CHAPTER FIVE

Induction Motors – Rotating Field, Slip and Torque

1. INTRODUCTION

Judged in terms of fitness for purpose coupled with simplicity, the induction motor must rank alongside the screwthread as one of mankind's best inventions. It is not only supremely elegant as an electromechanical energy converter, but is also by far the most important, with something like half of all the electricity generated being converted back to mechanical energy in induction motors. Despite playing a key role in industrial society, it remains largely unnoticed because of its workaday role driving machinery, pumps, fans, compressors, conveyors, hoists, and a host of other routine but vital tasks. It will doubtless continue to dominate fixed-speed applications, but, thanks to the availability of reliable variable-frequency inverters, it is now also the leader in the controlled-speed arena.

Like the d.c. motor, the induction motor develops torque by the interaction of axial currents on the rotor and a radial magnetic field produced by the stator. But whereas in the d.c. motor the 'work' current has to be fed into the rotor by means of brushes and a commutator, the torque-producing currents in the rotor of the induction motor are induced by electromagnetic action, hence the name 'induction' motor. The stator winding therefore not only produces the magnetic field (the 'excitation'), but also supplies the energy which is converted to mechanical output. The absence of any sliding mechanical contacts and the consequent saving in terms of maintenance is a major advantage of the induction motor over the d.c. machine.

Other differences between the induction motor and the d.c. motor are first that the supply to the induction motor is a.c. (usually 3-phase, but in smaller sizes single-phase); secondly that the magnetic field in the induction motor rotates relative to the stator, while in the d.c. motor it is stationary; and thirdly that both stator and rotor in the induction motor are non-salient (i.e. effectively smooth) whereas the d.c. motor stator has projecting poles or saliencies which define the position of the field windings.

Given these differences we might expect to find major contrasts between the performance of the two types of motor, and it is true that their inherent characteristics exhibit distinctive features. But there are also many aspects of behavior which are similar, as we shall see. Perhaps most important from the user's point of view is that there is no dramatic difference in size or weight between an induction motor and a d.c. motor giving the same power at the same base speed, though the induction

© 2013 Austin Hughes and William Drury.
Published by Elsevier Ltd.
All rights reserved.
Electric Motors and Drives
http://dx.doi.org/10.1016/B978-0-08-098332-5.00005-X

motor will usually be cheaper. The similarity in size is a reflection of the fact that both types employ similar amounts of copper and iron, while the difference in price stems from the simpler construction and production volume of the induction motor.

1.1 Outline of approach

Throughout this chapter we will be concerned with how the induction motor behaves in the steady state, i.e. the supply voltage and frequency are constant, the load is steady, and any transients have died away. We will aim to develop a sound qualitative understanding of the steady-state behavior, based on the ideas we have discussed so far (magnetic flux, m.m.f., reluctance, electromagnetic force, motional e.m.f.). But despite many similarities with the d.c. motor, most readers will probably find that the induction motor is more difficult to understand. This is because we are now dealing with alternating rather than steady quantities (so, for example, inductive reactance becomes very significant), and also because (as mentioned earlier) a single winding acts simultaneously as the producer of the flux and the supplier of the converted energy.

In the next chapter, we will extend our qualitative understanding to look at how motor performance depends on design parameters: we will again be following an approach that has served well since the early days of the induction motor, and was developed to reflect the fact that motors were operated at a fixed voltage and frequency. It turns out that under these 'utility supply' conditions, the transient performance is poor and fast control of torque is not possible, and so the induction motor was considered unable to compete with the d.c. motor in controlled-speed drives.

All this changed rapidly beginning in the 1970s. The full set of governing equations (describing not only the steady state but also the much more complex dynamic behavior) had become tractable with computer simulation, which in turn led the way to understanding how the stator currents would have to be manipulated to obtain fast control of torque. The hardware for implementing rapid current control became available with the development of pulse-width modulation (PWM) inverters, but it was not until digital signal processing finally became cheap and fast enough to deal with the complex control algorithms that so-called 'field-oriented' or 'vector' control emerged as a practicable commercial proposition. We will defer consideration of this spectacularly successful system until later, because experience has shown that a solid grounding based on the classical approach is invaluable before getting to grips with more demanding ideas, which are introduced in Chapter 7.

2. THE ROTATING MAGNETIC FIELD

To understand how an induction motor operates, we must first unravel the mysteries of the rotating magnetic field. We will see later that the rotor is effectively dragged along by the rotating field, but that it can never run quite as fast as the field.

Our look at the mechanism of the rotating field will focus on the stator windings because they act as the source of the flux. In this part of the discussion we will ignore the presence of the rotor conductors. This makes it much easier to understand what governs the speed of rotation and the magnitude of the field, which are the two factors that most influence motor behavior.

Having established how the rotating field is set up, and what its speed and strength depend on, we move on to examine the rotor, concentrating on how it behaves when exposed to the rotating field, and discovering how the induced rotor currents and torque vary with rotor speed. In this section we assume – again for the sake of simplicity – that the rotating flux set up by the stator is not influenced by the rotor.

Finally we turn attention to the interaction between the rotor and stator, verifying that our earlier assumptions are well justified. Having done this we are in a position to examine the 'external characteristics' of the motor, i.e. the variation of motor torque and stator current with speed. These are the most important characteristics from the point of view of the user.

Readers who are unfamiliar with routine a.c. circuit theory, including reactance, impedance, phasor diagrams (but not, at this stage, 'j' notation) and basic ideas about 3-phase systems will have to do some preparatory work[1] before tackling the later sections of this chapter.

Before we investigate how the rotating magnetic field is produced, we should be clear what it actually is. Because both the rotor and stator iron surfaces are smooth (apart from the regular slotting), and are separated by a small air-gap, the flux produced by the stator windings crosses the air-gap radially. The behavior of the motor is dictated by this radial flux, so we will concentrate first on establishing a mental picture of what is meant by the 'flux wave' in an induction motor.

The pattern of flux in an ideal 4-pole motor supplied from a balanced 3-phase source is shown in Figure 5.1(a). The top sketch corresponds to time $t = 0$; the middle one shows the flux pattern one-quarter of a cycle of the supply later (i.e. 5 ms if the frequency is 50 Hz); and the lower one corresponds to a further quarter-cycle later. We note that the pattern of flux lines is repeated in each case, except that the middle and lower ones are rotated by 45° and 90°, respectively, with respect to the top sketch.

The term '4-pole' reflects the fact that flux leaves the stator from two N poles, and returns at two S poles. Note, however, that there are no physical features of the stator iron that mark it out as being 4-pole, rather than say 2-pole or 6-pole. As we will see, it is the layout and interconnection of the stator coils which set the pole-number.

If we plot the variation of the radial air-gap flux density with respect to distance round the stator, at each of the three instants of time, we get the patterns shown in

1 The revised book *Electrical and Electronic Technology*, 10th Edition by Edward Hughes (no relation) is a tried and tested favorite.

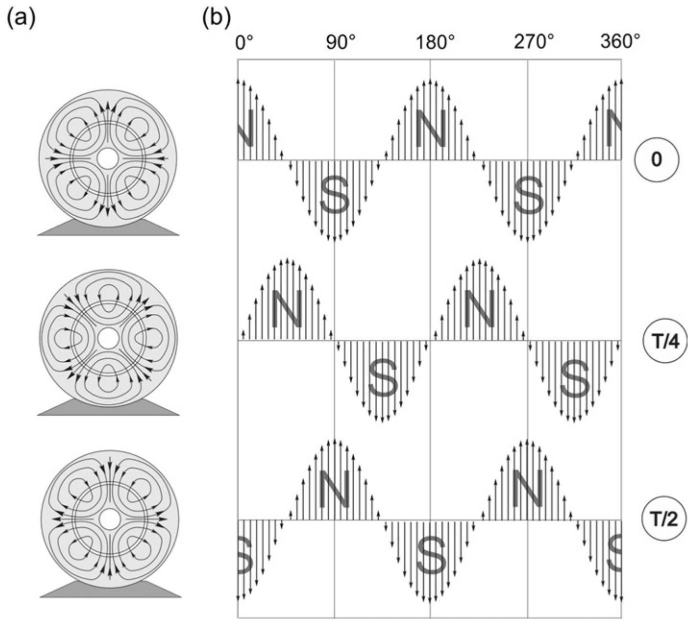

Figure 5.1 (a) Flux pattern in a 4-pole induction motor at three successive instants of time, each one-quarter of a cycle apart; (b) radial flux density distribution in the air-gap at the three instants shown in Figure 5.1(a).

Figure 5.1(b). The first feature to note is that the radial flux density varies sinusoidally in space. There are two N peaks and two S peaks, but the transition from N to S occurs in a smooth sinusoidal way, giving rise to the term 'flux wave'. The distance from the center of one N pole to the center of the adjacent S pole is called the pole-pitch, for obvious reasons.

Staying with Figure 5.1(b), we note that after one-quarter of a cycle of the mains frequency, the flux wave retains its original shape, but has moved round the stator by half a pole-pitch, while after half a cycle it has moved round by a full pole-pitch. If we had plotted the patterns at intermediate times, we would have discovered that the wave maintained a constant shape, and progressed smoothly, advancing at a uniform rate of two pole-pitches per cycle of the supply. The term 'traveling flux wave' is thus an appropriate one to describe the air-gap field.

For the 4-pole wave here, one complete revolution takes two cycles of the supply, so the speed is 25 rev/s (1500 rev/min) with a 50 Hz supply, or 30 rev/s (1800 rev/min) at 60 Hz. The general expression for the speed of the field (which is known as the synchronous speed) N_s, in rev/min is

$$N_s = \frac{120f}{p} \tag{5.1}$$

Table 5.1 Synchronous speeds, in rev/min

Pole-number	50 Hz	60 Hz
2	3000	3600
4	1500	1800
6	1000	1200
8	750	900
10	600	720
12	500	600

where p is the pole-number. The pole-number must be an even integer, since for every N pole there must be an S pole. Synchronous speeds for commonly used pole-numbers are given in Table 5.1.

We can see from the table that if we want the field to rotate at intermediate speeds, we will have to be able to vary the supply frequency, and this is what happens in inverter-fed motors, which are dealt with in Chapter 7.

2.1 Production of rotating magnetic field

Now that we have a picture of the field, we turn to how it is produced. If we inspect the stator winding of an induction motor we find that it consists of a uniform array of identical coils, located in slots. The coils are in fact connected to form three identical groups or phase-windings, distributed around the stator, and symmetrically displaced with respect to one another. The three phase-windings are connected either in star (wye) or delta (mesh), as shown in Figure 5.2.

The three phase-windings are connected to a balanced 3-phase a.c. supply, and so the currents (which produce the m.m.f. that sets up the flux) are of equal amplitude but differ in time-phase by one-third of a cycle (120°), forming a balanced 3-phase set.

2.2 Field produced by each phase-winding

The aim of the winding designer is to arrange the layout of the coils so that each phase-winding, acting alone, produces an m.m.f. wave (and hence an air-gap flux

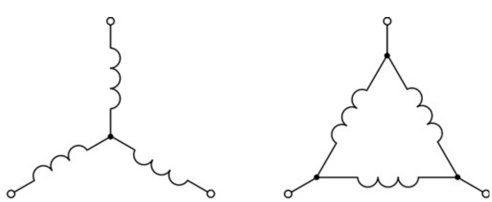

Figure 5.2 Star (wye) and delta (mesh) connection of the three phase-windings of a 3-phase induction motor.

wave) of the desired pole-number, and with a sinusoidal variation of amplitude with angle. Getting the desired pole-number is not difficult: we simply have to choose the right number and pitch of coils, as shown by the diagrams of an elementary 4-pole winding in Figure 5.3.

In Figure 5.3(a) we see that by positioning two coils (each of which spans one pole-pitch) 180° apart we obtain the correct number of poles (i.e. 4). However, the air gap field – shown by only two flux lines per pole for the sake of clarity – is uniform between each go and return coil-side, not sinusoidal.

A clearer picture of the air-gap flux wave is presented in the developed view in Figure 5.3(b), where more equally spaced flux lines have been added to emphasize the uniformity of the flux density between the 'go' and 'return' sides of the coils. Finally, the plot of the air-gap flux density underlines the fact that this very basic arrangement of coils produces a rectangular flux density wave, whereas what we are seeking is a sinusoidal wave.

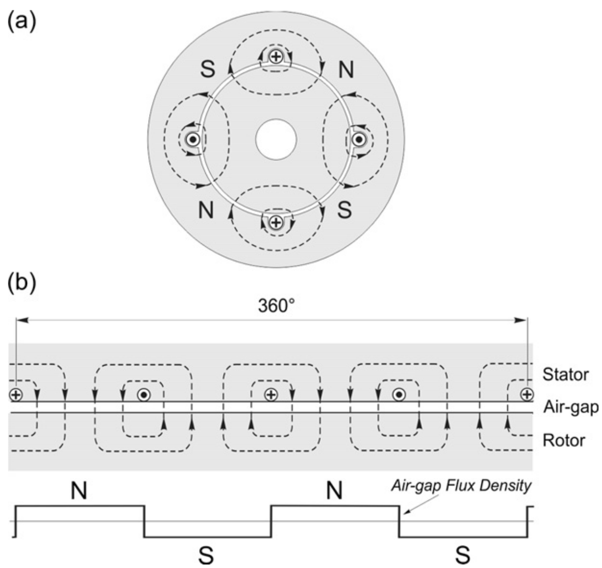

Figure 5.3 Arrangement (a) and developed diagram (b) showing elementary 4-pole, single-layer stator winding consisting of four conductors spaced by 90°. The 'go' side of each coil (shown by the plus symbol) carries current into the paper at the instant shown, while the 'return' side (shown by the dot) carries current out of the paper.

We can improve matters by adding more coils in the adjacent slots, as shown in Figure 5.4. All the coils have the same number of turns, and carry the same current. The addition of the extra slightly displaced coils gives rise to the stepped waveform of m.m.f. and air-gap flux density shown in Figure 5.4. It is still not sinusoidal, but is much better than the original rectangular shape.

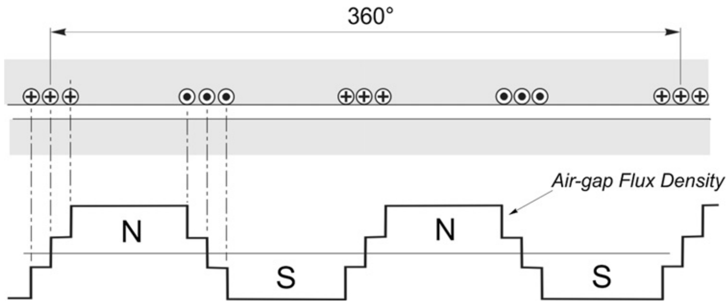

Figure 5.4 Developed diagram showing flux density produced by one phase of a single-layer winding having three slots per pole per phase.

It turns out that if we were to insist on having a perfect sinusoidal flux density waveform, we would have to distribute the coils of one phase in a smoothly varying sinusoidal pattern over the whole periphery of the stator. This is not a practicable proposition, first because we would also have to vary the number of turns per coil from point to point, and secondly because we want the coils to be in slots, so it is impossible to avoid some measure of discretization in the layout. For economy of manufacture we are also obliged to settle for all the coils being identical, and we must make sure that the three identical phase-windings fit together in such a way that all the slots are fully utilized. (See Plate 5.1)

Despite these constraints we can get remarkably close to the ideal sinusoidal pattern, especially when we use a 'two-layer' winding (in which case the stator slots may contain turns from more than one phase winding). A typical arrangement of one phase is shown in Figure 5.5. The upper expanded sketch shows how each coil sits with its 'go' side in the top of a slot while the 'return' side occupies the bottom of a slot rather less than one pole-pitch away. Coils which span less than a full pole-pitch are known as short-pitch or short-chorded: in this particular case the coil pitch is six slots and the pole-pitch is nine slots, so the coils are short-pitched by three slots.

This type of winding is almost universal in all but small induction motors, the coils in each phase being grouped together to form 'phase-bands' or 'phase-belts'. Since we are concentrating on the field produced by only one of the phase-windings (or 'phases'), only one-third of the coils in Figure 5.5 are shown carrying current. The remaining two-thirds of the coils form the other two phase-windings, as discussed below.

Returning to the flux density plot in Figure 5.5 we see that the effect of short-pitching is to increase the number of steps in the waveform, and that as a result the field produced by one phase is a fair approximation to a sinusoid.

The current in each phase pulsates at the supply frequency, so the field produced by say phase A, pulsates in sympathy with the current in phase A, the axis of each 'pole' remaining fixed in space, but its polarity changing from N to S and back once

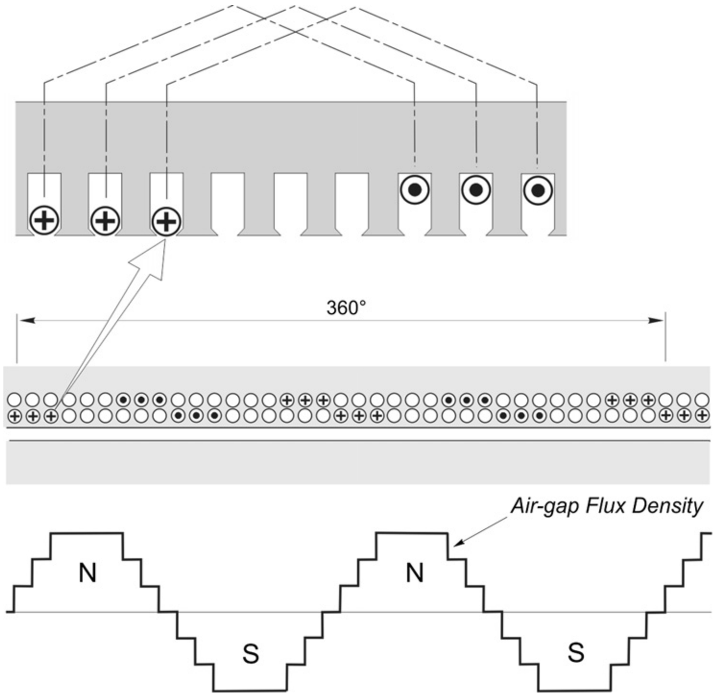

Figure 5.5 Developed diagram showing layout of windings in a 3-phase, 4-pole, two-layer induction motor winding, together with the flux density wave produced by one phase acting alone. The upper detail shows how the coil-sides form upper and lower layers in the slots.

per cycle. There is no hint of any rotation in the field of one phase, but when the fields produced by each of the three phases are combined, matters change dramatically.

2.3 Resultant 3-phase field

The layout of coils for the complete 4-pole winding is shown in Figure 5.6(a). The 'go' sides of each coil are represented by the capital letters (A, B, C) and the 'return' sides are identified by bars over the letters ($\overline{A}, \overline{B}, \overline{C}$). (For the sake of comparison, a 6-pole winding layout that uses the same stator slotting is shown in Figure 5.6(b): here the pole-pitch is six slots and the coils are short-pitched by one slot.)

Returning to the 4-pole winding, we can see that the windings of phases B and C are identical with that of phase A apart from the fact that they are displaced in space by plus and minus two-thirds of a pole-pitch, respectively. Phases B and C therefore also produce pulsating fields, along their own fixed axes in space. But the currents in phases B and C also differ in time-phase from the current in phase A, lagging by one-third and two-thirds of a cycle, respectively. To find the resultant

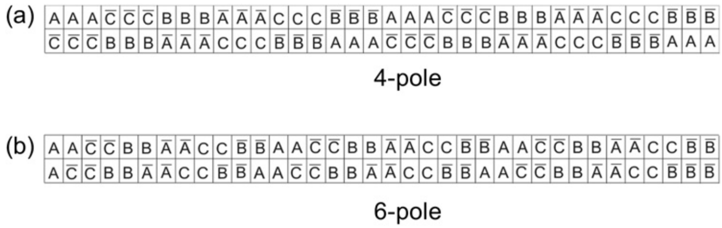

Figure 5.6 Developed diagram showing arrangement of 3-phase, two-layer windings in a 36-slot stator. A 4-pole winding with three slots/pole/phase is shown in (a), and a 6-pole winding with two slots/pole/phase is shown in (b).

field we must therefore superimpose the fields of the three phases, taking account not only of the spatial differences between windings, but also the time differences between the currents. This is a tedious process, so the intermediate steps have been omitted and instead we move straight to the plot of the resultant field for the complete 4-pole machine, for three discrete times during one complete cycle, as shown in Figure 5.7.

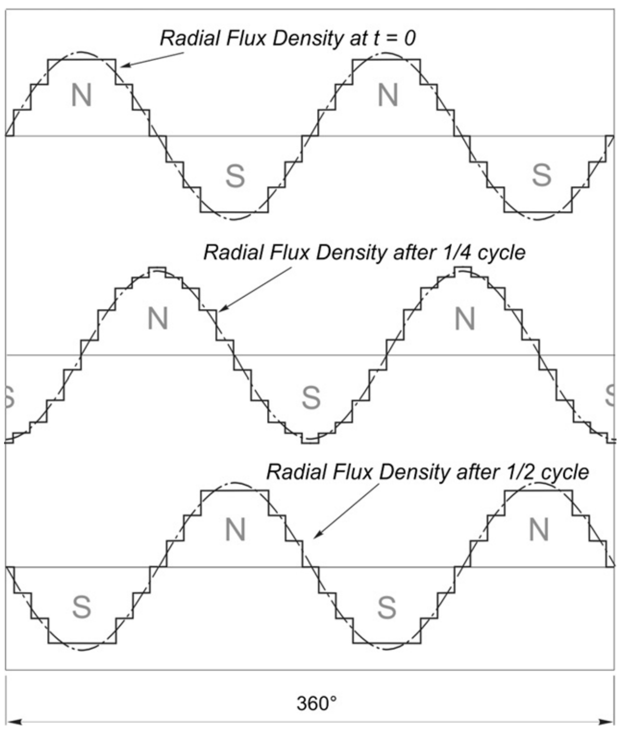

Figure 5.7 Resultant air–gap flux density wave produced by a complete 3-phase, 4-pole winding at three successive instants in time.

We see that the three pulsating fields combine beautifully and lead to a resultant 4-pole field which rotates at a uniform rate, advancing by two pole-pitches for every cycle of the supply. The resultant field is not exactly sinusoidal in shape (though it is actually more sinusoidal than the field produced by the individual phase-windings), and its shape varies a little from instant to instant; but these are minor worries. The resultant field is amazingly close to the ideal traveling wave and yet the winding layout is simple and easy to manufacture. This is an elegant engineering achievement, however one looks at it.

2.4 Direction of rotation

The direction of rotation depends on the order in which the currents reach their maxima, i.e. on the phase-sequence of the supply. Reversal of direction is therefore simply a matter of interchanging any two of the lines connecting the windings to the supply.

2.5 Main (air-gap) flux and leakage flux

Broadly speaking the motor designer shapes the stator and rotor teeth to encourage as much as possible of the flux produced by the stator windings to pass right down the rotor teeth, so that before completing its path back to the stator it is fully linked with the rotor conductors (see later) which are located in the rotor slots. We will see later that this tight magnetic coupling between stator and rotor windings is necessary for good running performance, and the field which provides the coupling is of course the main or air-gap field, which we are in the midst of discussing.

In practice the vast majority of the flux produced by the stator is indeed main or 'mutual' flux. But there is some flux which bypasses the rotor conductors, linking only with the stator winding, and known as stator leakage flux. Similarly not all the flux produced by the rotor currents links the stator, but some (the rotor leakage flux) links only the rotor conductors.

The use of the pejorative-sounding term 'leakage' suggests that these leakage fluxes are unwelcome imperfections, which we should go out of our way to minimize. However, while the majority of aspects of performance are certainly enhanced if the leakage is as small as possible, others (notably the large and unwelcome current drawn from the mains when the motor is started from rest directly on the utility supply) are made much worse if the coupling is too good. So we have the somewhat paradoxical situation in which the designer finds it comparatively easy to lay out the windings to produce a good main flux, but is then obliged to juggle the detailed design of the slots in order to obtain just the right amount of leakage flux to give acceptable all-round performance. (In contrast, as we will see later, an inverter-fed induction motor can avoid such issues as excessive starting current and, ideally, could be designed with much lower leakage than its utility-fed counterpart. It has to be said, however, that the majority of induction motors are still designed for general-purpose use, and in this respect they lose out in comparison with other forms of motor that are specifically designed for operation with a drive.)

The weight which attaches to the matter of leakage flux is reflected in the prominent part played by the associated leakage reactance in equivalent circuit models of the induction motor (see Appendix 2). However, such niceties are of limited importance to the user, so in this and the next chapters we will limit references to leakage reactance to well-defined contexts, and, in general, where the term 'flux' is used, it will refer to the main air-gap field.

2.6 Magnitude of rotating flux wave

We have already seen that the speed of the flux wave is set by the pole-number of the winding and the frequency of the supply. But what is it that determines the amplitude of the field?

To answer this question we can continue to neglect the fact that under normal conditions there will be induced currents in the rotor. We might even find it easier to imagine that the rotor conductors have been removed altogether: this may seem a drastic assumption, but will prove justified later. The stator windings are assumed to be connected to a balanced 3-phase a.c. supply so that a balanced set of currents flow in the windings. We denote the phase voltage by V, and the current in each phase by I_m, where the subscript m denotes 'magnetizing' or flux-producing current.

From the discussion in Chapter 1 we know that the magnitude of the flux wave (B_m) is proportional to the winding m.m.f., and is thus proportional to I_m. But what we really want to know is how the flux density depends on the supply voltage and frequency, since these are the only two parameters over which we have control.

To guide us to the answer, we must first ask what effect the traveling flux wave will have on the stator winding. Every stator conductor will of course be cut by the rotating flux wave, and will therefore have an e.m.f. induced in it. Since the flux wave varies sinusoidally in space, and cuts each conductor at a constant velocity, a sinusoidal e.m.f. is induced in each conductor. The magnitude of the e.m.f. is proportional to the magnitude of the flux wave (B_m), and to the speed of the wave (i.e. to the supply frequency f). The frequency of the induced e.m.f. depends on the time taken for one N pole and one S pole to cut the conductor. We have already seen that the higher the pole-number, the slower the field rotates, but we found that the field always advances by two pole-pitches for every cycle of the supply. The frequency of the e.m.f. induced in the stator conductors is therefore the same as the supply frequency, regardless of the pole-number. (This conclusion is what we would have reached intuitively, since we would expect any linear system to react at the same frequency at which we excited it.)

The e.m.f. in each complete phase winding (E) is the sum of the e.m.f.s in the phase coils, and thus will also be at supply frequency. (The alert reader will realize that while the e.m.f. in each coil has the same magnitude, it will differ in time phase, depending on the geometrical position of the coil. Most of the coils in each phase-band are close together, however, so their e.m.f.s – though slightly out of phase – will more or less add up directly.)

If we were to compare the e.m.f.s in the three complete phase windings, we would find that they were of equal amplitude, but out of phase by one-third of a cycle (120°), thereby forming a balanced 3-phase set. This result could have been anticipated from the overall symmetry. This is very helpful, as it means that we need only consider one of the phases in the rest of the discussion.

So we find that when an alternating voltage V is applied, an alternating e.m.f. E is induced. We can represent this state of affairs by the primitive a.c. equivalent circuit for one phase shown in Figure 5.8.

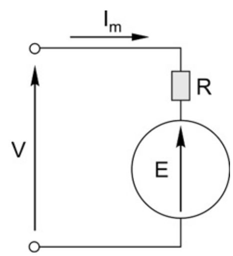

Figure 5.8 Simple equivalent circuit for the induction motor under no-load conditions.

The resistance shown in Figure 5.8 is the resistance of one complete phase-winding. Note that the e.m.f. E is shown as opposing the applied voltage V. This must be so, otherwise we would have a runaway situation in which the voltage V produced the magnetizing current I_m which in turn set up an e.m.f. E, which added to V, which further increased I_m and so on *ad infinitum*.

Applying Kirchhoff's voltage law to the a.c. circuit in Figure 5.8 yields

$$V = I_m R + E \tag{5.2}$$

We find in practice that the term $I_m R$ (which represents the volt-drop due to winding resistance) is usually very much less than the applied voltage V. In other words most of the applied voltage is accounted for by the opposing e.m.f. E. Hence we can make the approximation

$$V \approx E \tag{5.3}$$

But we have already seen that the e.m.f. is proportional to B_m and to f, i.e.

$$E \propto B_m f \tag{5.4}$$

So by combining equations (5.3) and (5.4) we obtain

$$B_m = k\frac{V}{f} \tag{5.5}$$

where the constant k depends on the number of turns per coil, the number of coils per phase and the distribution of the coils.

Equation (5.5) is of fundamental importance in induction motor operation. It shows that if the supply frequency is constant, the flux in the air-gap is directly proportional to the applied voltage, or in other words the voltage sets the flux. We can also see that if we raise or lower the frequency (in order to increase or reduce the speed of rotation of the field), we will have to raise or lower the voltage in proportion if, as is usually the case, we want the magnitude of the flux to remain constant. (We will see in Chapters 7 and 8 that the early inverter drives used this so-called 'V/f control' to keep the flux constant at all speeds.)

It may seem a paradox that having originally homed-in on the magnetizing current I_m as being the source of the m.m.f. which in turn produces the flux, we find that the actual value of the flux is governed only by the applied voltage and frequency, and I_m does not appear at all in equation (5.5). We can see why this is by looking again at Figure 5.8 and asking what would happen if, for some reason, the e.m.f. (E) were to reduce. We would find that I_m would increase, which in turn would lead to a higher m.m.f., more flux, and hence to an increase in E. There is clearly a negative feedback effect taking place, which continually tries to keep E equal to V. It is rather like the d.c. motor (Chapter 3) where the speed of the unloaded motor always adjusted itself so that the back e.m.f. almost equaled the applied voltage. Here, the magnetizing current always adjusts itself so that the induced e.m.f. is almost equal to the applied voltage.

Needless to say this does not mean that the magnetizing current is arbitrary, but to calculate it we would have to know the number of turns in the winding, the length of the air-gap (from which we could calculate the gap reluctance) and the reluctance of the iron paths. From a user point of view there is no need to delve further in this direction. We should, however, recognize that the reluctance will be dominated by the air-gap, and that the magnitude of the magnetizing current will therefore depend mainly on the size of the gap: the larger the gap, the bigger the magnetizing current. Since the magnetizing current contributes to stator copper loss, but not to useful output power, we would like it to be as small as possible, so we find that induction motors usually have the smallest air-gap which is consistent with providing the necessary mechanical clearances. Despite the small air-gap the magnetizing current can be appreciable: in a 4-pole motor, it may be typically 50% of the full-load current, and even higher in 6-pole and 8-pole designs.

2.7 Excitation power and volt-amps

The setting up of the traveling wave by the magnetizing current amounts to the provision of 'excitation' for the motor. Some energy is stored in the magnetic field, but since the amplitude remains constant once the field has been established, no net power input is needed to sustain the field. We therefore find that under the

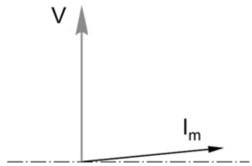

Figure 5.9 Phasor diagram for the induction motor under no-load conditions, showing magnetizing current I_m.

conditions discussed so far, i.e. in the absence of any rotor currents, the power input to the motor is very small. (We should perhaps note that the rotor currents in a real motor are very small when it is running light, so the hypothetical situation we are looking at is not as far removed from reality as we may have supposed.)

Ideally the only source of power losses would be the copper losses in the stator windings, but to this must be added the 'iron losses' which arise from eddy currents and hysteresis in the laminated steel cores of rotor and stator. However, we have seen that the magnetizing current can be quite large, its value being largely determined by the air-gap, so we can expect an unloaded induction motor to draw appreciable current from the supply, but very little real power. The volt-amps will therefore be substantial, but the power-factor will be very low, the magnetizing current lagging the supply voltage by almost 90°, as shown in the time phasor diagram (Figure 5.9).

Viewed from the supply the stator looks more or less like a pure inductance, a fact which we would expect intuitively given that – having ignored the rotor circuit – we are left with only an arrangement of flux-producing coils surrounded by a good magnetic circuit.

2.8 Summary

When the stator is connected to a 3-phase supply, a sinusoidally distributed, radially directed rotating magnetic flux density wave is set up in the air-gap. The speed of rotation of the field is directly proportional to the frequency of the supply, and inversely proportional to the pole-number of the winding. The magnitude of the flux wave is proportional to the applied voltage, and inversely proportional to the frequency.

When the rotor circuits are ignored (i.e. under no-load conditions), the real power drawn is small, but the magnetizing current itself can be quite large, giving rise to a significant reactive power demand from the utility supply.

3. TORQUE PRODUCTION

In this section we begin with a brief description of rotor types, and introduce the notion of 'slip', before moving on to explore how the torque is produced, and investigate the variation of torque with speed. We will find that the behavior of the rotor varies widely according to the slip, and we therefore look separately at low and high values of slip. Throughout this section we will assume that the rotating magnetic field is unaffected by anything which happens on the rotor side of the air-gap. Later, we will see that this assumption is pretty well justified.

3.1 Rotor construction

Two types of rotor are used in induction motors. In both, the rotor 'iron' consists of a stack of silicon steel laminations with evenly spaced slots punched around the circumference. As with the stator laminations, the surface is coated with an oxide layer which acts as an insulator, preventing unwanted axial eddy-currents from flowing in the iron.

The cage rotor is by far the most common: each rotor slot contains a solid conductor bar and all the conductors are physically and electrically joined together at each end of the rotor by conducting 'end-rings' (Figure 5.10, plate 5.2 and see also Figure 8.4). In the larger sizes the conductors will be of copper, in which case the end-rings are brazed on. In small and medium sizes, the rotor conductors and end rings may be of copper or die-cast in aluminum.

The term squirrel cage was widely used at one time and the origin should be clear from Figure 5.10. The rotor bars and end-rings are reminiscent of the rotating cages used in bygone days to exercise small rodents (or rather to amuse their human captors).

The absence of any means for making direct electrical connection to the rotor underlines the fact that in the induction motor the rotor currents are induced by the air-gap field. It is equally clear that because the rotor cage comprises permanently short-circuited conductor bars, no external control can be exercised over the resistance of the rotor circuit once the rotor has been made. This is a significant drawback which can be avoided in the second type of rotor, which is known as the 'wound-rotor' or 'slipring' type.

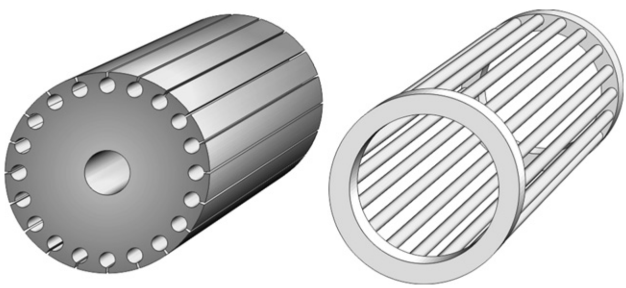

Figure 5.10 Cage rotor construction. The stack of pre-punched laminations is shown on the left, with the copper or aluminum rotor bars and end-rings on the right.

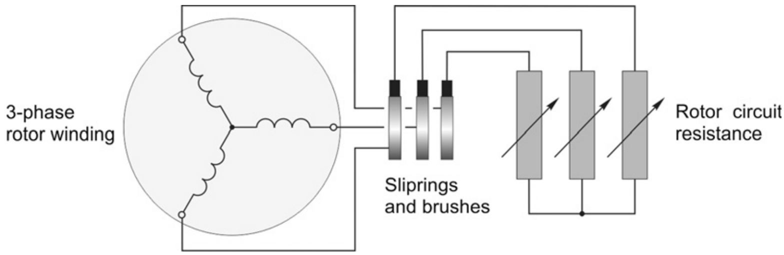

Figure 5.11 Schematic diagram of wound rotor for induction motor, showing sliprings and brushes to provide connection to the external (stationary) 3-phase resistance.

In the wound rotor, the slots accommodate a set of three phase-windings very much like those on the stator. The windings are connected in star, with the three ends brought out to three sliprings (Figure 5.11). The rotor circuit is thus open, and connection can be made via brushes bearing on the sliprings. In particular, the resistance of each phase of the rotor circuit can be increased by adding external resistances, as indicated in Figure 5.11. Adding resistance can be beneficial in some circumstances, as we will see.

Cage-rotors are usually cheaper to manufacture, and are very robust and reliable. Until the advent of variable-frequency inverter supplies, however, the superior control which was possible from the slipring type meant that the extra expense of the wound rotor and its associated control gear was frequently justified, especially for high-power machines. Nowadays comparatively few are made, and then only in large sizes. But many old motors remain in service, so they are included in Chapter 6.

3.2 Slip

A little thought will show that the behavior of the rotor depends very much on its relative velocity with respect to the rotating field. If the rotor is stationary, for example, the rotating field will cut the rotor conductors at synchronous speed, thereby inducing a high e.m.f. in them. On the other hand, if the rotor was running at the synchronous speed, its relative velocity with respect to the field would be zero, and no e.m.f.s would be induced in the rotor conductors.

The relative velocity between the rotor and the field is known as the slip speed. If the speed of the rotor is N, the slip speed is $N_s - N$, where N_s is the synchronous speed of the field, usually expressed in rev/min. The slip (as distinct from slip speed) is the normalized quantity defined by

$$s = \frac{N_s - N}{N_s} \tag{5.6}$$

and is usually expressed either as a ratio as in equation (5.6), or as a percentage. A slip of 0 therefore indicates that the rotor speed is equal to the synchronous speed, while a slip of 1 corresponds to zero speed. (When tests are performed on induction motors with their rotor deliberately held stationary so that the slip is 1, the test is said to be under 'locked-rotor' conditions. The same expression is often used loosely to mean zero speed, even when the rotor is free to move, e.g. when it is started from rest.)

3.3 Rotor-induced e.m.f. and current

The rate at which the rotor conductors are cut by the flux – and hence their induced e.m.f. – is directly proportional to the slip, with no induced e.m.f. at synchronous speed ($s = 0$) and maximum induced e.m.f. when the rotor is stationary ($s = 1$).

The frequency of the rotor e.m.f. (the slip frequency) is also directly proportional to slip, since the rotor effectively slides with respect to the flux wave, and the higher the relative speed, the more times in a second each rotor conductor is cut by an N and an S pole. At synchronous speed (slip = 0) the slip frequency is zero, while at standstill (slip = 1), the slip frequency is equal to the supply frequency. These relationships are shown in Figure 5.12.

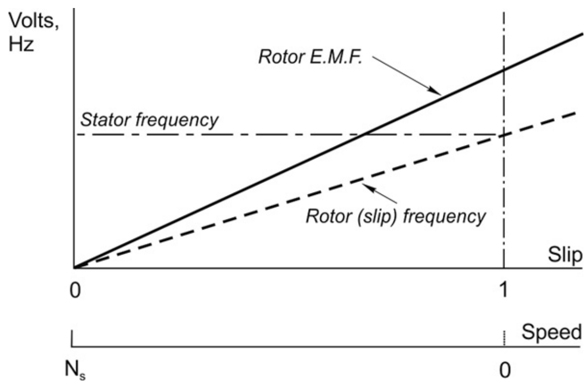

Figure 5.12 Variation of rotor-induced e.m.f. and frequency with speed and slip.

Although the e.m.f. induced in every rotor bar will have the same magnitude and frequency, they will not be in phase. At any particular instant, bars under the peak of the N poles of the field will have maximum positive voltage in them, those under the peak of the S poles will have maximum negative voltage (i.e. 180° phase shift), and those in between will have varying degrees of phase shift. The pattern of instantaneous voltages in the rotor is thus a replica of the flux density wave, and the rotor-induced 'voltage wave' therefore moves relative to the rotor at slip speed, as shown in Figure 5.13.

All the rotor bars are short-circuited by the end-rings, so the induced voltages will drive currents along the rotor bars, the currents forming closed paths through the end-rings, as shown in the developed diagram (Figure 5.14).

In Figure 5.14 the variation of instantaneous e.m.f. in the rotor bars is shown in the upper sketch, while the corresponding instantaneous currents flowing in the

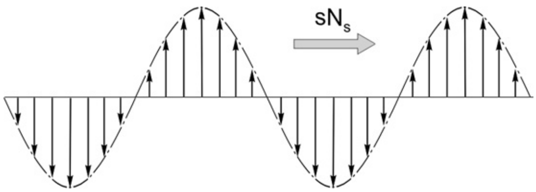

Figure 5.13 Pattern of induced e.m.f.s in rotor conductors. The rotor 'voltage wave' moves at a speed of sN_s with respect to the rotor surface.

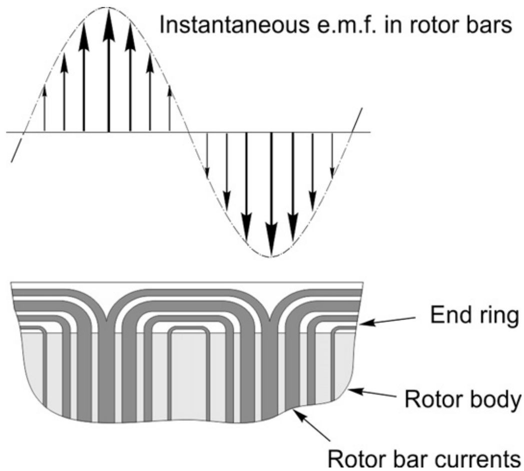

Figure 5.14 Instantaneous sinusoidal pattern of rotor currents in rotor bars and end-rings. Only one pole-pitch is shown, but the pattern is repeated.

rotor bars and end-rings are shown in the lower sketch. The lines representing the currents in the rotor bars have been drawn so that their width is proportional to the instantaneous currents in the bars.

3.4 Torque

The axial currents in the rotor bars will interact with the radial flux wave to produce the driving torque of the motor, which will act in the same direction as the rotating field, the rotor being dragged along by the field. We note that slip is essential to this mechanism, so that it is never possible for the rotor to catch up with the field, as there would then be no rotor e.m.f., no current, and no torque. The fact that motor action is only possible if the speed is less than the synchronous speed explains why the induction machine is described as 'asynchronous'. Finally, we can see that the cage rotor will automatically adapt to whatever pole-number is impressed by the stator winding, so that the same rotor can be used for a range of different stator pole-numbers.

3.5 Rotor currents and torque – small slip

When the slip is small (say between 0 and 10%), the frequency of induced e.m.f. is also very low (between 0 and 5 Hz if the supply frequency is 50 Hz). At these low frequencies the impedance of the rotor circuits is predominantly resistive, the inductive reactance being small because the rotor frequency is low.

The current in each rotor conductor is therefore in time-phase with the e.m.f. in that conductor, and the rotor current-wave is therefore in space-phase with the rotor e.m.f. wave, which in turn is in space-phase with the flux wave. This situation was assumed in the previous discussion, and is represented by the space waveforms shown in Figure 5.15.

To calculate the torque we first need to evaluate the '$BI_r l_r$' product (see equation (1.2)) in order to obtain the tangential force on each rotor conductor. The torque is then given by the total force multiplied by the rotor radius. We can see from Figure 5.15 that where the flux density has a positive peak, so does the rotor current, so that particular bar will contribute a high tangential force to the total torque. Similarly, where the flux has its maximum negative peak, the induced current is maximum and negative, so the tangential force is again positive. We don't need to work out the torque in detail, but it should be clear that the resultant will be given by an equation of the form

$$T = kBI_r \tag{5.7}$$

where B and I_r denote the amplitudes of the flux density wave and the rotor current wave, respectively. Provided that there are a large number of rotor bars (which is a safe bet in practice), the waves shown in Figure 5.15 will remain the same at all instants of time, so the torque remains constant as the rotor rotates.

If the supply voltage and frequency are constant, the flux will be constant (see equation (5.5)). The rotor e.m.f. (and hence I_r) is then proportional to slip, so we can see from equation (5.7) that the torque is directly proportional to slip. We must remember that this discussion relates to low values of slip only, but since this is the normal running condition, it is extremely important.

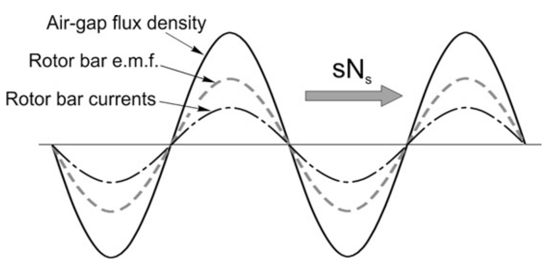

Figure 5.15 Pattern of air-gap flux density, induced e.m.f. and current in cage rotor bars at low values of slip.

The torque–speed (and torque–slip) relationship for small slips is thus approximately a straight line, as shown by the section of line AB in Figure 5.16.

If the motor is unloaded, it will need very little torque to keep running – only enough to overcome friction in fact – so an unloaded motor will run with a very small slip at just below the synchronous speed, as shown at A in Figure 5.16.

When the load is increased, the rotor slows down, and the slip increases, thereby inducing more rotor e.m.f. and current, and thus more torque. The speed will settle when the slip has increased to the point where the developed torque equals the load torque – e.g. point B in Figure 5.16.

Induction motors are usually designed so that their full-load torque is developed for small values of slip. Small ones typically have a full-load slip of 8%, large ones around 1%. At the full-load slip, the rotor conductors will be carrying their safe maximum continuous current, and if the slip is any higher, the rotor will begin to overheat. This overload region is shown by the dotted line in Figure 5.16.

The torque–slip (or torque–speed) characteristic shown in Figure 5.16 is a good one for most applications, because the speed only falls a little when the load is raised from zero to its full value. We note that, in this normal operating region, the torque–speed curve is very similar to that of a d.c. motor (see Figure 3.9).

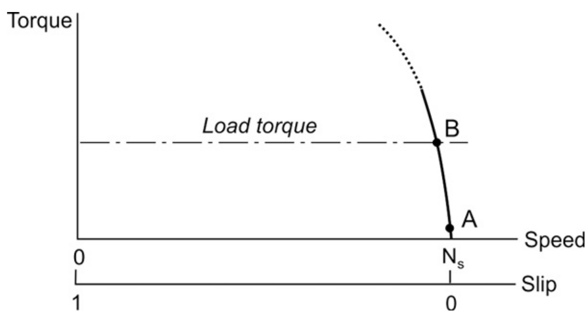

Figure 5.16 Torque–speed relationship for low values of slip.

3.6 Rotor currents and torque – large slip

As the slip increases, the rotor e.m.f. and rotor frequency both increase in direct proportion to the slip. At the same time the rotor inductive reactance, which was negligible at low slip (low rotor frequency), begins to be appreciable in comparison with the rotor resistance. Hence although the induced current continues to increase with slip, it does so more slowly than at low values of slip, as shown in Figure 5.17.

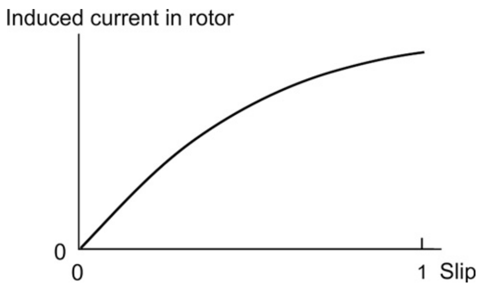

Figure 5.17 Magnitude of current induced in rotor over the full (motoring) range of slip.

At high values of slip, the rotor current also lags behind the rotor e.m.f. because of the inductive reactance. The alternating current in each bar reaches its peak well after the induced voltage, and this in turn means that the rotor current wave has a space-lag with respect to the rotor e.m.f. wave (which is in space-phase with the flux wave). This space-lag is shown by the angle ϕ_r in Figure 5.18.

The space-lag means that the peak radial flux density and peak rotor currents no longer coincide, which is bad news from the point of view of torque production, because although we have high values of both flux density and current, they do not occur simultaneously at any point around the periphery. What is worse is that at some points we even have flux density and currents of opposite sign, so over those regions of the rotor surface the torque contributed will actually be negative. The overall torque will still be positive, but is much less than it would be if the flux and current waves were in phase. We can allow for the unwelcome space-lag by modifying equation (5.7), to obtain a more general expression for torque as

$$T = kBI_r \cos \phi_r \tag{5.8}$$

Equation (5.7) is merely a special case of equation (5.8), which only applies under low-slip conditions where $\cos\phi_r \approx 1$.

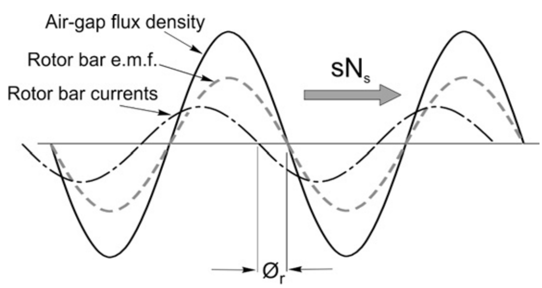

Figure 5.18 Pattern of air-gap flux density, induced e.m.f. and current in cage rotor bars at high values of slip. (These waveforms should be compared with the corresponding ones when the slip is small, see Figure 5.15.)

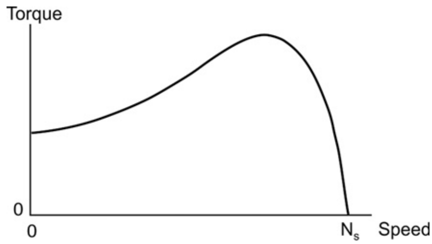

Figure 5.19 Typical complete torque–speed characteristic for motoring region of cage induction motor.

For most cage rotors, it turns out that as the slip increases the term cos ϕ_r reduces more quickly than the current (I_r) increases, so that at some slip between 0 and 1 the developed torque reaches a maximum value. This is illustrated in the typical torque–speed characteristic shown in Figure 5.19. The peak torque actually occurs at a slip at which the rotor inductive reactance is equal to the rotor resistance, so the motor designer can position the peak torque at any slip by varying the reactance to resistance ratio.

3.7 Generating – negative slip

When we explored the steady-state characteristics of the d.c. machine (see section 4 in Chapter 3) we saw that at speeds less than that at which it runs when unloaded the machine acts as a motor, converting electrical energy into mechanical energy. But if the speed is above the no-load speed (for example, when driven by a prime-mover), the machine generates and converts mechanical energy into electrical form.

The inherently bi-directional energy converting property of the d.c. machine seems to be widely recognized. But in the experience of the authors the fact that the induction machine behaves in the same way is far less well accepted, and indeed it is not uncommon to find users expressing profound scepticism at the thought that their 'motor' could possibly generate.

In fact, the induction machine behaves in essentially the same way as the d.c. machine, and if the rotor is driven by an external torque such that its speed is above the synchronous speed (i.e. the slip becomes negative), the electromagnetic torque reverses direction, and the power becomes negative, with energy fed back to the utility supply. It is important to note that, just as with the d.c. machine, this transition from motoring to generating takes place naturally, without intervention on our part.

When the speed is greater than synchronous, we can see from equation (5.6) that the slip is negative, and in this negative slip region the torque is also negative, the torque–speed curve broadly mirroring that in the motoring region, as shown in Figure 5.20. We will discuss this further in Chapter 6, but it is worth noting that for both, motoring and generating continuous operation will be confined to low values of slip, as indicated by the heavy line in Figure 5.20.

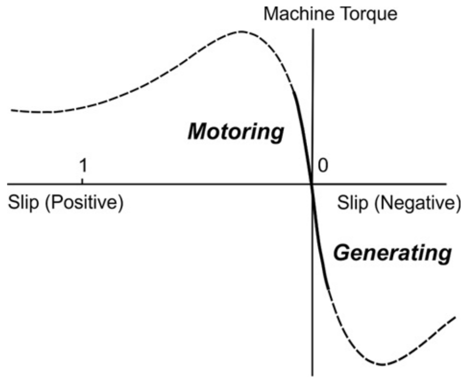

Figure 5.20 Typical torque–speed characteristic showing stable motoring and generating regions.

4. INFLUENCE OF ROTOR CURRENT ON FLUX

Up to now all our discussion has been based on the assumption that the rotating magnetic field remains constant, regardless of what happens on the rotor. We have seen how torque is developed, and that mechanical output power is produced. We have focused attention on the rotor, but the output power must be provided from the stator winding, so we must turn attention to the behavior of the whole motor, rather than just the rotor. Several questions spring to mind.

First, what happens to the rotating magnetic field when the motor is working? Won't the m.m.f. of the rotor currents cause it to change? Secondly, how does the stator know when to start supplying real power across the air-gap to allow the rotor to do useful mechanical work? And finally, how will the currents drawn by the stator vary as the slip is changed?

These are demanding questions, for which full treatment is beyond our scope. But we can deal with the essence of the matter without too much difficulty. Further illumination can be obtained from the equivalent circuit, which is discussed in Appendix 2.

4.1 Reduction of flux by rotor current

We should begin by recalling that we have already noted that when the rotor currents are negligible ($s = 0$), the e.m.f. which the rotating field induces in the stator winding is very nearly equal to the applied voltage. Under these conditions a reactive current (which we termed the magnetizing current) flows into the windings, to set up the rotating flux. Any slight tendency for the flux to fall is immediately detected by a corresponding slight reduction in e.m.f. which is reflected in a disproportionately large increase in magnetizing current, which thus opposes the tendency for the flux to fall.

Exactly the same feedback mechanism comes into play when the slip increases from zero, and rotor currents are induced. The rotor currents are at slip frequency, and they give rise to a rotor m.m.f. wave, which therefore rotates at slip speed (sN_s) relative to the rotor. But the rotor is rotating at a speed of $(1 - s)N_s$, so that when viewed from the stator, the rotor m.m.f. wave always rotates at synchronous speed, regardless of the speed of the rotor.

The rotor m.m.f. wave would, if unchecked, cause its own 'rotor flux wave', rotating at synchronous speed in the air-gap, in much the same way that the stator magnetizing current originally set up the flux wave. The rotor flux wave would oppose the original flux wave, causing the resultant flux wave to reduce.

However, as soon as the resultant flux begins to fall, the stator e.m.f. reduces, thereby admitting more current to the stator winding, and increasing its m.m.f. A very small drop in the e.m.f. induced in the stator is sufficient to cause a large increase in the current drawn from the supply because the e.m.f. E (see Figure 5.8) and the supply voltage V are both very large in comparison with the stator resistance volt-drop IR. The 'extra' stator m.m.f. produced by the large increase in stator current effectively 'cancels' the m.m.f. produced by the rotor currents, leaving the resultant m.m.f. (and hence the rotating flux wave) virtually unchanged.

There must be a small drop in the resultant m.m.f. (and flux) of course, to alert the stator to the presence of rotor currents. But because of the delicate balance between the applied voltage and the induced e.m.f. in the stator the change in flux with load is very small, at least over the normal operating speed-range, where the slip is small. In large motors, the drop in flux over the normal operating region is typically less than 1%, rising to perhaps 10% in a small motor.

The discussion above should have answered the question as to how the stator knows when to supply mechanical power across the air-gap. When a mechanical load is applied to the shaft, the rotor slows down, the slip increases, rotor currents are induced and their m.m.f. results in a modest (but vitally important) reduction in the air-gap flux wave. This in turn causes a reduction in the e.m.f. induced in the stator windings and therefore an increase in the stator current drawn from the supply. We can anticipate that this is a stable process (at least over the normal operating range) and that the speed will settle when the slip has increased sufficiently that the motor torque equals the load torque.

As far as our conclusions regarding torque are concerned, we see that our original assumption that the flux was constant is near enough correct when the slip is small. We will find it helpful and convenient to continue to treat the flux as constant (for given stator voltage and frequency) when we turn later to methods of controlling the normal running speed.

It has to be admitted, however, that at high values of slip (i.e. low rotor speeds), we cannot expect the main flux to remain constant, and in fact we would find in practice that when the motor was first switched on to the utility supply (50 or 60 Hz), with the rotor stationary, the main flux might typically be only half what it was when the motor was at full speed. This is because at high slips, the leakage fluxes assume a much greater importance than under normal low-slip conditions. The simple arguments we have advanced to predict torque would therefore need to be modified to take account of the reduction of main flux if we wanted to use them quantitatively at high slips. There is no need for us to do this explicitly, but it will be reflected in any subsequent curves portraying typical torque–speed curves for real motors. Such curves are of course used when selecting a motor to run directly from the utility supply, since they provide the easiest means of checking whether the starting and run-up torque is adequate for the job in hand. Fortunately, we will see in Chapter 7 that when the motor is fed from an inverter, we can avoid the undesirable effects of high-slip operation, and guarantee that the flux is at its optimum value at all times.

5. STATOR CURRENT–SPEED CHARACTERISTICS

To conclude this chapter we will look at how the stator current behaves, remembering that we are assuming that the machine is directly connected to a utility supply of fixed voltage and frequency. Under these conditions the maximum current likely to be demanded and the power factor at various loads are important matters that influence the running cost.

In the previous section, we argued that as the slip increased, and the rotor did more mechanical work, the stator current increased. Since the extra current is associated with the supply of real (i.e. mechanical output) power (as distinct from the original magnetizing current which was seen to be reactive), this additional ‘work’ component of current is more or less in phase with the supply voltage, as shown in the phasor diagrams (Figure 5.21).

The resultant stator current is the sum of the magnetizing current, which is present all the time, and the load component, which increases with the slip. We can see that as the load increases, the resultant stator current also increases, and moves more nearly into phase with the voltage. But because the magnetizing current is appreciable, the difference in magnitude between no-load and full-load currents may not be all that great. (This is in sharp contrast to the d.c. motor, where the no-load current in the armature is very small in comparison with the full-load current. Note, however, that in the d.c. motor, the excitation (flux) is provided by a separate field circuit, whereas in the induction motor the stator winding furnishes both the excitation and the work currents. If we consider the behavior of the work components of current only, both types of machine look very similar.)

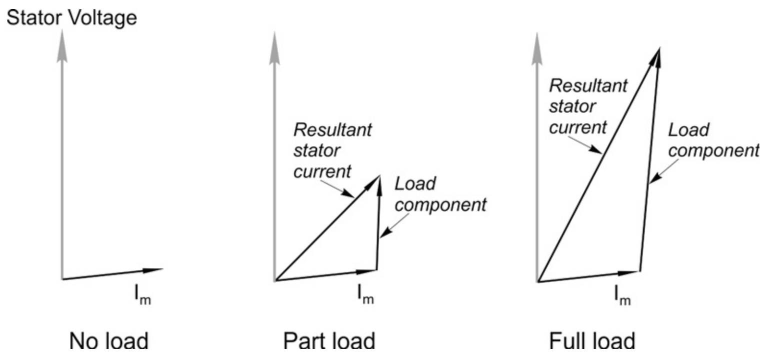

Figure 5.21 Phasor diagrams showing stator current at no-load, part-load and full-load. The resultant current in each case is the sum of the no-load (magnetizing) current and the load component.

The simple ideas behind Figure 5.21 are based on an approximation, so we cannot push them too far: they are fairly close to the truth for the normal operating region, but break down at higher slips, where the rotor and stator leakage reactances become significant. A typical current locus over the whole range of slips for a cage motor is shown in Figure 5.22. We note that the power factor is poor when the motor is lightly loaded, and becomes worse again at high slips, and also that the current at standstill (i.e. the 'starting' current) is perhaps five times the full-load value.

Very high currents when started direct-on-line are one of the worst features of the cage induction motor. They not only cause unwelcome volt-drops in the supply system, but also call for heavier switchgear than would be needed to cope with full-load conditions. Unfortunately, for reasons discussed earlier, the high starting

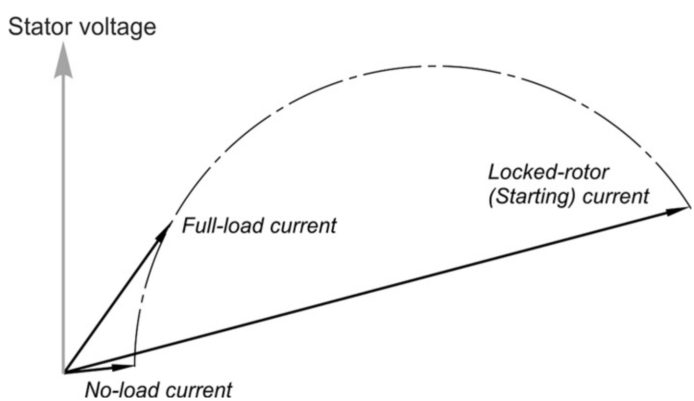

Figure 5.22 Phasor diagram showing the locus of stator current over the full range of speeds from no-load (full speed) down to the locked-rotor (starting) condition.

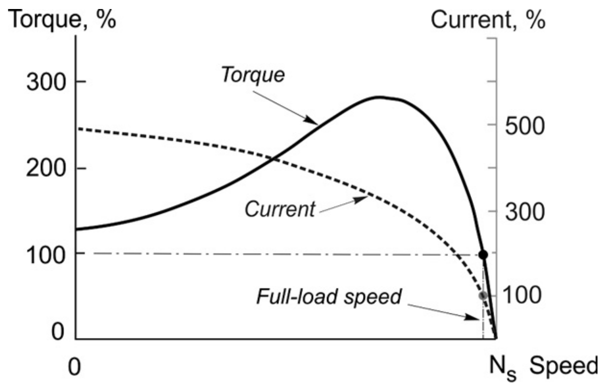

Figure 5.23 Typical torque–speed and current–speed curves for a cage induction motor. The torque and current axes are scaled so that 100% represents the continuously rated (full-load) value.

currents are not accompanied by high starting torques, as we can see from Figure 5.23, which shows current and torque as functions of slip for a general-purpose cage motor.

We note that the torque per ampere of current drawn from the mains is typically very low at start-up, and only reaches a respectable value in the normal operating region, i.e. when the slip is small. This matter is explored further in Chapter 6, and also in Appendix 2.