates that the quiescent anode current will be 1.05mA. We now know the voltage across and current through the bias resistor, so Ohm’s law suggests a value of:

$$\mathrm{Rk} = \frac{\mathrm{Vk}}{\mathrm{Ia}} = \frac{1.5\mathrm{V}}{1.05\mathrm{mA}} = 1.4\mathrm{k}\Omega$$

This is not a standard value, so we would probably use 1.5kΩ as a close standard. This is really just a first approximation, however, because as soon as we add Rk to the circuit it will alter the load line, since Rk is now in series with the valve along with Ra.

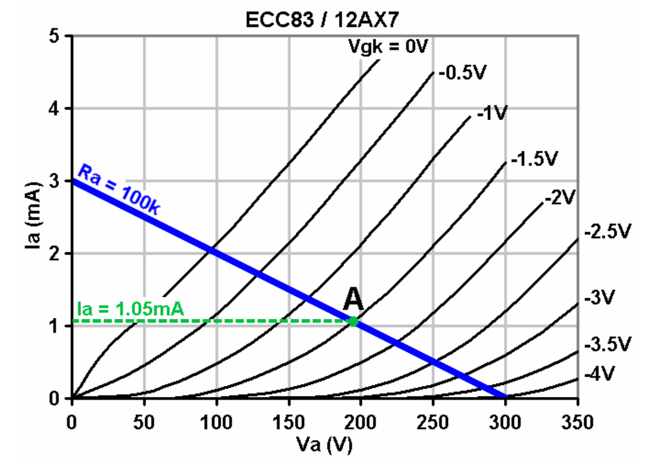

Fig. 1.14: Provided the cathode resistor is small compared with the anode resistor its value can be estimated as:
Rk = Vgk / Ia

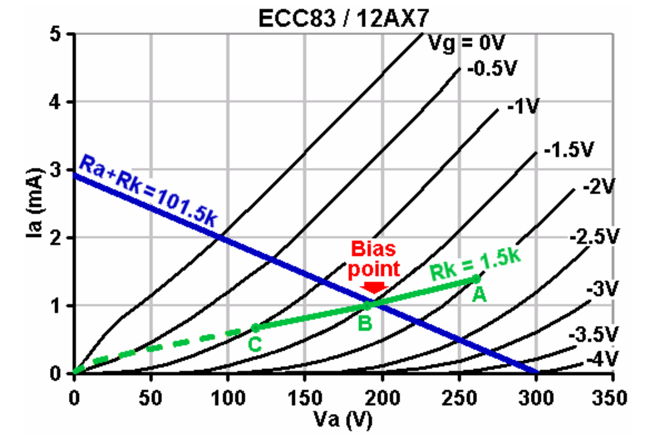

Fig. 1.15: Where the cathode load line crosses the main load line indicates the exact bias point.

Fortunately, we can double check our first approximation by redrawing the load line to correspond to the *total* resistance now in series with the valve, which is 101.5kΩ. This has been done in fig. 1.15 where it is obvious that because Rk is very small compared with Ra, it has made practically no difference to the slope of the load line, but we shall persevere for the sake of demonstration.

The main load line takes into account the total resistance in series with the valve and therefore shows the total range over which the valve can operate, but we can also draw a **cathode load line** to show where the bias point will be when part of that total resistance is used for Rk. This line will not be perfectly straight so we may need to plot a few points. First suppose that the cathode voltage were somehow 2V, what must the anode current be? Ohm's law guarantees:

$$Ia = \frac{Vk}{Rk} = \frac{2V}{1.5k\Omega} = 1.3mA$$

This point is plotted in fig. 1.15 at Vgk = –2V, Ia = 1.3mA (point A). Now repeat the process for different cathode voltages, say 1.5V and 1V:

$$Ia = \frac{Vk}{Rk} = \frac{1.5V}{1.5k\Omega} = 1.0mA \quad \text{(point B)}$$

$$Ia = \frac{Vk}{Rk} = \frac{1V}{1.5k\Omega} = 0.7mA \quad \text{(point C)}$$

Joining these points up gives us the cathode load line. Other points could be added to extend it, as shown dashed. Where the two load lines cross indicates the actual bias point. In this case we can see it is extremely close to our initial choice. With preamp valves we rarely need to go to the trouble of drawing a cathode load line because Rk will always be small compared with Ra, so the first approximation method will suffice. With power valves the bias is often more critical, so drawing cathode load lines can be useful.

1.11: Harmonic Distortion

Now that we understand what it means to bias the valve we can see what effect our choice of bias point has on the operation of the circuit. Consider fig. 1.16 which shows a 100kΩ load line and a bias point at –2V (B). Notice that this point is mid way between the 0V grid curve and the –4V grid curve where the valves reaches cut-off. Biasing in this area is called **centre biasing** and offers maximum headroom before the signal is clipped.

Now suppose we apply a 4Vp-p sine wave to the grid. When the grid voltage swings 2V positive the operating point will move up the load line and reach the 0V grid curve at point A; anode current peaks at 2mA and the anode voltage is pulled down to 95V. When the grid voltage swings 2V below the bias point the valve just reaches cut-off at point C, and the anode voltage hits 300V. Notice that we have swung the grid voltage symmetrically around the bias point, but the anode voltage has swung asymmetrically around its quiescent point because the grid curves are unequally spaced. The output waveform which has been plotted below the graph is bent out of shape and no longer a perfect sine

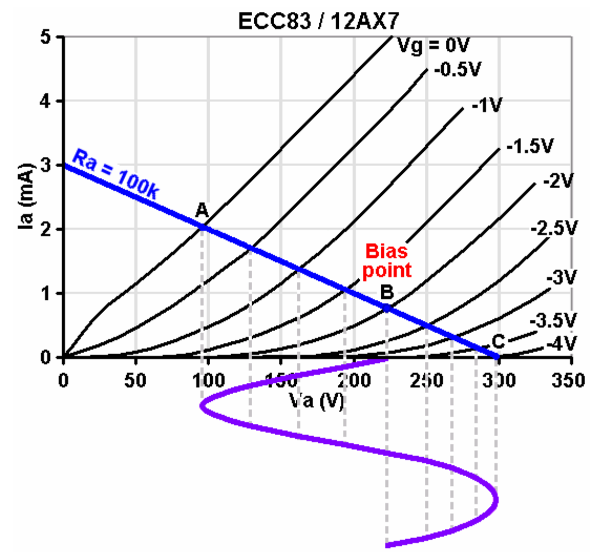

Fig. 1.16: Unequal spacing of the grid curves results in distortion, dominated by even harmonics.

wave; it is distorted. This natural bending of the waveform introduces new frequencies into the audio signal that were not there to begin with, and for the most part they will be integer multiples of the original audio frequency, which is called the **fundamental**. These new frequencies are called **harmonics** and are an important part of the classic valve tone.

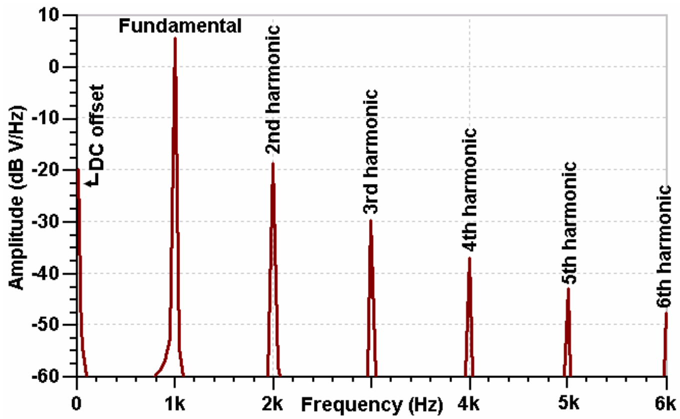

Fig. 1.17: The typical harmonic signature of a triode is a decaying series of all harmonics.

It is obvious from fig. 1.16 that the output waveform is stretched on the down-going half and squashed on the up-going half (don't forget that it is also inverted). Asymmetrical waveform shaping of this kind always introduces even-order harmonics (2nd, 4th etc.) more than odd-order harmonics (3rd, 5th etc.). The 2nd and 4th harmonics are two and four times the fundamental frequency respectively, or one and two octaves higher. They are therefore musically related to the original sound and tend to make it fuller and richer. Odd harmonics (and high-orders in general) are often not musically related to the fundamental and so are dissonant. In small amounts they add timbre and make the sound more cool and aggressive, while large amounts will tend towards a harsh and fuzzy sound. Triodes, therefore, can have a pleasant warming effect on tone because they produce a decaying series of all harmonics, dominated by the second, as shown in fig. 1.17. The larger the audio signal is, the more of the load line the operating point moves over and so the greater the distortion, and in most triodes the harmonic distortion is directly proportional to signal level. If multiple gain stages are cascaded then each will add another layer of distortion on top of the last, progressively building up a thicker tone by increasing the number and amplitude of harmonics.

The amount of distortion may be estimated from the load line. In this case the anode voltage has swung about 128V more negative but only about 77V more positive of its quiescent value. The percentage of second harmonic distortion is given by:[2]

---

[2] Langford-Smith, F. (1957) *Radio Designer's Handbook* (4th ed.), p491. Iliffe and Sons ltd., London. This formula is an approximation as it assumes *only* second harmonic distortion is produced. In reality other distortion products are present, though to much lesser degrees.

$$H_2\% = \frac{AB - BC}{2(AB + BC)} \times 100 \qquad (1.2)$$

$$H_2\% = \frac{128 - 77}{2 \times (128 + 77)} \times 100$$ = 12% second harmonic distortion.

In section 1.15 we will see how the choice of anode load resistance affects this figure.

We might wonder why fig. 1.17 shows a component at 0Hz (DC) even though it represents a pure AC signal. This is because the waveform has been squashed on one side and stretched on the other, so it no longer has equal power above and below the zero crossing. The *average* value of the waveform is therefore not zero as it would be for a symmetrical wave, but has an average DC offset, and this effect is called **self rectification**. On transient signals the offset does not have time to settle, so it is really a steady-state phenomenon. This has an important bearing on the tone of an amplifier because it can shift the average operating point of a valve dynamically with signal level, which can have an effect like a compressor. In the author's opinion this effect is the main contributor to 'the valve sound' that is held in such high esteem by musicians, and is seldom emulated in solid-state distortion circuits. However, if left unmanaged it can have unwanted effects too, which will crop up later in this book.

1.12: Intermodulation Distortion

When only a single input frequency is applied to a non-linear circuit the output will be distorted and will contain harmonics, just as described in the previous section. However, when *more* than one frequency is applied at the same time to a non-linear circuit the frequencies will interact with one another to produce new frequencies that are the *sum and difference* of the originals, and also the sum and difference of all the new frequencies produced! This is called **intermodulation distortion**, or **IMD**.

In a given gain stage harmonic distortion is the dominant form of distortion below the clipping point, but beyond clipping IMD rapidly takes over and becomes the stronger effect.* This kind of distortion is very objectionable for hifi reproduction as nearly all of the new frequency products will be dissonant with the original audio, and can rapidly make the sound 'mushy' or metallic. In other words, a little bit of IMD can sound a lot worse than a lot of harmonic distortion. In a guitar amp, however, IMD is an integral part of the overdriven tone, and arguably more influential to the 'valve sound' than simple harmonic distortion.

The mathematics of intermodulation is beyond the scope of this book, but there are some general rules that are worth pointing out. Circuits that produce even-order

* Strictly speaking, harmonic distortion is just a special case of intermodulation where the two input frequencies are the same.

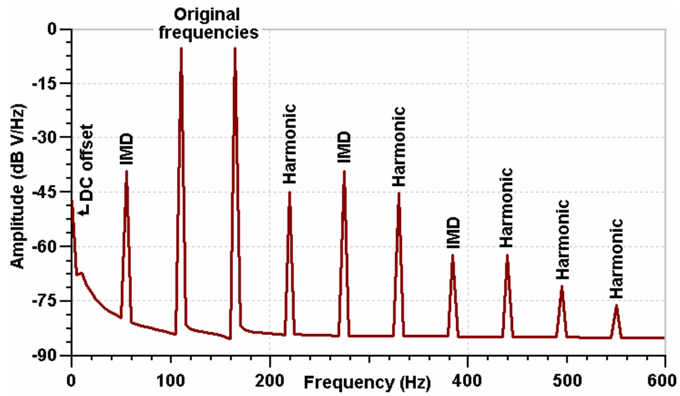

Fig. 1.18: Intermodulation distortion generates frequencies that are the sum and difference of all other frequencies, producing a complex distortion spectrum.

harmonic distortion (e.g., triodes and circuits that clip asymmetrically) produce both the sum and difference products. This means they will produce frequencies that are higher but also *lower* than the originals, thereby adding artificial bass to the sound. Fig. 1.18 shows an example distortion spectrum for a simulated ECC83 gain stage where the two original frequencies form the major fifth dyad A2 and E2 (110Hz and 165Hz). The strongest IMD products are at $110+165=275\text{Hz}$ and $165-110=55\text{Hz}$, which happen to be close to C♯3 (a perfect third) and A1 (a sub-octave)! This is another mechanism by which triodes and asymmetrical clipping circuits can make the tone seem fatter. Circuits that produce only odd-order harmonic distortion (e.g., symmetrical clipping circuits) do not produce the difference frequency and so lack this artificial bass.

IMD is an unavoidable and intrinsic part of an overdriven sound. However, the more notes we play at once the more dissonant the distortion products become, which is why dyads and power chords work well with distorted sounds while larger chords can sound muddied and indistinct. Generally, the heavier the distorted sound we want, the more we must rely on circuits that distort symmetrically, as these produce less overall IMD.

1.13: Cut-Off Clipping

Having looked at what happens for small signals when amplified by a non-linear circuit like a triode, we will now push things beyond the normal 'hifi' realm. Driving the valve beyond its normal limits of amplification is called **overdrive**, which is a word every guitarist knows. One way to overdrive the valve is to drive the grid very negative, towards the bottom of the load line where the grid curves become bunched up and gain begins to fall. If we continue to drive the grid negative then the valve will eventually reach cut-off and stop conducting completely, so the output waveform will level off. At this point we can say that the output wave is **clipped** (on its positive side since the stage is inverting). The grid voltage may continue to swing more negative but the valve can do nothing but remain in cut-off.

Fig. 1.16 suggests that cut-off ought to occur around Vgk = –4V, but it will usually take a little more than this. This is because electrons can still pass by the extremities of the grid where its electric field has less influence even when, in theory, current should already have ceased to flow. This is known as the **island effect**. Because cut-off is delayed in this way, the onset of cut-off clipping is not sudden but instead leaves a rounded edge on the waveform. This is quite unlike solid-state circuits that nearly always clip very hard. Sharp corners on any waveform indicate the presence of high frequencies, which are generally unpleasant, but since the valve soft-clips the signal this high frequency content is naturally constrained, even for high levels of overdrive. It is extremely difficult to emulate this soft clipping using solid-state electronics, and this is another reason why valve overdrive remains the distortion of choice for discerning guitarists.

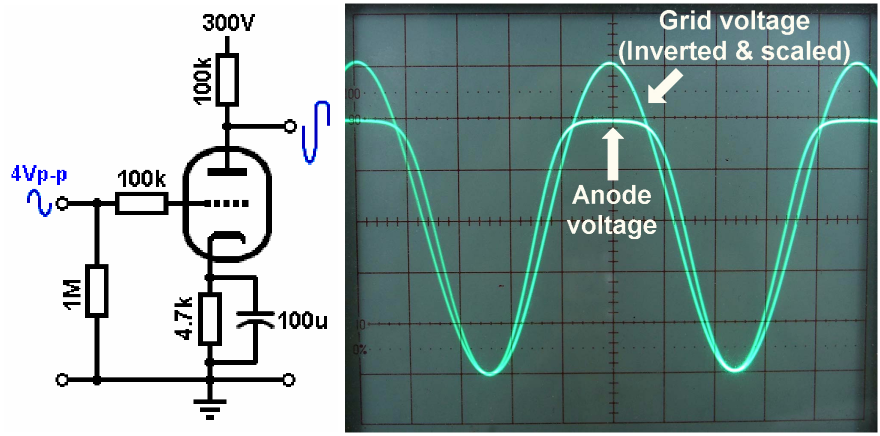

Fig. 1.19: Cut-off clipping in a cold-biased ECC83.

Fig. 1.19 shows a test circuit along with the waveforms viewed on the grid and anode. The grid voltage has been inverted and scaled to allow easy comparison with the anode waveform. The quiescent bias was –2.5V, but the DC shift in the distorted waveform caused the *average* anode current to increase, so the cathode voltage rose to about 2.8V in this case. Note that the transition into and out of clipping is relatively soft. Biasing a valve close to cut-off like this results in less quiescent power dissipation in the valve, so is called **cool biasing** or **cold biasing**.

1.14: Grid-Current Clipping

There is an obvious question that some readers have probably been asking since the start: why are there no more grid curves shown to the left of the Vgk = 0V curve? Curves for positive Vgk values do indeed exist, but they are not usually included on the graph because it is not practical to operate the valve in this region. The reason is that when the grid voltage approaches the cathode voltage, electrons being drawn from the space charge get attracted to the grid rather than to the anode (since the grid is very much closer), rather as if we had placed a diode between grid and cathode. Conventional current flows in the opposite direction –into the grid and down through the cathode– so we have forward grid current. This current causes a voltage drop across whatever resistance happens to be in series with the grid, making it harder to drive the grid beyond Vgk = 0V. To put it crudely, less of our input voltage actually makes it to the grid. The more we attempt to make the grid positive the more current flows into it to prevent us from doing so, as if some invisible volume control is being suddenly turned down. In more technical terms, the input impedance of the valve suddenly falls from many megohms to a few kilohms or less, and this effect is known as **grid-current limiting**.[*] The grid signal will therefore be clipped on the positive side, but since the stage is inverting the anode signal will appear clipped on the negative side. But it must be fully understood that it is actually the *grid signal* which is being clipped, while the valve continues to amplify what appears on its grid quite normally.

Grid-current limiting does not happen instantly. Forward grid current actually begins to flow even before the grid reaches the cathode voltage, around Vgk = –1V, though this varies considerably between valve types. The data sheet may even quote this as ‘Vgk$_{(max)}$’, indicating the point at which grid current will begin to exceed 0.3μA, and for the ECC83 this is –0.9V. The more resistance we put in series with the grid the greater the voltage drop caused by this grid current, and therefore the harder and

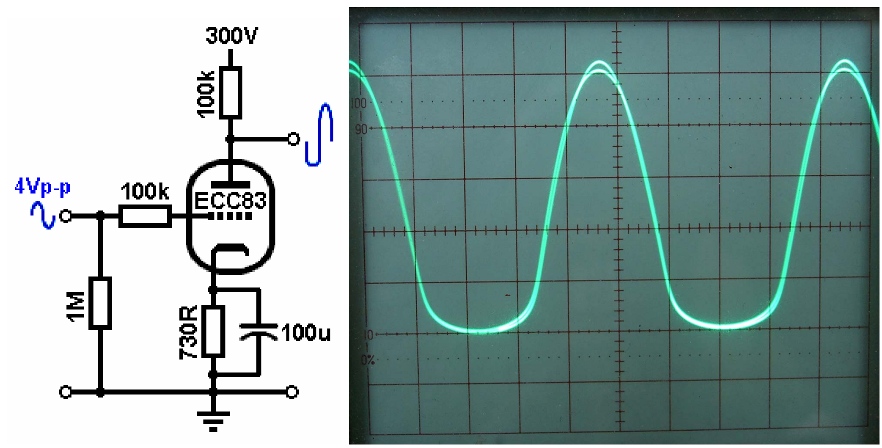

Fig. 1.20: Grid-current clipping in a warm-biased ECC83. The grid voltage waveform has been inverted and scaled for comparison with the anode waveform; they are nearly identical.

[*] Some modern texts may misleadingly refer to this as **saturation** because it resembles the clipping effect produced in a saturating transistor. However, saturation has a more specific meaning for valves, as we saw in section 1.2; it is the absolute maximum current that can be drawn from the cathode. This is well beyond the current levels ever found in a normal triode circuit.

more abrupt the clipping. Additionally, the amount of grid current increases as anode voltage falls since the anode has less ability to draw electrons away from the grid at low voltages. We can therefore control the hardness of grid-current clipping with circuit resistances, at least to some extent. This is a very important observation that we will revisit throughout this book.

Fig. 1.20 shows another test circuit along with the waveforms viewed on the grid and anode. The quiescent bias was –1V but this time the waveform is clipped on the opposite side from cut-off clipping so the DC shift is in the opposite direction, reducing the average anode current, so the cathode voltage fell to about 0.9V in this case. The grid voltage is inverted and scaled as before, and it can be seen that both signals are nearly identical because the valve is amplifying just what it sees on the grid. It is also clear that the grid-current clipping is slightly softer than the cut-off clipping seen earlier, but the opposite may be true for triodes with higher grid current than the ECC83, such as the ECC81. Biasing a valve close to grid-current limiting results in more quiescent power dissipation in the valve, so is called **warm biasing** or **hot biasing**.

1.15: The Effect of Load on Distortion

It has been hinted at already that the value of anode load resistance has some influence over the distortion generated by the stage. To illustrate this, fig 1.21 shows load lines for three possible resistances: 220kΩ, 100kΩ and 47kΩ. Clearly, a smaller resistance produces a steeper line. If we made the resistance very large then the load line would rotate down towards the horizontal axis where anode current is very low and the grid curves become bunched up. In this region performance becomes unpredictable, so it is generally avoided. With preamp valves we are unlikely ever to use anything larger than about 500kΩ. A bias of –1.5V has been chosen for this example, as indicated by the dot on each load line, which is somewhat warmer than centre. The stage will therefore reach grid-current limiting before cut-off.

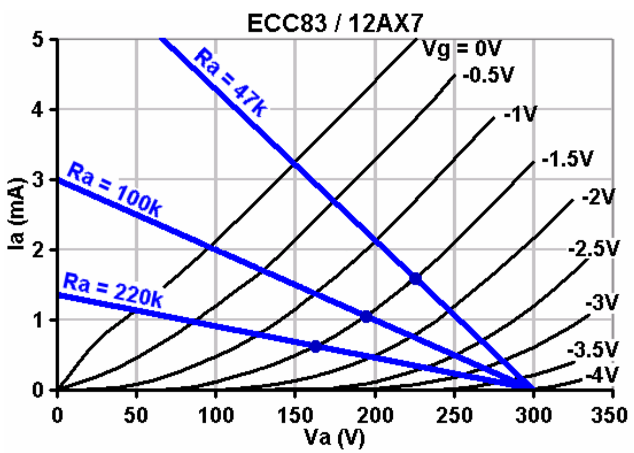

Fig 1.21: Load lines for three different values of anode resistor. Larger resistances reduce harmonic distortion.

Examining the figure we see that the greater the load resistance the larger the possible output voltage swing. In

Load resistance	Voltage gain	Maximum output swing	2^{nd} harmonic distortion
47kΩ	43	130Vp-p	7.7%
100kΩ	60	180Vp-p	4.3%
220kΩ	68	205Vp-p	3.7%

Table 1.1: Figures relating to fig. 1.21.

each case the *range* of possible bias voltages is exactly the same (because the HT is the same for each), from which we deduce that the voltage gain must be greater when using larger load resistances. Secondly, the grid curves are more evenly spaced along the 220kΩ load line than along the 47kΩ load line. More equal spacing means less distortion, at least up until clipping. Table 1.1 summarises the figures. In general, a larger load resistance will increase gain and available output swing, but reduce harmonic distortion.

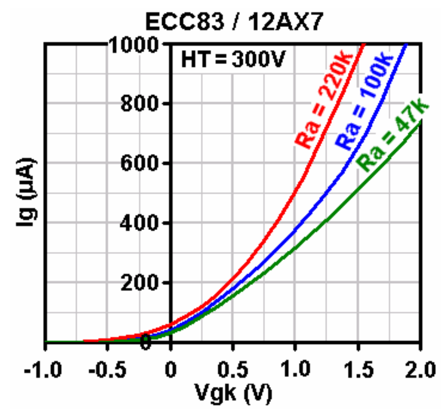

Fig. 1.22: Increasing the anode load resistance causes grid current to increase, which may lead to harder grid-current clipping.

Another effect which is not obvious from fig. 1.21 is that with a larger load resistance the anode voltage is pulled down lower as the 0V grid curve is approached, and therefore grid current will be greater when it does eventually start flowing. This is shown by fig. 1.22 which gives the average grid current characteristics of the ECC83 with a 300V supply voltage and different values of anode load. Increased grid current will result in slightly harder grid-current clipping, which likewise will lead to a slightly harder or crisper overdrive sound.

1.16: The Golden Ratio

When the anode load resistor is made equal to twice the internal anode resistance r_a, something special happens. Consider fig. 1.23 which shows the anode characteristics of a 'perfect' triode and a load line corresponding to $Ra = 2r_a$. As indicated, the peak anode current that can flow (when Vgk = 0V) is now exactly:

$$I_{peak} = \frac{HT}{(Ra + r_a)} = \frac{2}{3} I_{max}$$

The centre-bias point which gives maximum headroom is therefore half this value or $(1/3)I_{max}$ and results in a quiescent anode voltage of $(2/3)HT$. The voltage gain rather neatly becomes exactly $(2/3)\mu$ and it so happens that this is *also* the condition

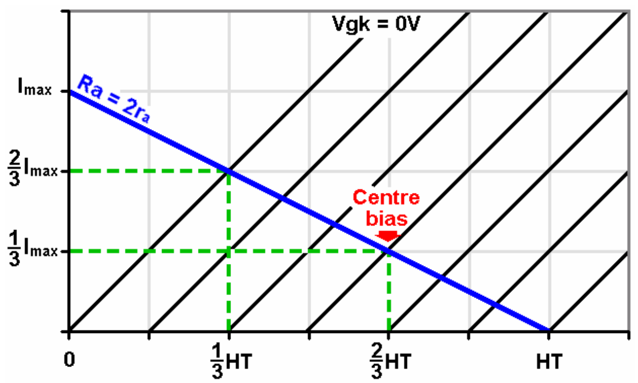

Fig 1.23: Characteristics of a theoretically ideal triode. When Ra = $2r_a$ the centre bias point occurs where the anode voltage is (2/3)HT, and the voltage gain becomes (2/3)μ.

that results in maximum power being delivered to the load! This last point is only of importance in a power amp of course, not a voltage amplifier.

A real triode is not perfectly linear and so in practice this rule does not hold exactly, but in the days when a designer might have only a volt meter and no access to a full data sheet and characteristic curves, he could still use this 'golden ratio' of Ra = $2r_a$ to build a fairly respectable hifi circuit. After choosing Ra the bias could be adjusted until the anode voltage was (2/3)HT at which point he could be reasonably sure of maximum headroom, and that the gain was two-thirds μ (the valve constants could be found in quick-reference tables that were published more widely than full data sheets). Of course, we do not have to obey this rule, especially not in a guitar amp, but it is mentioned here as a historical reason why 100kΩ is so universally used with the ECC83 / 12AX7; it is a nice, round, standard value that is also roughly twice r_a.

1.17: The AC Load Line

The signal voltage appearing at the anode is superimposed on the average anode voltage, so our output signal is really a varying DC voltage. But in most cases what we actually want is a pure AC signal voltage to feed to the next stage of the amp, so an output coupling capacitor Co is added to block the DC voltage while letting the AC signal pass through, as shown in fig. 1.24. This is called **CR coupling**, **capacitor coupling** or usually **AC coupling**. Co also forms a high-pass filter with the following resistance to ground, Rl, so signals below the cut-off frequency will be attenuated. For now we will assume Co it is large enough[†] to pass all audio frequencies.

As far as DC currents are concerned nothing has changed, and the anode load is simply Ra. AC currents, however, can now flow in Rl via Co, so the total anode load for AC is the parallel combination of Ra and Rl, or:

[†] By 'large' the author means high capacitance, not physically large.

$$R_{ac} = Ra \,||\, Rl = \frac{RaRl}{Ra + Rl} = \frac{100k \times 39k}{100k + 39k} = 28k\Omega$$

We can show this on the static anode characteristics by first drawing the ordinary DC load line and choosing a bias point somewhere on it, and then draw an **AC load line**. Since the coupling capacitor does not affect the DC conditions the AC load line must pass through the bias point. Fig. 1.25 shows the DC load line and a bias point of –1.5V. One way to find the AC load line is to suppose the anode voltage changes by 100V. From Ohm's law we know the current in the AC load must change by: 100V / 28k = 3.6mA . We can now plot a point that is 100V lower than the quiescent anode voltage, and 3.6mA higher than the quiescent anode current. This is labelled A, and the AC load line is then drawn through A and the bias point.

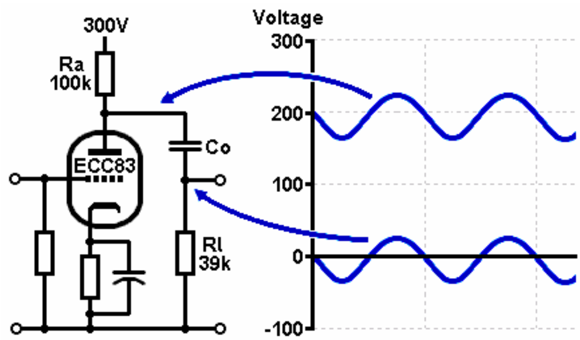

Fig 1.24: Coupling capacitor Co blocks the DC anode voltage. The DC load on the valve is simply Ra, but the total AC load is the parallel combination of Ra and Rl.

The AC load line only applies while an AC signal is being amplified or, in other words, while the valve is working under dynamic (rather than static) conditions. It is therefore sometimes called the **working load line**. We can see from the figure that the DC load line has a slope of –1/Ra, but the addition of the AC load causes the working load line to rotate clockwise around the bias point, and it now has a steeper slope of –1/(Ra||Rl). This affects operation in several ways. It is obvious that the harmonic distortion is increased while the output swing and gain

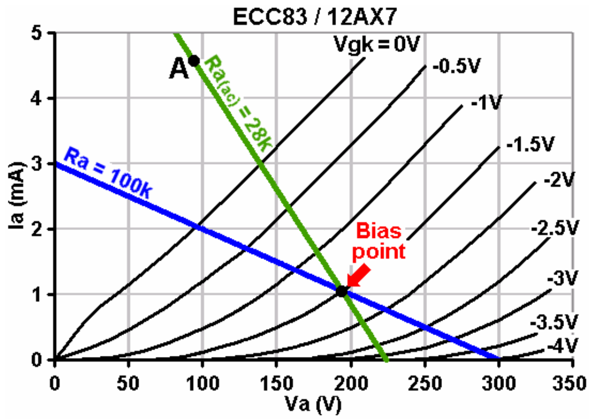

Fig 1.25: The AC load line rotates around the bias point, causing loss of gain and a change in output signal symmetry.

are reduced. In fact, the gain is reduced by exactly half in this case, and the reason for this will become clear in section 1.21. In this case an extreme value for Rl was chosen to illustrate the changes in operation clearly, but in practice Rl will normally be made much larger to avoid loading the valve so heavily, so the working load line will not rotate much.

A further important change is that the point of centre-bias has moved. Although the bias appears warm when looking at the DC load line, it is actually about central on the AC load line because the anode voltage cannot swing so high any more. Table 1.2 outlines these figures.

Load line	Voltage gain	Maximum output swing	2nd harmonic distortion	Centre bias point
DC (100kΩ)	60	180Vp-p	4.3%	–2V
AC (28kΩ)	30	90Vp-p	11%	–1.5V

Table 1.2: Figures relating to fig. 1.25.

Astute readers may wonder how the peak current in the AC load can be greater than the current actually supplied by Ra. This is possible because when the circuit is first switched on Co charges up and stores energy. When the anode voltage swings low the ‘extra’ current is sourced from the coupling capacitor, and when the anode swings high again current in Ra is diverted back into Co to recharge it.

1.18: The Cathode Bypass Capacitor

Now that we have a better understanding of how a triode amplifier works we can turn our attention back to the cathode bias circuit. We already know that when the grid voltage is increased the anode current also attempts to increase, and this in turn can be used to develop a much larger voltage signal at the anode. But the anode current also flows in Rk, so when the anode current increases, the voltage across Rk also increases, and likewise it must fall when the anode current falls. The cathode voltage, therefore, attempts to follow the grid voltage, so the voltage difference *between* grid and cathode will be less than the applied signal voltage. But what the valve actually amplifies *is* the difference between grid and cathode, so by adding Rk the valve now has a smaller signal to amplify and the gain of the circuit will be less than the load line suggests. This effect is called **cathode current feedback** or **cathode degeneration**, and as well as reducing the overall gain it also reduces distortion and increases headroom.

Less distortion and more headroom may be just what we want in some circuits, but not always. To eliminate the effect of cathode current feedback we need to stop the cathode voltage from following the grid, but we don’t want to alter the DC bias voltage. This is easily done by adding a **cathode bypass capacitor**, Ck, in parallel with Rk, as shown in fig. 1.26. Any rapid rise in cathode current will now be diverted into charging the capacitor, and if cathode current falls the capacitor will supply the deficit from its own charge. Another way of looking at it is to say that the capacitor **decouples** or bypasses to ground any AC signals on the cathode, so signal current does not flow in the cathode resistor and the DC bias voltage remains unchanged. With either explanation the result is the same; the cathode bypass capacitor smoothes out changes in cathode voltage, helping to hold the cathode voltage constant and thereby preventing cathode feedback.

Now, a capacitor will allow greater current flow at high frequencies than it will at low frequencies. If we want the stage to have maximum gain at all audible frequencies then the capacitor must be large enough to smooth out the lowest frequencies of interest, and the stage is then described as **fully bypassed**. If the capacitor is made relatively small then only high frequencies will be smoothed out while lower frequencies will not. Therefore the stage will have maximum gain at high frequencies and minimum gain at low frequencies, producing a treble boost, and the stage would then be called **partially bypassed**. If the stage has no cathode bypass capacitor then it is said to be **unbypassed** and will have minimum gain. This ability to control the gain between lower and upper frequencies is an extremely useful consequence of using cathode bias.

The exact relationship between the gain of the stage with frequency and the size of the cathode bypass capacitor is not straightforward,[3] but the shape of the frequency response *is*; it forms a 'shelved' response like that shown in fig. 1.26. The frequency at which the gain starts to rise from its lower level is given by:

$$f_{lo} = \frac{1}{2\pi \cdot Rk \cdot Ck} \tag{1.3}$$

The response then rises at a rate of 20dB/decade (6dB/octave) until it levels off at the

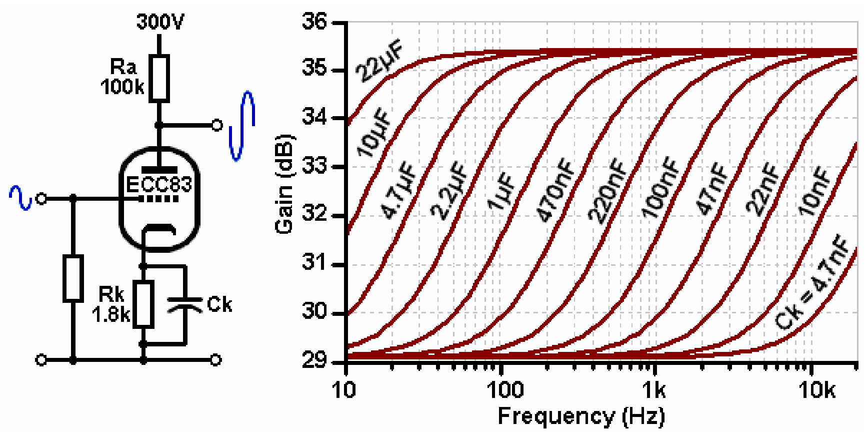

Fig 1.26: Typical gain stage and its frequency response with different values of cathode bypass capacitor Ck.

[3] Blencowe, M & James, D. I., (2008). Choosing Cathode Bypass Capacitors. *Audio Xpress*, (August), pp19-20.

maximum available gain of the circuit, which is the value we would get from the load line. Fig. 1.26 shows how the frequency response of a typical ECC83 gain stage changes with different values of Ck. Notice that halving or doubling the value of Ck shifts the whole response up or down by exactly one octave.

Equation 1.3 indicates that both Rk and Ck affect the frequency response. Making either of them larger will shift the response down in frequency, but altering Rk will also affect the bias. Rk also controls the minimum gain level, and a formula for finding this is given in the next section. Other types of triode will give different values of maximum and minimum gain depending on their own anode and cathode resistors, but the *shape* of the frequency response is always the same.

1.19: Equivalent Circuits

Using graphical methods to work out gain and general circuit operation is highly instructive, and with practice it is possible to 'see' what affects a circuit by imagining load lines. However, it is also convenient to have some more universal equations that we can use for quick calculations, or even a general method for generating equations that apply to our own particular circuits. This can be done by drawing a simplified version of the circuit in which the rather complicated active devices like valves or transistors are instead replaced by imaginary voltage or current generators. At this stage we are really interested in AC signals rather than DC conditions, so all our imaginary generators will be purely AC signal generators. Then, using fundamental formulae, we will be able to work out what signals will appear elsewhere in the circuit and so calculate things like gain and frequency response. Without going into the fundamental theorems which can be found in plenty of other textbooks, we will simply state the facts:

- A valve can be modelled as a *voltage* generator equal to $-\mu \times v_{gk}$, in *series* with a resistance equal to r_a. This is called a **Thévenin equivalent**.
- A valve can also be modelled as a *current* generator equal to $-g_m \times v_{gk}$, in *parallel* with a resistance equal to r_a. This is called a **Norton equivalent**.

Both models will ultimately give the same results, but sometimes one is a little easier to handle than the other, depending on the circuit.

Fig. 1.27 shows an example of the steps we might take to generate an equivalent circuit. The image in **a.** is a simple gain stage. A power supply smoothing capacitor Cs has been included for reasons that will emerge. The first step is to convert the valve into either its Thévenin or Norton equivalent, and the Thévenin model is often the most convenient for triodes, which takes us to **b.** Now, we are only interested in AC signals, so if we imagine the frequency to be high enough that any capacitors have very low reactance then we can treat them as short circuits. This takes us to **c.**, and notice that by replacing Ck with a short circuit we have also eliminated Rk. The last step is really just tidying up except for a small modification: since Rk has been eliminated we find that v_{in} is in fact the same as the grid-to-cathode voltage, so it has been relabelled v_{gk} in **d.**

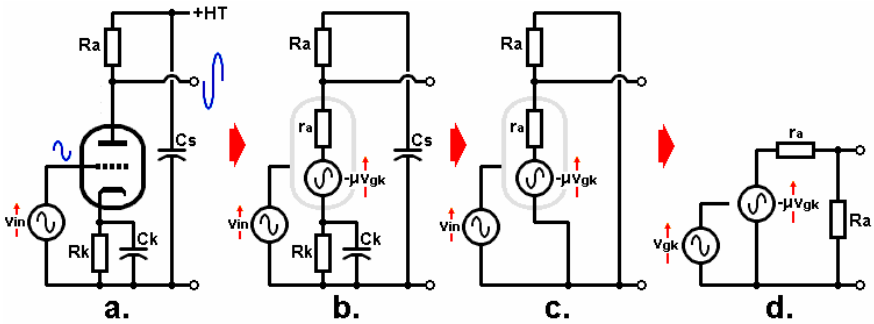

Fig. 1.27: Fully bypassed gain stage broken down into its Thévenin equivalent.

The equivalent circuit is obviously a lot simpler than what we started with. The valve now looks like a voltage generator that takes whatever signal voltage appears between grid and cathode (v_{gk}) and multiplies it by $-\mu$ (the minus sign appears because the stage is inverting). Because the power supply capacitor Cs looks like a short circuit we find that Ra –which is in reality connected between HT and the anode– appears to be connected between anode and ground. In other words, as far as AC is concerned the power supply and ground are *one and the same*, so r_a and Ra form a potential divider. The final output signal is therefore the generator voltage multiplied by the gain of this potential divider:

$$v_{out} = -\mu v_{gk} \frac{Ra}{Ra + r_a}$$

In this case v_{gk} is exactly equal to the input voltage so we can also write:

$$v_{out} = -\mu v_{in} \frac{Ra}{Ra + r_a}$$

The voltage gain of a circuit is v_{out}/v_{in}, so by dividing both sides by v_{in} we quickly obtain the gain of a fully bypassed triode gain stage:

$$A_{(bypassed)} = -\mu \frac{Ra}{Ra + r_a} \qquad (1.4)$$

We must not forget the limitations of this model, however. Firstly we said that the frequency was high enough for capacitors to look like short circuits, but at lower frequencies we may need another model that includes the capacitors, in which case we get into filter theory. The model also assumes that the valve is linear, which is not true in reality. But as long as the signal amplitudes are small the valve constants don't vary much and the valve appears to be at least approximately linear, which is why this is also referred to as **small-signal** equivalent circuit.

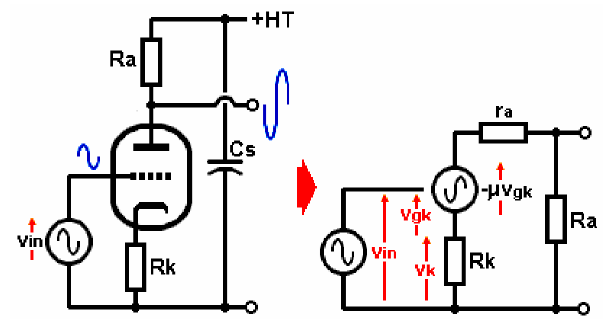

Fig. 1.28: Unbypassed gain stage and its Thévenin equivalent circuit.

The previous model assumed that Rk was fully bypassed. If no bypass capacitor is used or if the frequency is low enough for it to act like an open circuit, then Rk will cause cathode current feedback and the gain will be reduced. For this we need a new equivalent circuit, as shown in fig. 1.28. From Ohm's law the current flowing around the circuit must be:

$$i = \frac{v}{R} = \frac{-\mu v_{gk}}{Ra + r_a + Rk}$$

But unlike the previous circuit v_{in} is not equal to v_{gk} because some of it is cancelled out by the opposing voltage appearing across the cathode resistor, or in other words: $v_{gk} = v_{in} - v_k = v_{in} + iRk$. Substituting this into the previous equation and simplifying gives:

$$i = \frac{-\mu v_{in}}{Ra + r_a + Rk + \mu Rk}$$

(notice that the cathode current feedback has an effect like adding an extra resistor μRk to the output circuit). By multiplying this current by Ra we get the output voltage, and then also dividing both sides by v_{in} we find the overall gain of the circuit:

$$A_{(unbypassed)} = -\mu \frac{Ra}{Ra + r_a + Rk(\mu + 1)} \quad (1.5)$$

In case readers are now worrying that they will have to go through this process every time they design a guitar amp, they needn't. Although circuit analysis is a valuable skill it is possible to avoid it by using approximations, general equations given in this book, and a little real-life experimentation. Equations 1.4 and 1.5 were derived explicitly simply to show the usefulness of equivalent circuits.

1.20: Input and Output Impedance

Whenever one circuit is coupled to another it is important to know what effect this will have on the circuit being driven, or on the circuit doing the driving. The **output impedance** of the driving circuit, when feeding the **input impedance** of the next circuit, forms a potential divider as illustrated in fig. 1.29. Depending on the context these impedances may also be called the **source impedance** and **load impedance** respectively.

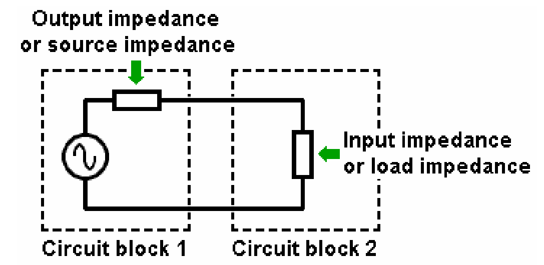

Fig. 1.29: Output or source impedance forms a potential divider with input or load impedance.

For maximum *power* transfer the two impedances need to be the same, which is called **impedance matching**. However, in a preamp we are really interested in signal *voltage*, and for maximum voltage transfer the input impedance must be as large as possible with respect to the source impedance, which is called **impedance bridging**. Although it wasn't stated at the time, we have already seen a prime example of this in fig. 1.27. There we saw that the output impedance of the valve, r_a, formed a potential divider with the load resistor Ra, so the smaller Ra is, the lower the voltage transfer (i.e. gain) becomes.* However, it is perhaps more usual to consider the *whole* gain stage as a single functional block that will be connected to other circuit blocks, so it is important to know what the input and output impedances of a gain stage are before we go connecting it to anything else. These things can also be found with equivalent circuits.

1.21: Output Impedance of a Triode Gain Stage

Fig. 1.30a shows the Thévenin equivalent circuit of a fully bypassed gain stage (repeated from fig. 1.27d). To find the output impedance we first set any **independent sources** to zero. In this case the only 'fixed' or independent source is v_{gk}, so setting this to produce zero voltage makes it a short circuit, taking us to **b**. This leaves the **dependent source** $-\mu v_{gk}$, but since there is now nothing for it to amplify this source also becomes zero, as in **c**. Finally we can tidy up the drawing as in **d.** and find the total impedance between the output terminals. In this case it is simply r_a and Ra in parallel, so we immediately know from the product-over-sum

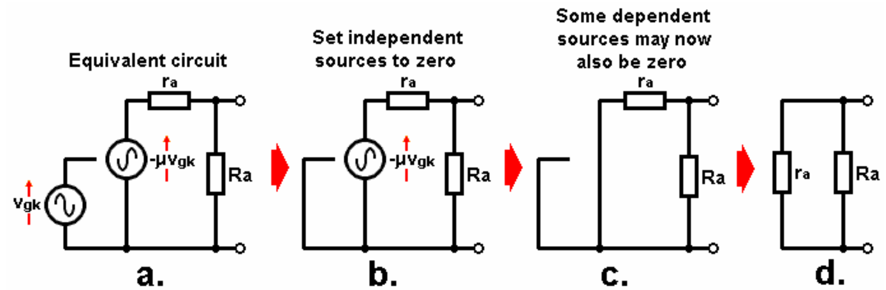

Fig. 1.30: Finding the output impedance of a fully bypassed gain stage.

*One way to imagine output impedance is to suppose we applied a 1V signal to the *output* of the circuit; how much current would flow into it? The ratio of this enforced voltage and the resulting current is the output impedance.

rule (see preface) that the output impedance of a fully bypassed gain stage is:

$$Zo_{(bypassed)} = \frac{Ra \times ra}{Ra + ra} \tag{1.6}$$

As an example, let us try some typical values for an ECC83 / 12AX7. If $Ra = 100k\Omega$ and $r_a = 65k\Omega$ the output impedance becomes:

$$Zo_{(bypassed)} = \frac{100 \times 65}{100 + 65} = 39k\Omega$$

This means that if the circuit is connected to one with a 39kΩ input impedance we will have a 1:1 potential divider, so half the available signal voltage will be lost across the output impedance or, in other words, the gain of the overall circuit would appear halved. This is exactly what the AC load line in fig. 1.25 showed, since it represented an extra 39kΩ load (which is exactly why it was chosen as an example!).

For an unbypassed gain stage the outcome is a little different because one of the properties of feedback is that it alters the input and output impedances. Fig. 1.31a shows the equivalent circuit, and setting v_{in} to zero takes us to **b**. However, we cannot immediately eliminate the dependent source because v_{gk} is not zero; it is in fact equal to the voltage across Rk. Without going through the full treatment, we already saw in the previous section that Rk has the effect of adding yet another resistor to the circuit which is μ times larger than Rk. This is represented in **c.**, and we find that the output impedance is $Ra \parallel (r_a + Rk + \mu Rk)$ or:

$$Zo_{(unbypassed)} = \frac{Ra(r_a + Rk(\mu + 1))}{Ra + r_a + Rk(\mu + 1)} \tag{1.7}$$

So an unbypassed cathode resistor causes the output impedance to increase.

Strictly speaking, what we have actually found is the output *resistance*. The output *impedance* would include the effects of stray capacitance and inductance, but these things are so small that we usually ignore them for guitar purposes.

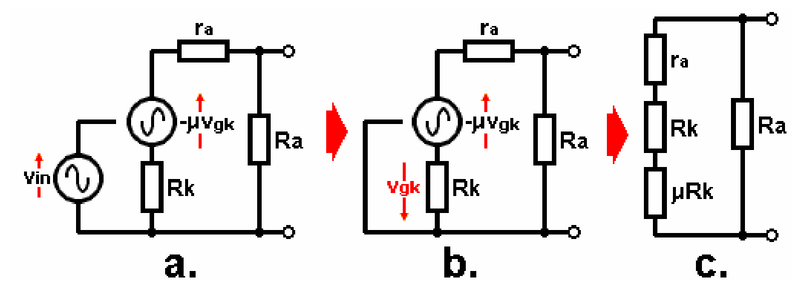

Fig. 1.31: Finding the output impedance of an unbypassed gain stage.

1.22: Input Impedance of a Triode Gain Stage

Provided a valve is biased sufficiently negative that grid current does not flow, the input *resistance* of the grid is practically infinite; it is an open circuit. However, a grid leak resistance of some form is always required, so this sets the input resistance of the stage as a whole, as shown in fig. 1.32. This is commonly 1MΩ, which is large enough that it should not cause undue attenuation with ordinary source impedances.

The input *capacitance* of a triode is another matter, and one that is not small enough to ignore this time. Any two conductors separated by an insulator form a capacitor, so there are unavoidable capacitances inside the valve, called the **interelectrode capacitances**. The important ones are between grid and anode (Cga) and between grid and cathode (Cgk), which are represented by the dashed lines in fig. 1.32a.* Input current will flow in these capacitances just as if they were actual components connected between the grid and ground, which is how they appear in the equivalent circuit in **b**. However, because the anode voltage signal is an amplified version of the signal on the grid, the voltage across Cga will be much larger than if the anode voltage were kept mostly constant, so the resulting input current must also be larger. For example, if the gain of the stage were –60 and we input a +1V signal, the anode voltage would fall by 60V, so the total signal voltage appearing across Cga would be 61V. As far as the input is concerned it looks like we have an additional capacitive load that is 60 times larger than the Cga we already have. This multiplication of the input capacitance is called the **Miller effect**† and this additional imaginary capacitor is called the **Miller capacitance**. The total input capacitance seen between the input terminals is therefore:

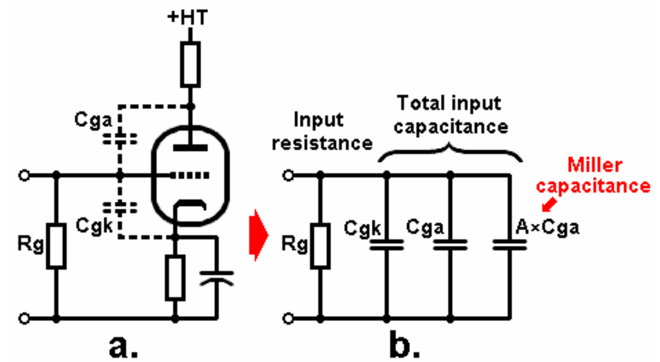

Fig. 1.32: The input capacitance of a triode gain stage is dominated by the Miller effect.

$$Cin = Cgk + Cga(A + 1) \tag{1.8}$$

Where A is the magnitude of the voltage gain.

* Some datasheets may not provide Cgk but will instead give the "grid to all except anode" capacitance, which can be used instead.

† The general term for electronic impedance multiplication or division is **bootstrapping**.

For example, the ECC83 data sheet quotes 1.6pF for both Cga and Cgk, so if the voltage gain is 60 (ignoring any minus sign):

$$\text{Cin} = 1.6 + 1.6 \times (60 + 1) = 99.2\text{pF}$$

It is usual to add some extra to this figure to allow for additional stray capacitances when the circuit is physically built –particularly those added by the valve socket itself– so we might as well round up to 100pF. This input capacitance will form an RC low-pass filter with the source impedance, thereby cutting high frequencies. This can be a hindrance or a help, as will be seen throughout this book.

1.23: Valve Ratings and the Safe Operating Area

Each valve type has certain limitations with regard to applied voltages, current capability and power dissipation, and it is important to be aware of these limits when designing a circuit. By plotting the limits on the static anode characteristics we can see clearly which areas are out-of-bounds and which is safe to work in; this is called the **safe operating area**, or **SOA**.

First there is the maximum voltage which the anode can withstand before it is likely to arc to other electrodes, causing serious failure. This value is given on the datasheet as $\mathbf{Va_0}$, and for the ECC83 it is 550V. The anode voltage must *never* exceed this value, so this is usually taken as the maximum allowable supply voltage. This is indicated in fig. 1.33 by shading everything to the right of Va = 550V in black.

Next there is a maximum allowable *quiescent* anode voltage, given on the data sheet as $\mathbf{Va_{(max)}}$. For the ECC83 it is 350V. The bias point must not lie to the right of this value, so this area is shaded grey in fig. 1.33. However, while amplifying a signal the *instantaneous* anode voltage is allowed to swing above $Va_{(max)}$ provided it does not exceed Va_0.

The upper limit of the SOA is determined by the maximum power that can be dissipated by the anode,* given on the data sheet as $\mathbf{Pa_{(max)}}$ or $\mathbf{Wa_{(max)}}$. Dissipating too much power will cause the anode to glow red hot, which is referred to as **red plating**. If left unchecked the heat may cause the internal parts to soften and deform out of their normal positions, and in extreme cases it can even melt the anode or glass envelope! For the ECC83 the dissipation limit is 1W, and this can be plotted by taking some example anode voltages and then calculating what maximum anode current is allowed to flow, using $I = P / V$:

1W / 400V = 2.5mA
1W / 350V = 2.9mA
1W / 300V = 3.3mA
1W / 250V = 4.0mA
1W / 200V = 5.0mA

---

* Some data sheets may quote a maximum allowable anode current which will put an additional upper limit on the SOA, but this is uncommon for audio valves.

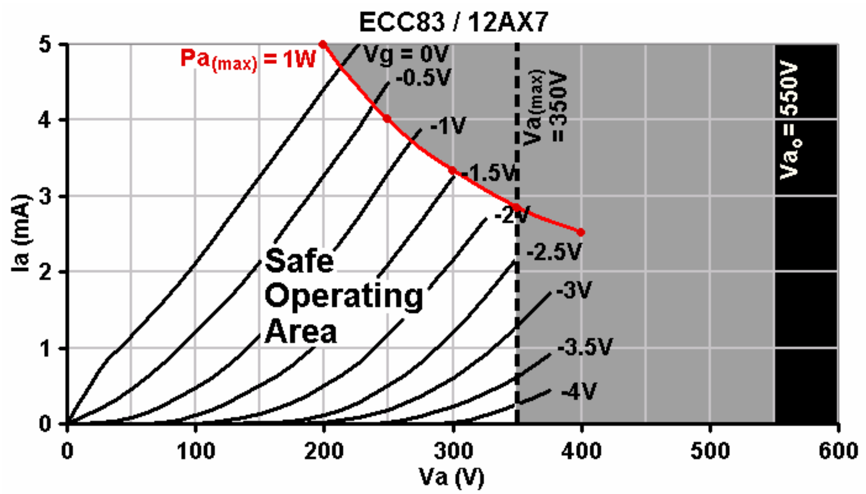

Fig. 1.33: Plotting a maximum anode dissipation curve and safe operating area. The bias point must not lie within the shaded areas.

A curve has been drawn through these points in fig. 1.33. Anywhere above this curve the anode dissipation will exceed 1W, so it too is shaded. The unshaded area is now the SOA, and the bias point must lie somewhere in here. In a preamp design it is fairly easy to stay within the SOA, and in most cases the supply voltage will be less than any of the voltage ratings of the valve, allowing us to forget about these limits. Power valves are often operated much closer to their maximum ratings so special attention must be paid to their data sheets.

There are some additional voltage limits that cannot be plotted on the anode characteristics. One of these is the maximum heater-to-cathode voltage, given as **$Vhk_{(max)}$**. If the potential difference between the heater and cathode becomes too great then leakage current will flow between the two, and if the heater is running on AC the leakage current will lead to a hum voltage on the cathode. Even if the heater is running on DC, sporadic leakage may still lead to popping or frying-bacon noises. $Vhk_{(max)}$ is not so much an absolute limit as an advised limit, and exceeding it may not cause immediate problems. Instead, leakage will increase over time until the valve becomes so noisy that it must be replaced prematurely. The recommended maximum limit for most preamp valves is ±90V between heater and cathode, although the ECC83 / 12AX7 optimistically quotes ±180V. In most cases the cathode is at a low voltage and the heater voltage is likewise, but in some special circuits like the cathode follower or cascode the cathode may rest at a high voltage. In such cases it may be necessary to elevate or ‘float’ the heater supply on a voltage that is close to the cathode voltage (see section 3.17).

Intuition suggests that there must also be a limit to the voltage we can apply between grid and cathode before it may arc. This limit is almost never quoted on data sheets, but having destroyed several devices the author suggests ±100V as a sensible limit for preamp valves. Such voltages are entirely possible in a preamp, and some of the circuits described later do benefit from grid-cathode voltage protection (see section 4.22).

As well as voltage ratings there some limitations on allowable circuit resistances, the most important of which is the maximum grid-leak resistance, $\mathbf{Rgk_{(max)}}$. During operation the grid becomes hot since it is close to the cathode, so it does in fact emit a few electrons of its own. The electrons leaving the grid constitute a tiny **reverse grid current** and, if the grid leak is very large, a positive voltage will therefore develop at the grid. Since this will bias the valve hotter the grid will heat up further and reverse grid current will increase, exacerbating the effect. In extreme cases the valve may go into thermal runaway and start red plating. A cathode-biased circuit is more resistant to this than a fixed-biased one, since any increase in anode current will cause an increase in the voltage across the bias resistor, which counteracts the increase in current. Data sheets will often quote two maximum allowable values of grid leak, one for cathode bias and a smaller one for fixed bias. For the ECC83 the maximum value is usually given as 2MΩ for cathode bias, although values up to 10MΩ can be used if the anode current is less than 1mA.

Many data sheets also quote a maximum allowable resistance between cathode and heater, which often takes the form of a bias resistance. This limit is determined by leakage current between heater and cathode; the larger the resistance between the two, the larger the noise voltage that will develop across it due to leakage current. For the ECC83 $\mathbf{Rhk_{(max)}}$ is given as 150kΩ. Again, this is more of an advised limit rather than an absolute one, and exceeding it does not usually cause problems provided $Vhk_{(max)}$ is not also exceeded.